

Assisting the Design of Consumer Internet of Things Applications Using User Experience Patterns Incorporating Machine Learning

By

Rubem Barbosa-Hughes

School of Creative Arts

Submitted to the University of Hertfordshire in partial fulfilment of the requirements for
the award of the degree of Doctor of Philosophy.

Supervisors

Dr. Ian Willcock

Dr. Yi Sun

May 2025

Abstract

User experience (UX) design has become a pervasive function in the modern digital product lifecycle, spanning initial concept through to post-launch optimisation. Within the consumer Internet of Things (IoT), continuous streams of behavioural and contextual data present both opportunities and challenges for UX professionals willing to craft more intelligent and adaptive experiences for end-users. Machine learning (ML) is a typical mechanism leveraged to take advantage of such a data availability, yet many designers lack the literacy required to specify ML-driven interactions.

This doctoral research explores an approach, grounded in Alexandrian Pattern Theory, that enables UX designers – regardless of their level of ML literacy – to conceptualise feasible adaptive interactions for the consumer IoT. The work begins by formalising the Personal Spark methodology pattern, which combines engagement with field experts and the “What Would a Human Do?” thought experiment to uncover and refine recurrent design knowledge. Employing Personal Spark in one-to-one workshops with UX professionals (N = 23) uncovered five interconnected UX patterns – Preference Learner, Geographic Trigger, Smart Shortcuts, Stop Learning Preferences, and the proto-pattern Intelligent Access Sharing – collectively forming a novel, extensible UX pattern language for consumer IoT that integrates ML.

A mixed-methods, proof-of-concept evaluation enlisted UX designers (N = 5) who, after approximately ten to fifteen-minute briefing on two patterns (Preference Learner and Geographic Trigger), each produced a comprehensive consumer-IoT user-flow concept within an average of twenty minutes, all of which were deemed technically feasible in an independent IoT-expert assessment.

The principal contribution to UX scholarship and practice lies in operationalising Alexandrian Pattern Theory for data-driven contexts, therefore providing a shared language that aids the collaboration between UX design and allied functions whilst reducing the expertise needed for the integration of ML into UX solutions.

Acknowledgments

This doctoral research would not have been possible without the generous support and constant guidance of my supervisors, Dr Ian Wilcock and Dr Yi Sun.

From literally day one, Dr Ian Wilcock offered valuable advice, suggesting that I investigate Pattern Theory, as my research proposal appeared well aligned with it – indeed it was. His insightful guidance continued throughout, leading me on an incredible research journey that has ultimately culminated in this thesis. I will be forever grateful to Dr Yi Sun for her relentless encouragement and for sharing her expertise in machine-learning and many other incredibly helpful aspects. Her razor-sharp observations and feedback played a key role in enabling my progress with this work.

I am incredibly grateful to both of you for your tremendous patience. My life faced many ups and downs during this long journey, and I am sure yours did too; at no point, however, this affected the commitment you showed to me and to my research project. Thank you very much.

I am also grateful to Dr Barbara Brownie and Dr Silvio Carta for their help in the initial phase of my academic research life. Activities such as the annual symposia held by the School of Creative Arts were extremely useful, as were their practical advice and the support provided by the Doctoral College team.

Furthermore, I would like to thank my wife and daughter for their understanding and support during the many evenings and weekends that I had to dedicate to this work over the years. I look forward to celebrating this achievement together. I love you both very much.

To my mother, thank you for instilling in me the resilience needed to accomplish this endeavour. I wish you were here; you are, and always will be, missed.

To my auntie, thank you for sharing the little you had and adopting us – three orphans who had nothing but an immense baggage of loss. You went through a great deal to provide us with a place we could call home. Your personal sacrifices, made to keep ‘the siblings’ together, did not go unnoticed. Any significant achievements in our lives, such as this one, are only possible because of your generous gesture back in 1997. In the decades that followed, I met many individuals who have made a significant positive impact on the world and whom I admire, but my respect for them represents only a fraction of what I feel for you. I will be eternally grateful, and no number of words can describe how much.

Declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

Rubem Barbosa-Hughes
29th May 2025

Contents

Abstract	ii
Acknowledgments.....	iii
Declaration	iv
List of Tables.....	viii
List of Figures	ix
Abbreviations	xi
Glossary.....	xiv
1. Introduction	1
1.1. CONTEXT, AIMS AND MOTIVATIONS.....	1
1.2. RESEARCH QUESTIONS.....	4
1.3. SCOPE	5
1.4. CONTRIBUTION.....	6
1.4.1. <i>Publications</i>	7
1.5. THESIS STRUCTURE	9
2. Literature Review	12
2.1. INTRODUCTION.....	12
2.2. ABOUT PATTERNS.....	13
2.3. PATTERN MINING AND WRITING	14
2.4. PATTERNS IN COMPUTER SCIENCE	16
2.5. PATTERNS IN USER EXPERIENCE DESIGN.....	17
2.6. THE INTEGRATION OF MACHINE LEARNING IN UX DESIGN SOLUTIONS.....	21
2.7. THE CONSUMER INTERNET OF THINGS AS A CONTEXT FOR USER INTERACTION.....	26
2.8. THE RESEARCH GAP	29
3. Methodology.....	31
3.1. INTRODUCTION.....	31
3.1.1. <i>Paradigmatic stance and case for mixing methods</i>	32
3.2. LINKING THE RESEARCH QUESTIONS TO THE METHODOLOGY	33
3.3. ETHICAL CONSIDERATIONS	34
3.4. VALIDATION STUDY	34
3.5. SUMMARY OF THE METHODOLOGIES	34
4. The Experiments Using Machine Learning	36
4.1. INTRODUCTION.....	36
4.2. SIMPLIFIED USER IDENTIFICATION.....	37
4.2.1. <i>Participants and Data Sources</i>	38
4.2.2. <i>Apparatus and Materials</i>	39
4.2.3. <i>Procedure</i>	39
4.3. SIMPLIFIED USER IDENTIFICATION EXPERIMENT RESULTS	42
4.3.1. <i>Deliberation on the level of ML-information to include in the patterns</i>	44
4.4. AUTOMATED LABELLING OF MATERIAL DESIGN UI SCREENS.....	45
4.4.1. <i>Data Sources</i>	47
4.4.2. <i>Apparatus and Materials</i>	47
4.4.3. <i>Procedure</i>	48
4.4.4. <i>Results for The Automated Labelling Approach</i>	54
4.4.5. <i>Relevance to the Central Argument</i>	56
4.5. DISCUSSION AND IMPLICATIONS FOR THIS RESEARCH.....	56

5. Pattern Mining	58
5.1. INTRODUCTION.....	58
5.2. METHODOLOGY	59
5.2.1. <i>Ethical Considerations</i>	59
5.2.2. <i>Approach Design and Rationale</i>	60
5.2.3. <i>Participants</i>	61
5.2.4. <i>Apparatus and Materials</i>	65
5.2.5. <i>Procedure</i>	66
5.2.6. <i>Data Collection and Analysis</i>	69
5.2.7. <i>Limitations of the Methodology</i>	75
5.3. RECRUITMENT RESULTS	75
5.3.1. <i>Results for the Recruitment Survey</i>	77
5.4. THE PATTERN MINING WORKSHOPS RESULTS	81
5.4.1. <i>Shortlisted Good Pattern Candidates</i>	82
5.4.2. <i>Selected Patterns for Documentation</i>	83
5.4.3. <i>Contribution to the Research Project</i>	86
6. Formal Pattern Documentation	87
6.1. INTRODUCTION.....	87
6.1.1. <i>Pattern Authoring and Review Process</i>	89
6.2. METHODOLOGY PATTERNS.....	92
6.2.1. <i>The Personal Spark Pattern – Identification and Naming</i>	92
6.2.2. <i>Personal Spark – Formal Documentation</i>	93
6.3. AUTOMATION PATTERNS	97
6.3.1. <i>Geographic trigger pattern – overview</i>	98
6.3.2. <i>Geographic trigger – formal documentation</i>	98
6.3.3. <i>Smart shortcut pattern – overview</i>	101
6.3.4. <i>Smart shortcut – Formal Documentation</i>	101
6.4. PERSONALISATION PATTERNS.....	106
6.4.1. <i>Preference Learner Pattern – Overview</i>	107
6.4.2. <i>Preference Learner – Formal Documentation</i>	107
6.5. PRIVACY AND CONTROL PATTERNS	110
6.5.1. <i>Stop Learning Preferences Pattern – Overview</i>	111
6.5.2. <i>Stop Learning Preferences – Formal Documentation</i>	111
6.5.3. <i>Intelligent Access Sharing Proto-Pattern – Overview</i>	113
6.5.4. <i>Intelligent Access Sharing – Formal Documentation</i>	114
6.6. DISCUSSION	117
7. Validation Study	121
7.1. INTRODUCTION.....	121
7.2. METHODOLOGY	122
7.2.1. <i>Ethical Considerations</i>	122
7.2.2. <i>Research Design</i>	123
7.3. PARTICIPANTS	124
7.4. APPARATUS AND MATERIALS	126
7.5. PROCEDURE	132
7.5.1. <i>The Workshops</i>	132
7.5.2. <i>Assessment of Designers Output using Scenario-Based Evaluation</i>	134
7.5.3. <i>IoT Expert Assessment</i>	141
7.6. DATA COLLECTION	142
7.7. RESULTS	145
7.7.1. <i>Screening Survey Results</i>	145
7.7.2. <i>Quantitative Analysis of UX Flow Evaluations</i>	148
7.7.3. <i>Qualitative Analysis: Individual UX Flow Case Studies</i>	153
7.7.4. <i>The IoT Expert's feasibility Assessment</i>	166
7.8. LIMITATIONS	169
7.9. DISCUSSION	171

8. Conclusions	173
8.1. INTRODUCTION.....	173
8.2. SUMMARY OF THE RESEARCH PROJECT	174
8.3. CONTRIBUTIONS TO THEORY AND PRACTICE	174
8.3.1. <i>Theoretical Contributions:</i>	174
8.3.2. <i>Practical Contributions:</i>	175
8.4. LIMITATIONS AND FUTURE RESEARCH	176
8.4.1. <i>Limitations of the Machine Learning Experiments</i>	176
8.4.2. <i>Limitations of the Pattern Mining Method</i>	177
8.4.3. <i>Limitations of the Validation Study</i>	177
8.4.4. <i>Limitations of the Consumer IoT Expert Assessment</i>	178
8.4.5. <i>Recommendations for Future Research</i>	178
8.4.6. <i>Final Reflections on Limitations</i>	179
8.5. FINAL CONCLUDING REMARKS	180
9. References.....	183
Appendix A	192
Appendix B	199
Appendix C.....	206
Appendix D	207
Appendix E.....	208
Appendix F.....	209
Appendix G	214
Appendix H	215
Appendix I	217

List of Tables

Table 1 - F-Scores breakdown for each classifier based on the confusion matrices	43
Table 2 - Results of attempts to use a computer-vision approach (CNN) for the automatic identification of Consumer-IoT UI screens based on Material Design.....	54
Table 3 - Challenges and their solutions for effective pattern mining	61
Table 4 - Results of the A/B test assessing the creative assets for the participant recruitment posts.....	77
Table 5 - The 18 human-AI interaction design guidelines (Amershi et al., 2019)	89
Table 6 - Results of the validation study's screening survey.....	147
Table 7 - Summary of the ratings for the UX Flows created by the designers. Each coded between D1-D5 to maintain the anonymity of the participants.	148
Table 8 - Median values for each criterion.....	150
Table 9 - Spearman Correlation Matrix for the evaluation criteria	151
Table 10 - P-values indicating the statistical significance of the Spearman's rank correlation coefficients.	152

List of Figures

Figure 1 - A Directed Acyclic Graph representing an HCI pattern language for interactive exhibits by Jan Borchers (Borchers, 2000a, 2000b).	18
Figure 2 - Comparison of confusion Matrixes showing the results for different ML approaches. Approach 3 (KNN) achieved the best results.	44
Figure 3 - The approach for identifying good candidate Consumer IoT UI patterns to be enhanced with ML.	46
Figure 4 - Initial setup screen for IBM Cloud Annotations.	50
Figure 5 - Examples of Floating Action Button (FAB) selection for annotation.	50
Figure 6 - Due to the small amount of training data, the model did not perform as expected and led to false positives.	51
Figure 7 - Rico Android apps numbers across 27 categories.	52
Figure 8 - Close up on IBM Cloud Annotations stats bar showing the thumbnails of screens that neither had FABs or Fab-like UI elements (not-FAB). In total, 279 FABs were annotated from a sample of 2,158 UI screens.	53
Figure 9 - Examples of UI Screens with a floating action button (FAB) located on the bottom right hand corner.	53
Figure 10 - The participant recruitment process initially leveraged social media, networking events, and industry contacts. Higher-quality channels were introduced in later stages, informed by insights gained from early recruitment efforts. Some channels shown on the map (e.g., Twitter) were expected to perform less effectively, whereas others (e.g., LinkedIn and in-person contact) were anticipated to have a greater likelihood of attracting the right participants.	63
Figure 11 - The recruitment strategy successfully reached a high-quality audience. In the example above, Donald Norman questions my use of the term 'humans,' as he now appears to prefer 'people' instead.	64
Figure 12 - The original workshop structure document. Initially focused on UI patterns, it evolved to more effectively identify UX patterns. After the pilot study, participants were no longer required to use the MD UI Patterns, though it remained available as an optional resource.	68
Figure 13 - Example of an excel line organising the notes from interaction concept to initial pattern seed.	72
Figure 14 - Example of material produced by a designer, which was also considered in the analysis process.	74

Figure 15 - A/B test results for two creative executions for the participant recruitment ad. The Ad featuring the baby's hand had greater engagement. The performance and engagement metrics for each variant are shown in the lower section.	76
Figure 16 - Age Distribution of Respondents.	78
Figure 17 - Years of Experience in UX Design.	79
Figure 18 - Extract of an exchange between myself (in purple) and reviewer during the pattern writing review process. It shows typical peer review aspects covering arguments presented throughout the entire paper, not restricted solely to the patterns.	90
Figure 19 – An implementation of the Smart Shortcut pattern used in the eBay iOS app.	106
Figure 20 - The resulting directed acyclic graph (DAG) representing this novel pattern language.	119
Figure 21 - Alex Smarty, the proto persona used in the workshop. Photo by Caique Nascimento via unsplash.com.	126
Figure 22 - Miro board with the pattern documentation for the Preference Learner pattern.	127
Figure 23 - Miro board with the pattern documentation for the Geographic Trigger pattern.	128
Figure 24 - Scenario and task associated with the Preference Learner pattern.	129
Figure 25 - Scenario and task associated with the Geographic Trigger pattern.	130
Figure 26 - Overlaid ratings for all UX flows produced by the participants – none received the lowest rating in any criterion.	149
Figure 27 - D1's proposed user experience flow.	154
Figure 28 - D2's proposed user experience flow.	156
Figure 29 - D3's proposed user experience flow.	158
Figure 30 - D4's proposed user experience flow.	160
Figure 31 - D5's proposed user experience flow.	163

Abbreviations

A

AI

Artificial Intelligence .. 7, 8, 21, 23, 24, 25, 26, 28, 35, 47, 89, 93, 103, 108, 109, 112, 116, 131, 168, 174, 178, 179, 192

API

Application Programming Interface.....48

C

CNN

Convolutional Neural Network54

CNNs

Convolutional Neural Networks.....46

CSS

Cascading Style Sheets.....49

CSV

Comma-separated values38

CTA

Call to Action162

D

DAG

Directed Acyclic Graph118, 119

E

ECHO

Electronic Computing Home Operator8, 27, 193

EUD

End User Development6

EuroPLoP

European Conference on Pattern Languages of Program60, 85, 86, 90, 91, 92

F

FAB

Floating Action Button.....48, 49, 50, 51, 52, 53

G

GP1

Guiding Proposition 14, 35, 60, 67, 73, 82, 86, 117, 120, 180

GP2

Guiding Proposition 24, 35, 73, 121, 123, 181

GPS

Global Positioning System.....98, 100, 109

GPU

Graphics Processing Unit39

GUI

Graphical User Interface45, 46

H

HCI

Human-Computer Interaction14, 17, 18, 20, 23, 96, 123, 134, 203

HR

Heart Rate154

I

IDE	
Integrated Development Environment.....	39
IFTTT	
If This Then That	100
IML	
Interactive Machine Learning.....	23
IoT	
Internet of Things 1, 2, 3, 4, 5, 7, 8, 9, 11, 12, 13, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 38, 42, 46, 48, 52, 54, 55, 56, 58, 59, 60, 61, 62, 65, 66, 67, 69, 71, 72, 73, 77, 78, 79, 81, 82, 83, 84, 86, 92, 93, 96, 97, 98, 101, 102, 105, 106, 110, 111, 113, 114, 117, 120, 121, 122, 123, 124, 125, 141, 142, 145, 147, 156, 161, 165, 166, 168, 170, 174, 175, 176, 177, 178, 179, 180, 181, 182, 192, 193, 195, 196, 197, 206, 215, 216	

J

JSON	
JavaScript Object Notation	47, 49, 51, 52

K

KNN	
K-Nearest Neighbours.....	40, 41, 43, 44, 45, 57

L

LNCS	
Lecture Notes in Computer Science	91

M

MD	
Material Design.....	46, 52, 67
ML	
Machine Learning ... 1, 2, 3, 4, 5, 8, 9, 10, 11, 12, 13, 14, 17, 20, 21, 22, 23, 24, 26, 28, 29, 30, 31, 32, 33, 34, 35, 37, 39, 41, 42, 44, 45, 46, 54, 56, 57, 58, 59, 60, 61, 67, 69, 71, 73, 75, 77, 78, 79, 80, 81, 83, 85, 86, 91, 101, 115, 116, 117, 119, 120, 123, 125, 142, 147, 174, 175, 176, 179, 180, 202	
MVC	
Model View Controller.....	16

N

NLP	
Natural Language Processing	70

O

OCR	
Optical Character Recognition	23

P

PARC	
Palo Alto Research Laboratory.....	16
PINs	
Personal Identification Numbers	38
PLoP	
Pattern Languages of Programs	15, 16, 17, 91, 203

R

RBF	
Radial Basis Function.....	41, 45
RQ1	

Research Question 1	4, 33, 56, 60, 67, 80, 86, 180
RQ2	
Research Question 2	4, 33, 81, 121, 123, 180
S	
SMS	
Short Message Service	115
SP 1.1	
Sub-Proposition 1.1	45
SP 1.2	
Sub-Proposition 1.2	59
SP 2.1	
Sub-Proposition 2.1	37
Spyder	
Scientific Python Development Environment	39, 132
SVMs	
Support Vector Machines	40, 41
T	
TPLoP	
Transactions on Pattern Languages of Programming	91
U	
UCI	
University of California Irvine	38
UI	
User Interface... 10, 20, 32, 45, 46, 47, 48, 49, 51, 52, 53, 54, 55, 56, 58, 66, 67, 80, 81, 95, 99, 108, 109, 111, 135, 169, 201, 209	
UX	
User Experience.. 1, 2, 3, 4, 5, 7, 8, 9, 10, 11, 12, 13, 14, 15, 17, 18, 20, 21, 22, 23, 24, 27, 28, 29, 30, 31, 32, 34, 35, 42, 44, 45, 48, 54, 55, 56, 57, 59, 60, 61, 62, 65, 66, 67, 69, 70, 72, 73, 77, 78, 79, 80, 81, 82, 86, 91, 93, 94, 97, 105, 117, 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 131, 133, 134, 135, 136, 137, 139, 140, 141, 142, 143, 144, 145, 146, 148, 149, 150, 153, 154, 155, 156, 157, 158, 159, 161, 162, 164, 165, 166, 168, 169, 170, 174, 175, 177, 178, 180, 181, 192, 195, 196, 197, 203, 207, 215, 216	
W	
WWHD	
What Would a Human Do or What Would a Helpful Human Do 34, 61, 67, 70, 72, 92, 94, 95, 96, 174, 177, 195, 196	
X	
XAI	
Explainable AI	25

Glossary

Adaptive user experiences: A UX flow that integrates machine learning to adjust its behaviour, layout or content to match a user's context, often based on previous interactions and preferences. 1, 5, 82, 138.

Algorithmic opacity: The challenge for end users or UX practitioners to understand how an algorithm makes decisions. 2, 20.

Automation Patterns: UX patterns whose primary goal is to reduce users' cognitive load by having the system to carry out routine tasks automatically. Typically, the system is enabled to 'watch' for contextual cues (location, schedule, user behaviour, etc.) and then trigger the desired action without further user input – see Geographic Trigger and Smart Shortcuts as documented examples. 10, 91, 97.

Designer-centric pattern mining: The approach of discovering design patterns by collaborating directly with practising UX designers, capturing their lived experience to form a usable pattern language rather than relying on automated, data-driven techniques. 57, 59, 81, 86.

Methodology Patterns: Category of patterns that capture methods or approaches that UX professionals can embed in their design process. They operate one step 'above' typical UX patterns, serving as a repeatable approach to assist with specific tasks in their design practice (e.g., the **Personal Spark** pattern). 10, 85, 91, 92.

Personalisation patterns: UX patterns that tailor behaviour to the individual by learning from prior interactions and feedback. 10, 91, 106, 120.

Privacy and Control Patterns: Category of patterns that focus on ensuring that users retain agency over data collection specifically in the context of automated behaviours in ML-driven Consumer-IoT systems. They feature explicit controls (pause, delete, temporary sharing) so that personalisation and automation never happen at the expense of user autonomy – see Stop Learning Preferences and the proto-pattern Intelligent Access Sharing as examples. 10, 91, 110.

Proto-persona: A lightweight, assumption-based user archetype often used in early stages of digital product design drive empathy, before full research-based personas exist. 131, 132.

1. Introduction

1.1. Context, Aims and Motivations

The Consumer Internet of Things (IoT) has dramatically changed how people interact with technology (Swan, 2012) by embedding connectivity and intelligence into everyday devices – ranging from smart thermostats to wearable fitness trackers. These devices capture and analyse large volumes of data, creating new possibilities for personalisation and automation of user experiences (Barbosa-Hughes, 2020a; Perera et al., 2014). Machine Learning (ML) supports these advances by enabling systems to learn from user interactions and adapt to changing (Jordan and Mitchell, 2015).

That combination of technological resources and context of use suggested an ideal opportunity for enhancing the user experience (UX) of smart home devices, making it possible to adapt to user's needs in a manner we have never been able before. Nonetheless, by late 2016, when this research proposal was first formulated, there were not many examples of adaptive user experiences leveraging ML to reduce

cognitive load from user interactions. In particular in the consumer-level Internet of Things context, ML was not being widely used to enable personalisation of user interactions.

My professional engagement with the UX design community revealed a general disconnect between ML's theoretical potential and practical design solutions. In these conversations, fellow professionals often shared that ML was rarely considered in their design processes, and it transpired that there was a pervasive lack of understanding regarding the nature of ML and its potential, when wisely applied, to enhance user experiences – this view is also supported by contemporary research findings (Dove et al., 2017; Yang, 2017).

Challenges in Integrating ML into UX Design

The potential for ML to improve user experiences is not without significant hurdles, particularly for practitioners accustomed to more predictable design processes. UX designers often lack the technical literacy needed to effectively integrate ML capabilities into their design frameworks (Dove et al., 2017). This skills gap may hinder the development of Consumer IoT design solutions that leverage the potential of ML to improve user interactions. Further, the incorporation of ML into UX design presents challenges that extend beyond mere technical proficiency.

Traditional UX design paradigms, which are often based on static user flows and predictable interactions, may prove inadequate when faced with the complexities introduced by ML technologies. The inherent algorithmic opacity of ML systems, in some cases characterised by their 'black box' decision-making processes, presents a significant challenge for designers (Amershi et al., 2019). Moreover, the unpredictability and adaptive nature of ML algorithms may push the design cognition processes away from conventional design practices, forcing designers to account for systems that learn and evolve over time, and which potentially can change their behaviour with its users in unpredictable ways.

The Need for a User-Centred Approach

Considering the complexities inherent in ML-driven applications, it becomes essential to reassert the role of user-centred methodologies. That is why a focus on user needs is not only advisable (from a commercial perspective), but crucial (from a user experience standpoint), when integrating ML into Consumer IoT.

Nonetheless, as it is often the case with emerging technologies, there is a significant risk that technological novelty, rather than genuine user needs, may drive the adoption of ML in design solutions. This technology-centred approach could lead to the implementation of ML features simply because they are technically impressive,

without a consideration of their relevance or impact to users. This phenomenon seems very similar to what happened during the early days of the web, when graphic designers applied print design principles to that new dynamic and interactive medium that needed its own emerging set of rules¹.

Motivation for a Pattern-Based Approach

To mitigate the aforementioned risks, it is essential for UX designers to adopt an approach that prioritises usability and user needs over technological novelty. As Norman (2013) emphasises, good design is centred around the user's needs and capabilities, ensuring that technology serves to enable better UX rather than complicate it. Nielsen (1993) also advocated for usability engineering principles that focus on making systems 'easy to learn and efficient to use' – a concept that remains valid after all these years, particularly when integrating complex technologies like ML.

Consequently, there is a pressing need for accessible tools and frameworks that enable UX designers to integrate ML into Consumer IoT applications without requiring deep ML expertise. Such resources would enable designers to harness ML capabilities effectively while maintaining a focus on user needs and usability (Dove et al., 2017), in a similar way that UX professionals today design web experiences without necessarily knowing the technicalities of JavaScript frameworks. By accounting for the complexities associated with ML while providing structured guidance, such a resource might provide the missing link between technological potential and practical design application.

Design patterns, as conceptualised by Alexander et al. (1977), offer a means to capture and communicate proven solutions to recurring design problems within a context. These patterns emerge from the collective knowledge and experience of practitioners and can exist informally within a community of practice. Formal documentation of these patterns enables their wider accessibility and usability, providing designers with structured guidance (Dearden and Finlay, 2006). UX designers can utilise design patterns to bridge the gap between ML technologies and user-centred design, ensuring that ML is applied in ways that improve the UX.

Through this research project, I investigate a pattern approach to address that pressing need, and a mechanism to enable UX designers who lack technical expertise in ML to harness it to design more intuitive and intelligent Consumer IoT applications.

¹ John Allsopp's influential 2000 essay, 'A Dao of Web Design' highlighted the differences of designing for the web because the fixed conventions of print were no longer suitable for the web's more fluid medium. <https://alistapart.com/article/dao/>

1.2. Research Questions

Following Yin's (2018) framework for qualitative inquiry design, in which propositions serve to direct and bound investigation rather than specify prediction for deductive testing, the central framing statements of this research project are called Guiding Propositions. Each proposition orients the inquiry towards a specific direction, focusing and guiding what should be examined. The aims and motivations described in the previous section led to me to converge on the following guiding propositions:

Guiding Proposition 1 (GP1):

A pattern language that encapsulates machine-learning-driven interaction guidance for Consumer IoT can be identified, thereby helping designers conceptualise adaptive and user-centred features.

Guiding Proposition 2 (GP2)

When employed by designers with no deep expertise in machine learning, this pattern language may support the creation of feasible Consumer IoT interaction concepts.

These propositions are explored through two research questions that define the scope and direction of the empirical work. These are:

RQ1. Can a pattern language that encapsulates the use of machine learning to support user interactions with Consumer IoT be identified?

This question seeks to determine whether it is possible to identify a collection of related UX design patterns that leverage ML to enhance user interactions within the Consumer IoT context. To answer this question, I followed the following steps:

1. Identification of UX patterns emerging from the community of practice that integrate ML and addressed UX problems in the Consumer IoT context but have not yet been formally documented.
2. Formal documentation of these patterns to make them available for UX practitioners and academic community.

RQ2. If identified, can such a pattern language support the effective design of user interaction concepts for Consumer IoT scenarios by designers with no machine learning expertise?

This question examines the practicality and usability of the documented pattern language for UX designers without specialised knowledge in ML. I addressed this question by using the following approach:

1. I assessed whether designers lacking deep ML expertise can successfully utilise these novel patterns in a typical design process and evaluated the usability these patterns.
2. I facilitated an independent technical feasibility assessment of the design solutions (derived from the use of the same patterns), which served as an evaluation of the effectiveness of these patterns in enabling UX designers to design technically feasible Consumer IoT interaction flows that incorporate ML functionality.

1.3. Scope

This research sits on the intersection of UX design, pattern theory, ML, and the Consumer IoT. It specifically addresses the challenges faced by UX designers in integrating ML into Consumer IoT applications without possessing deep technical understanding of ML algorithms and systems.

The scope includes:

- **Consumer level IoT applications:** The design process considered pertains to the smartphone controller apps that are used by consumers to interact with smart home devices.
- **Identification and Formal Documentation of UX Design Patterns:** Uncovering existing but undocumented good quality patterns that incorporate ML within the Consumer IoT context and formally documenting them for use by UX designers.
- **Evaluation of the Patterns' Potential Efficacy as a Design Tool:** If identified and documented, these patterns would be examined for their capacity to enable the feasible design of adaptive user experiences within the Consumer IoT context, drawing on designers' existing UX design expertise.

The scope excludes:

- **Technical implementation of ML Algorithms:** I focused on the UX design perspective rather than the technical implementation or optimisation of ML technologies.
- **Industrial IoT Applications:** I concentrate solely on consumer-facing products, this research project does not address industrial or commercial IoT applications.
- **Technical implementation of Consumer IoT technology:** Although the UX design level solutions take technical limitations into consideration, this work does not cover aspects of the Consumer IoT hardware or connectivity methods, nor go into any detail of troubleshooting issues in these areas.

- **End User Development Paradigms:** While acknowledging the relevance of End User Development (EUD) in empowering users, this research does not focus on EUD approaches.

1.4. Contribution

This research contributes to the fields of UX design and Pattern Theory in several ways:

Formal Documentation of a Set of UX Design Patterns

A key contribution of my research is the formal documentation of User Experience (UX) design patterns, which was the outcome of applying theoretical pattern frameworks to practical UX design needs. This formal documentation moves these patterns from the realm of 'abstract knowledge' that existed within the design community towards an accessible resource available to industry practitioners.

I systematically captured a set of recurring design solutions called UX patterns or UX design patterns, explicitly defining their contexts, forces, and implementation considerations. This resulted in a formal documentation that supports the effective application of these patterns in real-world settings, therefore establishing a reliable foundation for informed design decisions. Furthermore, the documented patterns can aid interdisciplinary collaboration by enabling UX practitioners to propose sound and technically viable concepts, which enhance dialogue with technical teams regarding design requirements and constraints. Consequently, this contributes to the ongoing advancement of best practices within contemporary UX design.

Identification of a Novel ML-Enhanced UX Pattern Language

As I identified patterns that integrated ML to address Consumer IoT UX challenges with adaptive solutions, a cohesive relationship between these patterns also began to emerge. This interrelationship aligns with the principles from Alexandrian pattern theory: a pattern language is not solely a collection of isolated patterns, but an interconnected network structured through meaningful relationships. Therefore, due to their interconnection, these patterns form a cohesive pattern language, which, although unlikely to be complete, in this first iteration establishes a sound foundation for further exploration and expansion by the pattern writing community.

Development of a Novel Approach for Pattern Mining of Adaptive UX Patterns

I further contribute methodologically by showing how UX professionals – who may have minimal machine learning expertise – can effectively participate in the

identification of adaptive design opportunities. I devised and tested a workshop format wherein participants employ a “What Would a Human Do” (WWHD) thought experiment to isolate recurring user interaction problems. By reframing digital tasks in terms of natural human interaction, I exposed steps in those interactions where ML-driven adaptivity could be used to improve the user experience.

When real-world implementations of these resulting concepts are identified and shown to employ good interaction design and Human-AI interaction principles, then they become suitable candidates for formal documentation. This process ensures that the resulting patterns are based on authentic, human-centric experiences, rather than on abstract technical considerations alone. I have documented this novel pattern-mining approach as a distinct methodology pattern titled Personal Spark.

Advancing Academic Discourse

This research expands the dialogue on the convergence of UX design and ML, particularly within the Consumer IoT domain, and contributes to the theoretical foundations of pattern languages in this context. While the use of Alexandrian Pattern Theory in digital domains is well documented, this research demonstrates its applicability in a novel context: the integration of machine learning into user experiences for the Consumer IoT. I have not sought to extend the fundamental theoretical framework itself, i.e. the Pattern Theory, but rather to operationalise it in a way that yields a dynamic pattern language.

1.4.1. Publications

Some of the work presented in this thesis has resulted in the following peer-reviewed publications and conference submissions:

A Pattern Approach for Identification of Opportunities for Personalisation and Automation of User Interactions for the IoT (Barbosa-Hughes, 2019):

In this paper I introduce the ‘Personal Spark’ pattern, aimed at helping UX designers identify opportunities for personalisation and automation within user interactions for Consumer IoT. The pattern provides a structure for systematically isolate and use relevant aspects of human-human interactions to inform the design of adaptive user interfaces. I tested this pattern with 23 UX designers, who found it a useful tool for the early conceptual phase of interaction design in Consumer IoT. Moreover, the resulting interaction concepts created by those UX designers were used to uncover the other patterns documented through this research work, as detailed in chapters 5 and 6.

Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT (Barbosa-Hughes, 2020a):

In this paper I introduce three patterns that target UX design challenges within Consumer IoT applications, leveraging ML and geolocation. The patterns – Geographic Trigger, Preference Learner, and Stop Learning Preferences – address the problems of automation, personalisation, and privacy/control, respectively. These patterns encapsulate guidance that assist the UX Design of connected devices by using well-established interaction principles and emerging human-AI interaction guidelines. The methodology for identifying these patterns involved workshops using the Personal Spark pattern described above, with experienced UX designers.

Professional Doctorates in Design and the Use of Computational Tools” for the 7th International Conference on Professional & Practice-Based Doctorates. In this work, my contribution focused on the section covering my use of the ‘Pattern Writing Sheet’ – see Appendix B.

I presented a case study involving the application of the Pattern Writing Sheet in the early stages of my research. The Pattern Writing Sheet, as introduced by Iba (2014), serves as a structured tool to assist in documenting design patterns by encouraging authors to deeply explore the problem, the solution, and the resulting consequences of the pattern.

Portions of this research were presented at peer-reviewed conferences but subsequently not accepted or included in the proceedings, however they significantly informed the research documented here. Some of this work is available at open platforms such as ResearchGate.

I presented relevant perspectives at the 17th Annual International Conference of the Architectural Humanities Research Association through a paper titled “The Role of the Consumer IoT Experience in Extending the Notion of Home”² (Rubem Barbosa-Hughes, 2020) – see Appendix A.

In this paper I examine the evolving role of Consumer IoT in extending the notion of the modern home beyond traditional boundaries. I traced the historical development of smart homes, starting with early automation concepts like the 1960’s ECHO IV, to current applications involving single smart devices and interconnected systems. I explored how UX design for smart devices influences

² Available at ResearchGate

https://www.researchgate.net/publication/344754834_The_Role_of_the_Consumer_IoT_Experience_in_Extending_the_Notion_of_Home

people's experience of home and highlighted the importance of personalisation, efficient automation, and human-centric design in enhancing user interaction, ensuring that smart home technologies contribute meaningfully to users' quality of life and emotional connection to their homes. The paper also points out the challenges faced by UX designers in integrating smart technologies smoothly and effectively, emphasising the need for user-centricity.

Additionally, I had the article titled '**A step towards a much richer and personalised human experience with the Internet of Things**'³ published in the 2019 edition of the **Vision & Voice** virtual magazine, where I briefly explained some of the benefits of my research.

1.5. Thesis Structure

This thesis is structured into eight chapters, each contributing to a systematic exploration of UX patterns incorporating machine learning (ML) for the Consumer Internet of Things (IoT) context. The structure follows a logical progression, beginning with research motivations and explanations on foundational theories, moving through methodological choices and empirical investigations, and culminating in the synthesis of findings and contributions.

Each chapter plays a unique role in addressing the research questions.

Following the Introduction (Chapter 1), where I establish the context, motivations, research questions, and contributions, the thesis proceeds with a structured exploration of mechanisms for identification, documentation and validation of UX patterns incorporating Machine Learning (ML) to aid the design of Consumer Internet of Things (IoT) interactions.

As I critically review previous research in the Literature Review (Chapter 2), I chart the development of design patterns from architecture to Human-Computer Interaction (HCI), emphasising their use in UX design. The difficulties of designing for consumer IoT environments, the incorporation of machine learning into UX design, and pattern mining and documentation techniques are among topics I cover. The chapter concludes by identifying a research gap, namely the lack of a pattern language that aids UX designers in integrating machine learning into consumer Internet of Things applications.

In the Methodology (Chapter 3), I outline the methodological approach, structured into four key phases:

³ Vision & Voice 2019 edition

https://issuu.com/uniofherts/docs/vision_and_voice_2019_the_next_generation_of_resea

1. **Machine Learning Experiments**, where I explore ML's approaches to aid identification of UI patterns. I also use gait recognition as a use case to understand which level of ML implementation information would be relevant to added to the documented patterns.
2. **Pattern Mining**, where I describe the structured workshops and expert input I used to identify the novel UX design patterns;
3. **Pattern Documentation**, where I cover the process for formally documenting the UX patterns.
4. **Validation Study**, where I evaluate the usability and effectiveness of these patterns. I also detail the ethical considerations, participant recruitment, and data collection strategies.

In The Experiments Using Machine Learning (Chapter 4), I present results of two exploratory studies using ML. I discuss the performance of The Automated Labelling experiment and the comparison of confusion matrices containing results for different ML approaches used for recognising users through gait data. I discuss how this realistic ML workflow enabled me to inform my decision regarding the adequate levels ML information to be included in the formal pattern documentation, especially when considering the primary target audience – UX professionals.

In Pattern Mining Results (Chapter 5), I detail how UX patterns were extracted from expert workshops, including outcomes of the participant recruitment strategy and data analysis. I describe the criteria for selecting high-quality patterns, followed by a discussion on the shortlisted early pattern concepts that were further developed and formally documented.

In Formal Pattern Documentation (Chapter 6), I present the structured documentation of the newly identified UX patterns, categorised into four categories:

- Methodology Patterns, guiding designers in conceptualising ML-driven UX solutions.
- Automation Patterns, supporting interaction automation through ML.
- Personalisation Patterns, enabling adaptive, user-centred experiences.
- Privacy and Control Patterns, addressing trust and transparency in ML-powered UX interactions.

Each pattern is documented using a standard, formal pattern authoring structure that includes problem definition, context, forces, solution, and consequences.

In Validation Study Outcomes (Chapter 7), I evaluate the usability and effectiveness of the documented UX patterns through a structured validation study. I involve UX

designers and an IoT expert, using both statistical and qualitative methods to assess the impact of these patterns on typical UX design processes.

Finally, in Conclusions (Chapter 8), I synthesise the key findings of the research, discussing its theoretical and practical contributions. I reflect on and highlight limitations of this research project and suggest future research directions. The chapter concludes with a final reflection on the significance of the research in bridging the gap between ML capabilities and the UX design practice.

I conclude the thesis with a comprehensive list of references, followed by the appendices, which include supplementary materials quoted in parts of this document.

2. Literature Review

2.1. Introduction

This chapter critically examines the theoretical and practical contexts underpinning this research, with a focus on design patterns, user experience (UX) methodologies, and the integration of machine learning (ML). The aim is twofold. First, I explore how the concept of patterns has evolved from architecture to software engineering, eventually influencing UX design practices. Second, I highlight how recent developments in ML have yet to be sufficiently captured by pattern-based approaches – particularly for designers who lack deep technical expertise. By placing my investigation within the Consumer Internet of Things (IoT) domain, I highlight a critical gap: despite the rich history of pattern research, there remains no widely recognised pattern language dedicated to supporting UX designers to incorporate ML functionalities seamlessly into consumer IoT experiences. This gap is especially crucial given the rapid growth of connected devices in everyday life and the demand for personalised, intelligent interactions. Patterns

tailored for ML-driven consumer IoT can serve as a structured way for UX practitioners to address these design challenges without requiring a deep background in data science.

2.2. About Patterns

The concept of patterns emerged from the field of architecture when Christopher Alexander proposed an approach to encapsulate the design expertise found in solutions for recurring problems. Confronted with the challenges of creating spaces that resonate with human needs and aspirations, Alexander proposed an innovative approach at the time that sought to distil the essence of successful design solutions for recurrent architectural problems.

Alexander and his colleagues produced a series of books intended to disseminate this new approach to architecture, ultimately leading to a trilogy of seminal texts: Volume 3 'The Oregon Experiment' (Alexander et al., 1975), volume 2 'A Pattern Language: Towns, Buildings, Construction' (Alexander et al., 1977), and finally volume 1 'The Timeless Way of Building' (Alexander, 1979). Combined, these works marked the beginning of what became known as Alexander's Second Theory of Architecture (Dawes and Ostwald, 2017), which would later branch into the wider field of Pattern Theory (Leitner, 2015).

In The Oregon Experiment, Alexander and his colleagues cover ambitious project to apply the ethos and principles of their newly proposed approach in a real-world setting – the campus planning of the University of Oregon. They advocated for a participatory design process, very similar to co-creation methods adopted today in digital product design, where users of the space participated and had a say on its creation, thus ensuring that the resulting environment genuinely reflected their needs (Alexander et al., 1975, pp. 9-15). This experiment was critical in demonstrating how patterns could serve as a common language between architects and the community of users and therefore facilitate design.

Central to Alexander's philosophy was the idea that certain design problems emerge across different contexts and that the solutions to these problems could be encapsulated into patterns, which are a structured way of capturing the essence of a problem and its solution. Here is Alexander et al. on patterns:

"Each pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice" (Alexander et al., 1977)

This definition places patterns not merely as design templates to be replicated verbatim but as adaptable frameworks that can guide the design process, allowing for many different executions while remaining aligned to the underlying solution.

Alexander's vision was inherently democratic. By codifying design expertise into an accessible format (a pattern language), he hoped to enable individuals to actively participate in the creation of their own spaces. This approach challenged the traditional top-down paradigm of architectural design, such as the typical concept of a 'Master plan', advocating instead for a more inclusive and human-centred methodology which he called 'The Principle of Participation' (Alexander, 1975, pp. 38-66).

The implications of Alexander's work became noticed beyond the field of architecture. The idea of encapsulating expertise along with the efficiency of reusable design resonated with practitioners in software engineering and human-computer interaction (HCI), domains grappling with their own versions of complex, recurrent problems.

Building on Alexander's vision of participatory design and adaptable solutions, this research extends the pattern paradigm to contemporary user interaction challenges that involve increasingly intelligent systems. In essence, the democratic ethos of Alexander's approach aligns with the modern need to make sophisticated technological capabilities – like those of machine learning – accessible to designers without demanding specialised expertise. The next section explores how pattern mining and writing processes formalise these insights, setting the stage for integrating ML-driven solutions into UX frameworks.

2.3. Pattern Mining and Writing

The process of formally documenting patterns can be seen as divided in two phases: Mining and Writing.

Mining is a term used in the pattern community which describes the process of discovery of patterns in a specific field, and it is essential for documentation of patterns. There are many approaches for pattern mining (Akado et al., 2015; Hentrich et al., 2015; Iba and Yoder, 2014), however they all involve interacting with domain experts. In this research project, such experts are the UX professionals. This approach may also be interpreted as a variation of a Delphi Study (Chuenjitwongsa, 2017), as part of its aim is to gather consensus from a group of experts around a particular theme. However, the pattern mining process is further specialised as it will focus on extracting key essential 'sections' needed for a formal pattern documentation such as context or forces for example.

In a typical pattern mining process, expert knowledge is sought to understand how they would address a particular problem in their field of expertise, then further clarification is collected to cover all that is necessary to document a pattern. As mentioned earlier, this interrogation process can take place in many forms such as interviews (Iba and Yoder, 2014), workshops (Akado et al., 2015) or using Grounded Theory Techniques for example (Hentrich et al., 2015). The objective is to “try to identify a solid and reliable invariant, which relates context, problem, and solution, in an unchanging way.” (Alexander, 1979) (pp257).

In UX pattern mining, authors look for consistent user interface or interaction problems (e.g. how to guide onboarding or manage complex settings), in a similar manner software architects look for common coding challenges. Through interviews, observations, and expert feedback, UX pattern writers uncover underlying forces and constraints, then document a pattern’s context, problem, and solution – mirroring closely how software design patterns are documented.

Patterns, as problem solving abstractions extracted from successful designs, are the product of explicit design decisions made, in this case, by UX Designers. In Software engineering, this application of ‘solutions to specific recurring problems’ (Gamma et al., 1995) , translates into concerns such as code architecture or modularity for example, which helps developers to reuse working solutions across different projects.

‘[...] software patterns to be newly discovered are successful mental design constructs, which have not yet been discovered as such, but are rather represented as unexplored and unconscious knowledge of expert software designers that is hidden in the software implementations.’ (Hentrich et al., 2015)

This process – identification of recurrent problems and matching abstracted and structured solutions – has been naturally adapted to UX design challenges.

When it comes to writing, the most widely adopted approach proposes the following sequence: First identify the solution, then the problem, and finally the context (Iba, 2014; Wellhausen and Fießer, 2012). Authors may use specialised resources such as the Pattern Writing Sheet (Iba, 2014) to aid the organisation of the information in the unique way that ultimately result in a pattern. This clearly suggests that Pattern Mining and Pattern Writing are intimately linked.

Once documented, authors have the option to subject their patterns to the rigorous pattern community review at Pattern Languages of Programs PLoP

conferences. Submissions to the PLoP conferences go through a uniquely thorough process which ensures that only the highest standard of pattern authorship gets published in the final proceedings. Four of the patterns documented in this research project has already gone through the PLoP review.

2.4. Patterns in Computer Science

Myers (1998) emphasised that “more than anything else, improvements to interfaces have triggered this explosive growth” of computing. He also stated that as computers become faster, “more of the processing power is being devoted to the user interface” highlighting the growing importance of user-interface considerations in software design between 1960s and late 1970s. In this context, a software architecture pattern called Model View Controller (MVC) emerged. Trygve Reenskaug first implemented it at Xerox Palo Alto Research Laboratory (PARC) in 1978/79 (Reenskaug, 1979). This pattern logically separated the user interaction components (View) from the business logic (Model). The controller element acted as a bridge between the View and the Model.

The MVC pattern was a successful example of how patterns could encapsulate complex design principles into more accessible forms (Buschmann et al., 1996). It provided a structure for developing software in a way that addressed challenges of maintainability and modularity, supporting a more collaborative development effort. By abstracting the solution to a common problem, MVC allowed developers to leverage established best practices, reducing the cognitive load associated with designing complex systems.

Later, the notion of patterns found fertile ground, particularly with the advent of object-oriented programming. The seminal work ‘Design Patterns: Elements of Reusable Object-Oriented Software’ by Gamma et al. (1995) , often referred to as the "Gang of Four", was key in increasing the popularity of the pattern approach in software engineering. In the Object-Oriented paradigm, reusability as well as the notion of inheritance and abstraction were very important, making it an ideal use case for patterns.

In addition, patterns allowed the documentation of expertise so that experienced designers and novices alike could reuse the same successful designs, despite the many years of experience that separated them (Gamma et al., 1995) pp. 1-5. At the time, Gamma et al. also discussed patterns for ‘Embellishment’ of the user interface, however, these were targeted at software engineers rather than human-computer interaction designers.

Some researchers claim that the application of the patterns themselves may introduce further complexity to the software engineering process as it requires deeper understanding of the limitations and consequences of using them (Prechelt et al., 2001; Zhang and Budgen, 2012). Nonetheless, the use of design pattern in software engineering became a standard practice in the industry and it is constantly evolving. Perhaps the most remarkable aspect of the pattern approach is the fact that the solutions and knowledge embodied in them are from practice, rather than theory (Dix et al., 2003, pp. 285).

As modern software development embraces agile methods and rapidly changing technologies, pattern adoption must continuously adapt, balancing the clarity of reusable solutions with the intricacies of novel frameworks such as ML-driven architectures.

2.5. Patterns in User Experience Design

It is the use of patterns in HCI and its related domains that is of interest in this research project. Following the adoption of patterns in software engineering in the late 80s and early 90s, the concept started to find its way into HCI. In the first Pattern Languages of Program (PLoP) conference in 1995, Riehle and Züllighoven presented a pattern language containing four patterns that featured the use of metaphor in user interface design (Riehle and Züllighoven, 1995), in addition to seven other software engineering patterns associated with the implementation of the interfaces related solutions.

During the 1990's and early 2000's, the PLoP conferences became the common place for researchers in HCI to present their pattern collections and subject their work to the adequate level of rigour required for authoring patterns (Erickson, 2000; Tidwell, 1998).

This research adopts the views of Sharp et al. (2019), that the terms interaction design and UX design can be used interchangeably. Despite the term UX Design having been coined by Donald Norman in the early 90's⁴ and it becoming more popular towards the late 90's, terms such as human(or user)-computer interaction (Card et al., 1980) and interaction design had been used in the literature much earlier than that to discuss what later would become part of the UX concerns.

As such, it could be said that the UX Design domain did not emerge in the moment it was given a name. Around the same time, Jenifer Tidwell published her widely adopted Common Ground pattern language (Tidwell, 1998). These patterns,

⁴ <https://www.youtube.com/watch?v=9BdtGjoIN4E>

however, did not focus on user interaction needs, but on artefact design instead. Two years later, Welie and Tr  tteberg published a pattern collection which focused primarily on user interaction problems (Welie and Tr  tteberg, 2000). These patterns were also aligned with Norman’s interaction principles (Norman, 1988), setting a precedent for this practice in UX pattern authoring.

Other significant contributions appeared at this time, namely, Borchers’s A Pattern Approach to Interaction Design (Borchers, 2001) and The Design of Sites: Patterns, Principles, and Processes for Crafting a Customer-centered Web Experience (van Duyne et al., 2003). Shortly later Jenifer Tidwell’s Designing Interfaces: Patterns for Effective Interaction Design (Tidwell, 2005). These works not only contributed to the theoretical basis of pattern utilisation in design but also provided practical frameworks and pattern languages that heavily influenced the contemporary UX design practices.

Borchers proposed a Formal Model of Pattern Languages (Borchers, 2001, pp. 52), stating that ‘A pattern language is a directed acyclic graph $PL(\wp, \mathfrak{R})$ with nodes $\wp = \{P1,...,Pn\}$ and edges $\mathfrak{R} = \{R1,...,Rm\}$ ’. In this notation, ‘R’ represents the relationships or references between patterns in the pattern language. Moreover, to understand Borchers’s model one must consider a Directed Graph as being vertices (patterns) connected by edges (relationships), where each edge has a direction from one vertex to another. The Acyclic aspect means that there is no way to start at a vertex (or specific pattern) and follow a sequence of directed edges that eventually loops back to itself. This is a key property of this representation that facilitates a logical sequence and hierarchical organisation of patterns.




Figure 1 - A Directed Acyclic Graph representing an HCI pattern language for interactive exhibits by Jan Borchers (Borchers, 2000a, 2000b).

Figure 1 shows the mind map style diagram representing a pattern language, made by Borchers (2000a) where the author uses solid lines to illustrate the relationships between patterns. In Chapter 6 I offer similar representation for the pattern language I document in this research project.

In addition to the formal proposition, Borchers argued the crucial benefit that patterns could serve as a 'lingua franca' between interdisciplinary teams, facilitating communication and collaboration among designers, developers, and stakeholders.

Beyond offering a convenient way for documentation of recurring solutions, Alexandrian Pattern Theory is also fundamentally explanatory. A pattern is a structured account of the problem, context, forces, solution and consequences alignment that repeats across multiple situations. Patterns explicitly articulate the tensions they resolve, making visible why a particular configuration of interaction elements tends to succeed or fail. When patterns are linked into a pattern language, as in Borchers' example, the relationships between them further clarify dependencies and preconditions – for instance, which patterns must be in place to safeguard user autonomy when adaptive automation is introduced. In this thesis, pattern theory is therefore used not solely as a static taxonomy of ML-enhanced UX solutions for Consumer IoT, but also as a generative and explanatory framework that helps to uncover the key aspects of ML capabilities, contextual constraints, and user needs which are considered within an ideal interaction flow.

In parallel, van Duyne, Landay, and Hong's *The Design of Sites* expanded the conversation by focusing on the web design context – then a rapidly growing field. Acknowledging the increasing complexity of websites and the necessity for user-centric design, they compiled a collection of design patterns specifically geared towards customer-centred web experiences (van Duyne et al., 2003, pp. 3-5). Their work provided a structured framework to web design, offering patterns that addressed a wide array of challenges, from navigation and content organisation to e-commerce.

Jenifer Tidwell's *Designing Interfaces* further established the role of patterns in interaction design by providing a large collection of patterns that addressed common user interface problems across various platforms (Tidwell, 2005, pp. xiii-xv). One of the key aspects of Tidwell's work is its accessibility. She presented patterns in a clear and organised manner, accompanied by visual examples and practical advice that made the concepts easily understandable to designers of varying expertise levels. This approach facilitated the widespread adoption of

patterns in the industry, as it enabled designers to apply these solutions directly to their projects without requiring extensive theoretical background.

Following that vein, another more practical and less formal materialisation of the pattern approach emerged in the industry, translating the relevant concepts and best practices into components and user interface (UI) pattern libraries. This trend set the right context for HCI guidelines from key names in the technology industry such as Apple⁵ and Google⁶ to become widely used and often adapted for other digital UIs.

Finally, one of the most recent manifestations of the Pattern Theory core concepts is called Design Systems – “A design system is a collection of reusable components, guided by clear standards, that can be assembled together to build any number of applications”⁷. Good design systems enable the consistency of several aspects of the user experience, specific for a particular company or brand use case, suitable to its aesthetics, behaviour, interaction models and development needs.

Design systems are not necessarily pattern collections or pattern languages; however, they are executions of patterns delivering solutions to recurring user interface and interaction design problems for businesses or brands. A design system operates as a comprehensive framework that encompasses design best practices, reusable components, and guidelines, enabling teams to create consistent user experiences across diverse products and platforms.

In addition, as design systems typically represent both the visual design and coded representation of the UI components, their adoption often represent efficiencies beyond the UX practice, but also to the entire digital product development lifecycle. This is a clear example of the ways in which the industry has organically embraced and operationalised formal pattern principles.

Despite these pivotal contributions, most established UX pattern collections emphasise interface design, navigation, and usability heuristics without delving deeply into machine learning aspects. While personalisation and adaptive features have been implicitly referenced in past works (e.g., Borchers, 2001; Tidwell, 2005), few patterns address the complexities of integrating ML algorithms into the user experience and what considerations should be at the forefront (e.g. personal data usage). Unlike traditional interface patterns, ML features require designers to account for user trust, data-driven adaptivity, and algorithmic opacity. These considerations demand an updated perspective that addresses not just visual or

⁵ <https://developer.apple.com/design/human-interface-guidelines/>

⁶ <https://material.io/components>

⁷ Source <https://www.invisionapp.com/inside-design/guide-to-design-systems/>

interaction flows, but also the ethical and explanatory dimensions of predictive systems.

Consequently, there remains a gap in literature for dedicated ML-focused UX patterns – a gap that becomes increasingly significant as AI-driven features grow more prevalent. The following section thus examines how machine learning intersects with user-centred design, revealing both the potential benefits and the challenges that underscore this project's objectives.

2.6. The Integration of Machine Learning in UX Design Solutions

ML can enhance user experiences by enabling personalisation, predictive insights, and adaptive automation. In broad terms, ML approaches typically fall into three main categories. The first is supervised learning, in which models learn from labelled examples to anticipate user actions or preferences. Secondly, there is unsupervised learning, where the system clusters or segments user data to detect hidden patterns. Finally, interactive or incremental learning relies on ongoing input from end users to refine outcomes over time (Amershi et al., 2014).

Each of these paradigms carries unique implications for UX: supervised classification can reduce cognitive load by automating routine tasks (such as credit card scanning or voice input), whereas interactive learning can make interfaces more responsive to real-time feedback, as illustrated by CueFlik (Fogarty et al., 2008). Nevertheless, these advantages are tempered by risks associated with reduced transparency in 'black box' models (Dix, 2017), as well as ethical concerns around data privacy and bias (Lew and Schumacher Jr, 2020). Overlooking these dimensions can erode user trust, introduce undesired outcomes, or undermine the perceived credibility of ML-based features. These are issues that become especially critical when systems operate on sensitive personal data.

From a design perspective, the critical question is not so much on how ML works but rather how it can be integrated into interactions in ways that is natural, using existing design methodologies by practitioners without data science expertise. Therefore, in this context, machine learning is treated as an interactive feature embedded within the user experience, rather than an element that users themselves configure or develop. Designers conceptualise interactions that leverage ML to provide personalisation or automation, while end users simply benefit from these capabilities as part of a seamless UX flow. This approach aligns with the premise that most UX practitioners (and certainly most end users) do not need to understand the inner mechanics of ML but should feel confident and in control when engaging with ML-powered functionalities.

Typically, in more classic works such as *The Elements of Statistical Learning* by Hastie et al. (2009, 2001), each algorithm is presented with a detailed explanation of the mathematical underpinnings and breakdown of functions to enable the reader a step-by-step overview. Other texts on ML discuss concepts alongside implementations in popular languages such as Python (Chollet, 2017) or open source statistical packages like R (James et al., 2013). These titles are critical for practitioners in the industry and help to bridge the gap between statistical concepts and practical implementation.

For the context of this research for example, both practical and theoretical understanding of ML algorithms were useful to evaluate implementations that support user interactions and to uncover and document the forces that lead to particular patterns (solutions) as well as the consequences of adopting such patterns (Harrison, 2003; Wellhausen and Fießer, 2012). However, this ML knowledge is not a reasonable skillset to be expected from UX professionals, in similar ways that is not reasonable to expect UX practitioners to have practical skills in software engineering⁸.

It is encouraging to notice effective examples of the use of machine learning to reduce the cognitive cost of user interactions. Google Hum to Search⁹ is an excellent example of machine learning used to improve the user interaction, inspired by aspects of the human-human interaction. The objective is not to mimic human-human interaction, but instead to bring those elements of the natural human interaction which are feasible to translate into digital interactions. The result is a significant reduction in the cognitive load required for that user interaction with the system. My investigation also brought to light a much earlier example of use of ML in interaction design. In their 2008 paper, Fogarty et al. introduced CueFlik, a web image search application that used machine learning techniques to enable end user to specify rules to rank images based on their visual characteristics (Fogarty et al., 2008).

⁸ It is important to highlight that the field of UX typically welcomes practitioners from a variety of backgrounds, including software engineering. In fact, early background in Front-End Development was critical in facilitating my technical understanding of machine learning methods. Nonetheless, the requirement of both UX design and Development skills in job advertisements is usually criticised in social channels as an unreasonable practice.

⁹ <https://ai.googleblog.com/2020/11/the-machine-learning-behind-hum-to.html>

Another mature example is the Credit card scanner¹⁰. This interaction pattern has been quietly available at least since 2014¹¹ and it relies on mature computer vision approaches¹² to deliver Optical Character Recognition (OCR). As a result, smartphone users can avoid the inconvenience of data-entry by automatically scanning their credit card details with a camera.

Nonetheless, the discussion on approaches for integrating machine learning into design solutions from a UX practitioner perspective was limited at the point I began this research project.

In "Power to the People: The Role of Humans in Interactive Machine Learning," Amershi et al. argue that involving end-users is crucial for creating more effective and user-friendly systems (Amershi et al., 2014). They highlight that Interactive Machine Learning (IML) is distinct from traditional machine learning by featuring 'rapid, focused, and incremental' learning cycles as a result of end user interactions. In contrast, the authors claim that in traditional machine learning workflows, the models are fine-tuned based solely on the interventions by machine learning practitioners.

Further, Amershi et al. describe how involving people in the process of training the models produce more accurate models. They suggest recommender systems as an example which integrates ML into digital experiences of end users. They observed significant performance improvement as a result of involving people to aid in the training of such models. However, introducing people in a user testing-like setup risks attracting unnatural behaviour, given the artificial nature of the setup, and therefore the feedback provided might not be as useful as it would be in a real-world usage scenario (Sharp et al., 2019, pp. 518-519).

Alan Dix drew attention to the resurgence of Artificial Intelligence (AI), commenting that although the technology and associated paradigms existed for many years, its early adoption was not widespread in main stream HCI (Dix, 2017). Nonetheless, the advent of big data and cloud computing reignited the interest in AI adoption for digital user experiences. He also lists the successful use of recommender systems and how natural language processing was being handled through large-scale initiatives of the time. However, he also highlights his concerns about an issue that to this date still worries both users and designers of systems that integrate AI/ML –

¹⁰ <https://www.linkedin.com/pulse/how-improve-payment-forms-using-four-interaction-parimala-hariprasad>

¹¹ <https://www.cnet.com/how-to/scan-your-credit-card-instead-of-keying-in-number-for-purchases-in-ios-8/>

¹² <https://cloud.google.com/vision/>

the lack of transparency inherent in some of the most successful architectures such as neural networks, classed as 'black box' algorithms. Dix mentioned the need for the process he called 'appropriate intelligence', which represents mechanisms to manage and present the results of black box algorithms in more comprehensible ways and to reduce impact of when they go wrong.

Nonetheless, expertise in machine learning is beyond the skillset of the average UX professional (Dove et al., 2017; Yang, 2018), thus expecting them to understand the technical stack supporting the examples above is not realistic. In recent years, some authors have begun to address this challenge by making the AI subject more accessible to UX practitioners. For instance, Gavin Lew and Robert M. Schumacher Jr. in 'AI and UX: Why Artificial Intelligence Needs User Experience' highlight the common issues of privacy concerns, impact of poor data quality, ethics and bias and lack of transparency whilst providing an overview on the classical machine learning techniques (Lew and Schumacher Jr, 2020). Unsurprisingly, the text introduces UX practitioners to the subject of AI and ML, but it does not provide technical detail to implement or train models as this aspect is typically handled by ML engineers or data scientists.

In more foundational works such as 'Human-Centered AI', Ben Shneiderman emphasises the point of conceiving AI systems that augment human abilities instead of replacing them. He argues that ML techniques can be brittle when facing novel situations as it lacks the human abilities of common sense and higher cognition (Shneiderman, 2022, pp. 20). This perspective aligns with the need for UX designers to understand ML algorithms not from a technical standpoint but in terms of how they impact the user experience.

Shneiderman also discusses the concepts of rationalism and empiricism in the context of human centred AI. While rationalism is based on logical analysis, relying on design decisions based on formal methods that can be accomplished 'in the comfort and familiarity of their office desk', empiricism proposes that the design activity must happen through one's insertion in the real world to sense 'all its contextual complexity, diversity, and uncertainty' (Shneiderman, 2022, pp. 19). However, the differences between these two approaches will be more nuanced than that in practice and the author does explain that there will be scenarios where both approaches should be used to achieve a successful outcome.

From a pattern theory perspective, Shneiderman briefly introduces the idea of a Human Centred AI (HCAI) Pattern Language (pp. 76-78), in which the author's own Eight Golden Rules (Shneiderman and Plaisant, 1987) is deeply associated with. The practice of linking patterns with design principles is not novel in the HCI pattern

community, with early seminal work such as van Welie et al (2001) using similar approach. However, the lack of referencing other well-established principles could be seen as biased. Further, the role of patterns in 'Human-Centered AI' is solely introductory and beyond its brief mention in pages 76–78, the HCAI pattern language is no longer mentioned in the book.

Admittedly, 'Human-Centered AI' is not positioned as a work with the intent of promoting a pattern approach to deliver a more humanised user experience with AI. Instead, the book focuses on principles and best practices as a framework that is useful to all typical functions in modern digital product development, not only the HCI and UX community, although, given the author's reputation in these fields, it is likely to be the primary audiences the book is aimed at. Aspects such as Reliability, safety and trustworthiness are central to that framework.

More recently, Malizia and Paternò (2023) examine why current explainable AI (XAI) approaches are not meeting the expectations of practitioners across application domains. They argue that explanations satisfying a technical audience may offer little value to non-technical users, and that the cross-disciplinary nature of the explainability challenge demands solutions that go beyond purely algorithmic transparency. Notably, they call for what they defined as 'Understandable AI' (AI that serves the needs of non-technical stakeholders) and identify home automation as one of the domains where such approaches are needed. This perspective aligns with the pattern-based approach adopted in this research, where the aim is to provide UX designers with actionable, audience-appropriate guidance rather than low-level algorithmic detail. Turchi et al. (2024) advance this line of inquiry by proposing a research agenda around adaptive XAI interfaces that tailor explanations to diverse user needs and contexts, with their target audience explicitly including UX designers – further reinforcing the relevance of audience-sensitive approaches to communicating AI capabilities.

With regards to Consumer IoT, 'Human-Centered AI' covers classic examples such as Roomba, Nest thermostat and the role of virtual assistants such as Amazon's Alexa and Apple's Siri along with that of metaphors in user adoption. It contains a detailed account of the history of Social Robots - which can be considered a more anthropomorphic addition to the Consumer IoT context. The author also provides his insights on how the future of such appliances is intimately linked to the user needs, for instance, elder care and accessibility needs. Nonetheless, the book does not venture into more practical mechanisms to support designers – or any other function - to design Human Centred AI experiences in the context of Consumer IoT. This could be seen as a missed opportunity for better positioning and expansion of the HCAI pattern language role as a practical tool.

Overall, Shneiderman's central argument in 'Human-Centered AI' is the fact that intelligent systems should augment rather than replace humans and this work has a timely relevance for the community of practice.

2.7. The Consumer Internet of Things as a Context for User Interaction

The consumer IoT environment – ranging from smart thermostats to wearables – presents a fertile ground for ML-powered user experiences for several interrelated reasons. Firstly, these devices often involve frequent, repetitive interactions, such as daily use of a smart speaker, which generates substantial data that can inform adaptive interfaces (Kuniavsky, 2010). Secondly, consumer IoT is comprised of devices with distinct hardware, software, and connectivity layers, supporting the need for robust design patterns that ensure usability across diverse contexts (Rowland et al., 2015). A third important aspect is that the expanding market for consumer IoT (Lee and Lee, 2015) creates growing demand for solutions that integrate ML effectively.

Some authors have focused on areas very closely related to this research. Vega-Barbas et al. (Vega-Barbas et al., 2018) proposed patterns addressing interaction problems with a Smart Space, which uses IoT technology to provide an ubiquitous computing environment whereas Kuniavsky (2010) highlights the opportunity of 'information processing as a tool for experience design'. This perspective directly aligns with the opportunity for personalisation and context-awareness

Nonetheless, despite the steady growth of IoT in the past decade, designing for IoT experiences is a relatively new problem for the industry. The IoT Design Deck (Dibitonto et al., 2018) is defined as 'a method for the co-design of connected products' and was developed to assist with three stages of the IoT design challenge: Design, Facilitation and Prototype. It anticipates that designing for IoT should be a multi-disciplinary effort of co-creation. The toolkit shares an objective of this research project of helping designers that are not technology experts and the patterns I documented can be used to complement it in the design and prototyping stages.

Past reports (PricewaterhouseCoopers, 2017; YouGov, 2018), indicate that people are acquiring single connected devices rather adopting fully ubiquitous computing environments. These users, however, did not seem to dislike the idea of greater interconnectivity. Nonetheless, there is a significant discrepancy between the reality of the consumer IoT and the state-of-the-art environments frequently cited in the literature (Chung et al., 2010; Kuniavsky, 2010; Poslad, 2011; Vega-Barbas et al.,

2018), and it is very unlikely that advanced interaction mechanisms such as Brain Computer Interfaces (Zhang et al., 2018) will be seen in the average household anytime soon.

In the past few years, IoT as a field has benefited mainly from a technically driven perspective (Guo et al., 2013; Rowland et al., 2015, pp. 3, 27–28). Still, unless consumer IoT users are willing to venture towards End-User development (Caivano et al., 2018; Fogli et al., 2016), the complexity of the multitude of vendors could only increase the risks of inconsistent and poorer experiences.

The present context of design for the consumer IoT experience is a far cry from the ubiquitous computing envisaged in the 1960's when Jim Sutherland created the first home automation device, the ECHO IV (an acronym for Electronic Computing Home Operator) (Deschamps-Sonsino, 2018, pp. 68–71) or in the more ambitious 1980's Smart House project (Smith, 1988).

The 2015 work by Rowland et al. titled 'Designing Connected Products: UX for the Consumer Internet of Things' provides a comprehensive examination of the complexities associated with designing user experiences for IoT products and, at the time of writing, remains unparalleled in its depth and breadth on the subject.

The authors' central argument is that designing for the consumer IoT demands a holistic and system-oriented approach that goes beyond traditional UX design methodologies. They assert that IoT products are not isolated entities but part of intricate ecosystems involving hardware, software, cloud services, and user interactions spread across multiple devices and contexts. Therefore, they advocate for a design perspective that considers the entire user journey, encompassing the physical and digital touchpoints that collectively contribute to the user's experience.

A key aspect of their argument is the concept of "interusability," which refers to the consistency and coherence of user experiences across different devices and platforms within an IoT ecosystem (Rowland et al., 2015, pp. 165-170). They emphasise that as users interact with various interconnected devices, maintaining a seamless and intuitive experience is critical for user satisfaction and the overall success of IoT products.

In addition to covering typical concerns around consumer IoT experience, such as fragmented user journeys, security concerns, and the management of complex data flows, context of use, amongst other, the authors highlight that designers must adopt new tools, methodologies, and mindsets to address the unique challenges posed by IoT.

To substantiate their arguments, Rowland et al. draw upon a well curated list of case studies, design principles, and practical frameworks. Similarly to the key principle behind the pattern theory that is central to this research project, they explore real-world examples like the Nest Learning Thermostat and the Philips Hue lighting system, illustrating how these products provide good UX through the use of intuitive controls and adaptive functionalities (Rowland et al., 2015, pp. 123-130). These examples demonstrate how IoT devices can learn from user behaviours and preferences to provide more relevant and efficient interactions.

Some of the examples and technological contexts, however, might not accurately reflect the current state of IoT and ML integration given the speed at which technology has advanced since the book's 2015 release. Advances in edge computing, voice assistants, and artificial intelligence have expanded the potential and expectations of Internet of Things devices. The fundamental ideas presented by the writers are still applicable, but their insights would be even more useful from a updated perspective that takes these developments into account.

The book was written at a time when the consumer IoT market was experiencing an increase in interest and optimism. Advances in wireless connectivity, sensor technologies, and mobile computing sparked the widespread adoption of smart devices in the mid-2010s (Gartner, 2015; Lee and Lee, 2015; Taylor et al., 2018). The authors' emphasis on innovation and the investigation of novel areas in UX design for connected products may have been impacted by this background.

However, since then, there has been a growing awareness of the ethical implications of IoT and ML technologies, particularly concerning data privacy, surveillance, and algorithmic bias. There has been an increase of interest in the societal impact of technology, which suggests that designers must deal with these issues more explicitly. The authors' work, while comprehensive for its time may benefit from a deeper focus on these emerging considerations to make it more relevant to the current industry context.

While Rowland et al. (2015) effectively capture the complexities of designing connected products, their text predates many of the significant advances in AI and cloud-based analytics that shape today's IoT. By weaving machine learning capabilities into interaction design methods, this research extends the foundational principles they proposed. Specifically, I propose a pattern-driven approach that encapsulates design solutions featuring ML-integration. In doing so, I aim to help UX professionals address emerging concerns – such as ethical data handling and user

control – while still benefiting from wisdom of contributions such as Rowland et al.'s holistic approach to interconnected products.

2.8. The Research Gap

Considering the body of work presented in this chapter, I am able to highlight a clear gap at the intersection of pattern-based design, machine learning integration, and consumer IoT. Most established design pattern libraries offer limited focus on opportunities for incorporating ML algorithms into user-centred solutions, leaving professionals seeking structured, reusable approaches with not many resources. Furthermore, although standard ML textbooks (e.g., Hastie et al., 2009) provide comprehensive algorithmic insights, they typically expect a level of technical aptitude that exceeds the typical skill set of many UX practitioners (Dove et al., 2017). The issue becomes even more evident in the fragmented, multi-device ecosystems characteristic of consumer IoT, where designers must manage concerns relating to user needs, privacy protections, and system dependability (Rowland et al., 2015). As ML-driven technologies continue to evolve, issues of ethics and transparency – such as explainability, the possibility of biases, and securing informed consent – grow increasingly salient, yet older pattern frameworks rarely engage with these topics more explicitly.

Liao and Vaughan (2024), writing at the intersection of HCI and responsible AI, argue that transparency should be designed to support appropriate human understanding for different stakeholders in different contexts. Their human-centred roadmap highlights the need for transparency approaches that are usable, context-sensitive, and attentive to issues such as explanation, uncertainty, recourse, and stakeholder needs. While they do not explicitly propose pattern-based design resources, their analysis is compatible with efforts to translate responsible-AI considerations into practical guidance for designers.

A related contemporary strand of work also points toward AI-supported pattern use within UX practice: Lu et al. (2024) explore AI-driven pattern annotation in multi-screen user flows, suggesting that structured pattern-based support is an active topic in the wider UX field, beyond the Consumer IoT context.

This thesis demonstrates how to mine, document, and validate a pattern language that enables UX designers lacking specialised ML expertise to conceptualise adaptive consumer IoT experiences. By linking overarching ML principles to concrete design strategies, the patterns I introduce here aspire to advance both the theoretical discourse of pattern authoring and the practical challenges of designing more intuitive, user-friendly IoT experiences.

The next chapter outlines the methodology through which I approached the discovery, documentation, and evaluation of a novel set of ML-oriented UX patterns for consumer IoT scenarios.

3. Methodology

3.1. Introduction

In this chapter I cover the methodological approach adopted for this research, which aimed to identify, document and validate user experience (UX) patterns incorporating machine learning (ML) within the context of Consumer Internet of Things (IoT) applications.

As a practising UX professional who is deeply involved in the wider UX-design community, I initiated this research project with three working assumptions. First, I observed that (although the reasons were unclear), UX practitioners were rarely proposing or conceptualising design solutions that made use of machine learning. Second, the profession is well versed in the notion of best practices and in reusing solutions that embed those best practices. Third, it is common for UX professionals to design for media whose technical details they do not fully understand, like complex web architectures, for example, yet, when having access to relevant design

principles and guidance, they often provide viable solutions. These assumptions guided the methodology described in this chapter.

3.1.1. Paradigmatic stance and case for mixing methods

My guiding paradigm is **pragmatism**: each approach for a given phase of the inquiry was chosen based on what Robson (2011) describes as “what works” rather than allegiance to a philosophical dogma. Pragmatism legitimates selecting and mixing methods if they help to address the research questions, and It endorses ‘strong and practical empiricism’ as a mechanism to test potential solutions (Johnson and Onwuegbuzie, 2004; Robson, 2011).

To answer the research questions, the methodology was structured into four key phases:

1. **Machine Learning Experiments:** I explored potential implementations of ML to understand the level of technical information the patterns should contain. In addition, I experimented with ML approaches to facilitate automated image labelling in an attempt to identify common UI patterns used in IoT apps.
2. **Pattern Mining:** This phase involved identification of UX design patterns, specifically tailored for Consumer IoT context, which incorporated machine learning to deliver adaptive experiences.
3. **Pattern Documentation:** The UX patterns identified were formally documented, with most of them gone through peer review and closely examined through the pattern writer workshops at the European Conference on Pattern Languages of Programs.
4. **Validation Study:** The final phase assessed the effectiveness of the novel UX patterns. Through a validation study involving experienced UX designers and an IoT expert, I provide a proof of concept approach to test the patterns’ potential applicability and usability in real-world scenarios.

These deliberately sequential and connected phases aligned with the following strands:

- Empirical/ feasibility testing (phases 1 and 4) – to ground the work in technical reality. In addition to informing the level of algorithmic description and technical detail suitable for the emerging patterns, these phases enhance credibility by demonstrating that the research was based on realistic rather than purely hypothetical scenarios.
- Delphi-style consensus building (phase 2) – to uncover and refine expert design knowledge. This is a recognition that implicit design knowledge is

distributed across a community of practice, I used a modified Delphi process (Chuenjitwongsa, 2017; Hasson et al., 2000; Iqbal and Pipon-Young, 2009), which involved a panel of UX practitioners and iterative feedback, while remaining time-efficient for these professionals.

- Qualitative analysis in phase 3 drew on grounded-theory coding procedures to develop a conceptual structure for the phenomena that emerged. Rather than imposing categories from external theories, the analysis sought to generate concepts grounded in participants' accounts through systematic comparison of workshop transcripts, sketches, and notes. These materials were analysed through open, axial, and selective coding (Corbin and Strauss, 2014).

By combining empirical testing, Delphi-style consensus and grounded-theory analysis in a deliberately interconnected sequence, I established a sequence of evidence. Quantitative experiment results were used to inform documentation suitability, the Delphi-style approach and qualitative coding helped to uncover and translate practitioners' knowledge into pattern form. Together these strands provided a conceptual framework aimed at drawing robust conclusions.

Each study chapter that follows reports its methods using a consistent structure comprising Participants (or data sources), Apparatus and Materials, Procedure, and Data Collection and Analysis, before presenting the corresponding findings.

3.2. Linking the Research Questions to the Methodology

This methodology was devised to address two core research questions:

RQ1: Can a pattern language that encapsulates the use of machine learning to support user interactions with Consumer IoT be identified?

RQ2: If identified, can this pattern language be used by designers with little expertise in machine learning to produce feasible Consumer IoT interaction concepts?

Phases 1, 2 and 3 (the machine learning experiments, pattern mining and the pattern documentation) primarily address RQ1 by generating empirical insights into feasible ML enhancements, capturing emerging solutions and formally documenting them. Phase 4 (validation study) address RQ2 by and evaluating their potential application in real-world design scenarios.

3.3. Ethical considerations

This research followed the institutional ethics guidelines from the University of Hertfordshire particularly regarding participant privacy and informed consent. Participants in the pattern mining workshops and the validation study were informed about the study aims, format, and their right to withdraw at any time. Each participant provided written or electronic consent prior to data collection.

Data Anonymisation

To maintain confidentiality, participants were assigned unique codes (e.g., D1-D5). Any recordings or transcripts used for analysis were stored on password-protected drives, accessible only to the researcher and supervision team. Similarly, any references to identifiable user data were de-identified before integration into the dissertation (with exception of the IoT Expert who was happy to be identified in this work).

Compliance and Approval

Ethical clearance was obtained from the University of Hertfordshire Social Sciences, Arts and Humanities Ethics Committee with Delegated Authority. The UH protocol numbers are aCTA/PGR/UH/03834 and aCTA/PGR/UH/05871. The data collection involved minimal risk and no vulnerable participants.

3.4. Validation study

The final phase assessed the effectiveness of the novel UX patterns through a proof-of-concept study involving experienced UX designers and an independent IoT expert. This phase primarily explores Guiding Proposition 2 (GP2). The full methodology, including ethical considerations, research design, participant recruitment, apparatus, procedure, data collection, and limitations, is reported in Chapter 7 alongside the results.

3.5. Summary of the Methodologies

The pattern mining process successfully addressed the challenges inherent in engaging UX professionals with limited ML and IoT expertise at the time. By focusing on familiar devices and leveraging the WWHD exercise, participants provided useful insights into how ML could enhance user interactions in consumer IoT applications. The data collected provided seeds (foundational concepts) for developing formal UX patterns, based on the collective wisdom of the community of practice.

The methodology, while not without limitations, offered a practical approach to pattern mining in a complex context. As the ML technologies becomes more widespread, future research may find feasible to restrict the participant pool to only professionals with deeper expertise in ML and IoT.

The comprehensive process of pattern identification, verification and peer review contributes to the documentation of sound UX design patterns for machine learning-enhanced (also known as adaptive) Consumer IoT experiences. By using appropriate aspects of human-human interaction as a form of 'filter' or starting point, and adherence to established UX and human-AI interaction principles, the resulting patterns could aim to facilitate more natural and effective user interactions in consumer IoT contexts.

Overall, Phases 1, 2, and 3 thoroughly explored the identification and documentation of an ML-enhanced UX pattern language – thereby aligning with **Guiding Proposition 1 (GP1)** – while Phase 4 validated these patterns in practice, **exploring Guiding Proposition 2 (GP2)** by examining the performance of the patterns in enabling non-expert designers in applying the documented solutions to create feasible concepts.

By providing a comprehensive account of the research design, participant recruitment, materials and tools used, data collection methods, data analysis techniques, and the limitations of each approach (detailed in the study chapters that follow), I am confident in the rigour and validity of this research project for the field of UX Patterns in the context of the Consumer IoT.

4. The Experiments Using Machine Learning

4.1. Introduction

In the initial phase of my research project, conducted in 2018, I concentrated on exploring the relationship between machine learning (ML) and graphical user interfaces (GUIs) within the context of Consumer Internet of Things (IoT) applications. My aim was to investigate how ML could be used to enhance GUI components, potentially leading to the development of a pattern language for UI design in Consumer IoT. Two key experiments were conducted during this phase, namely, Simplified User Identification and Automated Labelling using Convolutional Neural Networks (CNN), each addressing an exploratory sub-proposition linked to the research questions.

These exploratory studies served two purposes. Firstly, they enabled me to decide the level of technical ML detail to incorporate in a pattern language for non-ML

specialists. Secondly, they influenced the strategic decision to transition from narrowly focusing on screen-level components to a broader UX pattern scope. Establishing these constraints early on helped me to ensure that subsequent pattern mining efforts were practically grounded and avoided unnecessary technical considerations.

The first experiment (Automated Labelling) was associated with RQ1 – 'Can a pattern language that encapsulates the use of machine learning to support user interactions with Consumer IoT be identified?' leading to Sub-Proposition 1.1 (SP 1.1) – 'The Rico dataset, combined with computer vision techniques, may offer a viable approach for identifying and filtering UI screens of Consumer IoT apps, potentially facilitating the discovery of common UI patterns and providing a starting point for exploring machine learning-enhanced versions of these patterns.'

The second experiment (Simplified User Identification) was associated with RQ2 – 'If identified, can such a pattern language support the effective design of user interaction concepts for Consumer IoT scenarios by designers with no machine learning expertise?' and led to Sub-Proposition 2.1 (SP 2.1) – 'The formal documentation of novel patterns may be more beneficial to its users if it includes detailed information specifically aimed at supporting the implementation of associated machine learning methods.'

In the sections that follow I present each experiment in detail, beginning with the Simplified User Identification study and followed by the Automated Labelling experiment. For each, I describe the data sources, apparatus, procedure and results before discussing the combined implications for this research.

4.2. Simplified User identification

This experiment was aimed at exploring the level of detail which should be covered in the machine learning (ML) aspect of the pattern documentation. At this stage my goal was to determine how to provide more comprehensive technical documentation within the patterns, with the intention of including examples such as algorithmic performance results and other relevant information, thus addressing Sub-Proposition 2.1 (SP 2.1).

User identification is a fundamental aspect of personalisation in Consumer IoT environments. Traditional methods such as passwords, PINs, or biometric scans often introduce friction into the user experience, particularly in contexts where seamless interaction is expected. In multi-user settings, common in households with multiple occupants interacting with shared IoT devices, unobtrusive user identification becomes very relevant. A system that can identify users passively, without explicit input from them, can significantly enhance the user experience by reducing barriers to interaction.

I explored a typical use case of leveraging widely available hardware – specifically accelerometers in smartphones or wearable devices – to implement gait recognition as a method for simplified user identification. Gait recognition analyses the unique walking patterns of individuals, offering unobtrusive means of identification that can enhance user experience without compromising security. The uniqueness of an individual's gait has been studied extensively and is considered a viable biometric trait (Murray, 1967; Nixon et al., 2010).

4.2.1. Participants and Data Sources

To investigate gait recognition, I utilised a dataset from the University of California Irvine (UCI) Machine Learning Repository (Casale et al., 2012). The original dataset contains accelerometer readings from 22 participants who walked along a predefined path, capturing the nuances of their gait patterns. The data was collected using the SENSOR_DELAY_FASTEST setting, which configures the accelerometer to record data at the highest possible frequency. This setting ensures that the maximum number of data points is captured, providing a detailed representation of each participant's gait.

Each participant's data is stored in a distinct comma-separated values (CSV) file, containing the following attributes:

- **Time-step:** Representing a sequential measure indicating the position of each reading of the sensor data.
- **X-Acceleration:** Acceleration along the x-axis.
- **Y-Acceleration:** Acceleration along the y-axis.
- **Z-Acceleration:** Acceleration along the z-axis.

For practical reasons, I selected data from two randomly chosen participants, resulting in a combined dataset of 8,951 observations. Focusing on two participants allowed for a controlled binary classification task, simplifying the initial exploration while providing insights into the feasibility of the approach. I assigned labels to the two participants data (e.g., Person A and Person B) to facilitate binary classification.

4.2.2. Apparatus and Materials

I used a standard MacBook Air (13-inch, Mid 2012), 256GB flash storage, processor 1.8GHz dual-core Intel Core i5, 4GB of 1600MHz DDR3L onboard memory with Intel HD Graphics 4000. No external Graphics Processing Unit (GPU) was used. The operating system was macOS version 10.15 (Catalina), running Anaconda Navigator 1.6.7 – a desktop graphical user interface that facilitates the launch of applications especially aimed at data science tasks and the management of their extension packages and libraries.

I used the Scientific Python Development Environment (Spyder) IDE 3.2 which is included with Anaconda Distribution, and it was equipped with Python 3.6. I also used the following libraries to enable ML methods:

- NumPy 1.13.0
- Pandas 0.21.0
- Scikit-learn 0.19.0

Spyder is an open-source¹³ Integrated Development Environment (IDE) purposely tailored for scientific-computing workflows in Python and therefore well suited to exploratory machine-learning experiments – such as those conducted in this research project. Finally, spyder is distributed with key libraries needed in this study (e.g. NumPy, pandas and scikit-learn) pre-installed, therefore making it an ideal solution.

4.2.3. Procedure

Before training the machine learning models, the data underwent a key preprocessing step:

Data Splitting – I divided the dataset into a training set comprising 75% of the data (6,713 observations) and a test set comprising 25% (2,238 observations). This split approach is commonly used in machine learning to provide sufficient data for model training while reserving a portion for unbiased evaluation (Arifin and Yaacob, 2025; Hastie et al., 2009 pp. 222; James et al., 2013 pp. 176-177; Li et al., 2024).

The Machine Learning Models used

I explored three supervised machine learning algorithms to see how well they could distinguish between two individuals based on their walking patterns (gait). These three algorithms were selected because they are classic approaches in supervised machine learning, often recommended in authoritative texts (Hastie et al., 2009 pp. 17, 119; James et al., 2013 pp. 127). By using them, I was able to take advantage of the fact that these algorithms have their typical behaviour and

¹³ <https://github.com/spyder-ide/spyder>

implementation details well documented, allowing me to quickly produce realistic simulations while keeping computational and expertise challenges manageable. Here's a summary of each method I used and how they differ:

Support Vector Machines (SVMs) (Hastie et al., 2009):

SVMs are supervised machine learning algorithms used for classification tasks. The principal objective of a Support Vector Machine is to identify the optimal hyperplane (optimal boundary) that most effectively divides observations from distinct classes within a high-dimensional feature space. This hyperplane defines the distance between the hyperplane and the closest data points from each class, which are called support vectors. SVMs can handle complicated datasets by employing these support vectors as references, and when kernel functions are included, they can model non-linear connections. A kernel is a function that allows the method to operate on an implicit, high-dimensional feature space, avoiding the need to explicitly compute the coordinates of the data inside that space. The kernel function computes the inner product (called similarity measure) between pairs of data points in this transformed feature space. By using kernels, SVMs can find hyperplanes that separate different classes, even when the data cannot be separated with a linear boundary in the original input space.

This method enables SVMs to represent complicated, non-linear connections by efficiently changing the original input data into that higher-dimensional space in which linear separation between classes is possible.

Linear Kernel SVM assumes a simple, linear connection between the input variables (such as gait data) and output classes (the two persons). It works well with data separated by a straight line (or a flat surface in higher dimensions). If the data patterns are plain and unambiguous, a linear kernel is efficient and simple to interpret. In contrast to the linear version, the **Radial Basis Function (RBF) Kernel SVM** may capture more complicated, non-linear relationships by mapping the data into higher dimensions. The RBF kernel can provide flexible boundaries to separate data points that are not linearly separable in their original form.

K-Nearest Neighbours (KNN) (Hastie et al., 2009):

K-Nearest Neighbours (KNN) is a straightforward algorithm used for classification and regression tasks. In classification (this use case), it predicts the class of a new data point by finding the 'k' nearest data points (neighbours) in the feature space, using a distance metric such as Euclidean distance. The new data point is subsequently labelled with the class that is most common among its closest neighbours.

- **Using Limited Features:** Initially, I trained the KNN model using only a few features (the time-step and x-acceleration) to see how selecting certain features affects the model's performance. KNN classifies a data point based on its 'k' closest neighbours in the feature space. Using fewer features simplifies the model but might reduce its accuracy because it has less information to work with.
- **Using All Features:** Next, I included all available features (time-step, x, y, and z acceleration) to assess whether a broader set of data would improve the model's effectiveness

Logistic Regression (Hastie et al., 2009):

This is a widely used method for problems with two possible outcomes. It estimates the probability of one of the two classes by applying a logistic function to a combination of input features. While it's less complex than some other models, it's often a strong starting point for binary classification tasks because it's straightforward and interpretable.

Differences Between the Models:

The models differ in several key aspects. Firstly, they handle the complexity of relationships in the data differently. The RBF Kernel SVM and the KNN model using all features are better at capturing complex, non-linear relationships. They can model intricate patterns in the data that linear models might miss. In contrast, the Linear SVM and Logistic Regression assume that the relationships between variables are linear, which makes them simpler but possibly less effective if the true relationships are non-linear.

Moreover, the computational requirements vary among the models. Logistic Regression and Linear SVMs are generally faster and require less computational power, making them suitable for larger datasets or when resources are limited. On the other hand, the RBF Kernel SVM and KNN are more computationally intensive, especially as the dataset grows in size and dimensionality. This is because they involve more complex calculations.

This is relevant, given that at this point the aim was to identify the right amount of ML technical information to add to the formal documentation of the patterns. Depending on the implementation, computational demand could represent a critical aspect to consider and if relevant, it could be mentioned in the documentation.

Model Training and Evaluation

Each model was trained using the training dataset and tested on the test dataset. To analyse their performance, I utilised a table called the **Confusion Matrix**, which is used to evaluate classification models. This table displays the overall number of true positives, false positives, true negatives, and false negatives (Dankers et al., 2019). To further assess each model's effectiveness, I also calculated the **F-score** – a single value that represents how often the model's positive predictions are correct and how many of the actual positive cases it successfully identifies – which provides an additional indicator of predictive capability.

I used 10-fold cross-validation on the 75% training set with sklearn's `cross_val_score`. For each candidate value of **k** (number of neighbours) ranging from 2 to 11, the classifier was trained on nine of the ten folds and validated on the remaining fold; these steps were repeated ten times, and the reported accuracy is the mean across those ten validation folds. This process was used to identify the hyper-parameter **k** that delivers better performance, based on the reported accuracy.

It is important to emphasise that the purpose of this experiment was to replicate the key stages involved in applying machine-learning techniques. The simulation generated insights such as how to evaluate and select algorithms for example and aimed to determine whether that depth of technical detail would be appropriate for the pattern's intended audience: UX practitioners.

4.3. Simplified user identification experiment results

The main guiding question in this section was how complete the ML details ought to appear in a formal pattern documentation – especially if the patterns are intended for UX designers with limited ML expertise. To explore this, I conducted a **simplified user identification** experiment using common smartphone accelerometer data to recognise individuals by their walking patterns. This scenario is highly relevant because many Consumer IoT devices are shared among multiple users and can benefit from seamless, passive user identification – a realistic use case.

By testing the feasibility of gait recognition with straightforward ML techniques, I aimed to generate realistic technical details that would typically be necessary to document (e.g., algorithms selection process, tuning parameters, results) so that a final working solution could be replicated. It remains important to highlight that the scope of this experiment was solely at algorithmic performance level using a previously generated dataset - building end user applications was out of scope.

These are the results obtained from all classifiers tested in the experiment:
(TN = true negatives, FP = false positives, FN = false negatives, TP = true positives)

Support vector machine (SVM), using linear kernel (Hastie et al.,2001):
TN=845, FP=130, FN=209, TP=1054

SVM, using rbf kernel (Hastie et al.,2001):
TN=928, FP=47, FN=76, TP=1187

Logistic Regression (Hastie et al.,2001):
TN=784, FP=191, FN=197, TP=1066

K-nearest neighbours (KNN), using only two features – **time-step** and **X_Acceleration**:
TN=860, FP=115, FN=200, TP=1063

KNN using all four features:
TN=935, FP=40, FN=65, TP=1198

The confusion matrix comparison displayed in Figure 2 shows the results for the data from the two individuals selected for this experiment.

The best performance was recorded using the k-nearest neighbours' algorithm (Cover and Hart, 1967; Hastie et al., 2001) with 9 neighbours and 10-fold cross validation achieving **95% accuracy** and **F-Score 0.96** achieved from the test set (25% of the entire dataset). Results for all classifiers can be seen in Table 4.

Classifier	TP	FP	FN	Precision	Recall	F1-Score
SVM (Linear Kernel)	1054	130	209	0.89	0.83	0.86
SVM (RBF Kernel)	1187	47	76	0.96	0.94	0.95
Logistic Regression	1066	191	197	0.85	0.84	0.85
KNN (Two Features)	1063	115	200	0.90	0.84	0.87
KNN (All Features)	1198	40	65	0.97	0.95	0.96

Table 1 - F-Scores breakdown for each classifier based on the confusion matrices

These findings confirm that KNN with a more comprehensive feature set outperforms the other methods tested. Its confusion matrix displays the fewest misclassifications, further supporting its suitability for this experiment's use case: accurate user identification based on their gait data.

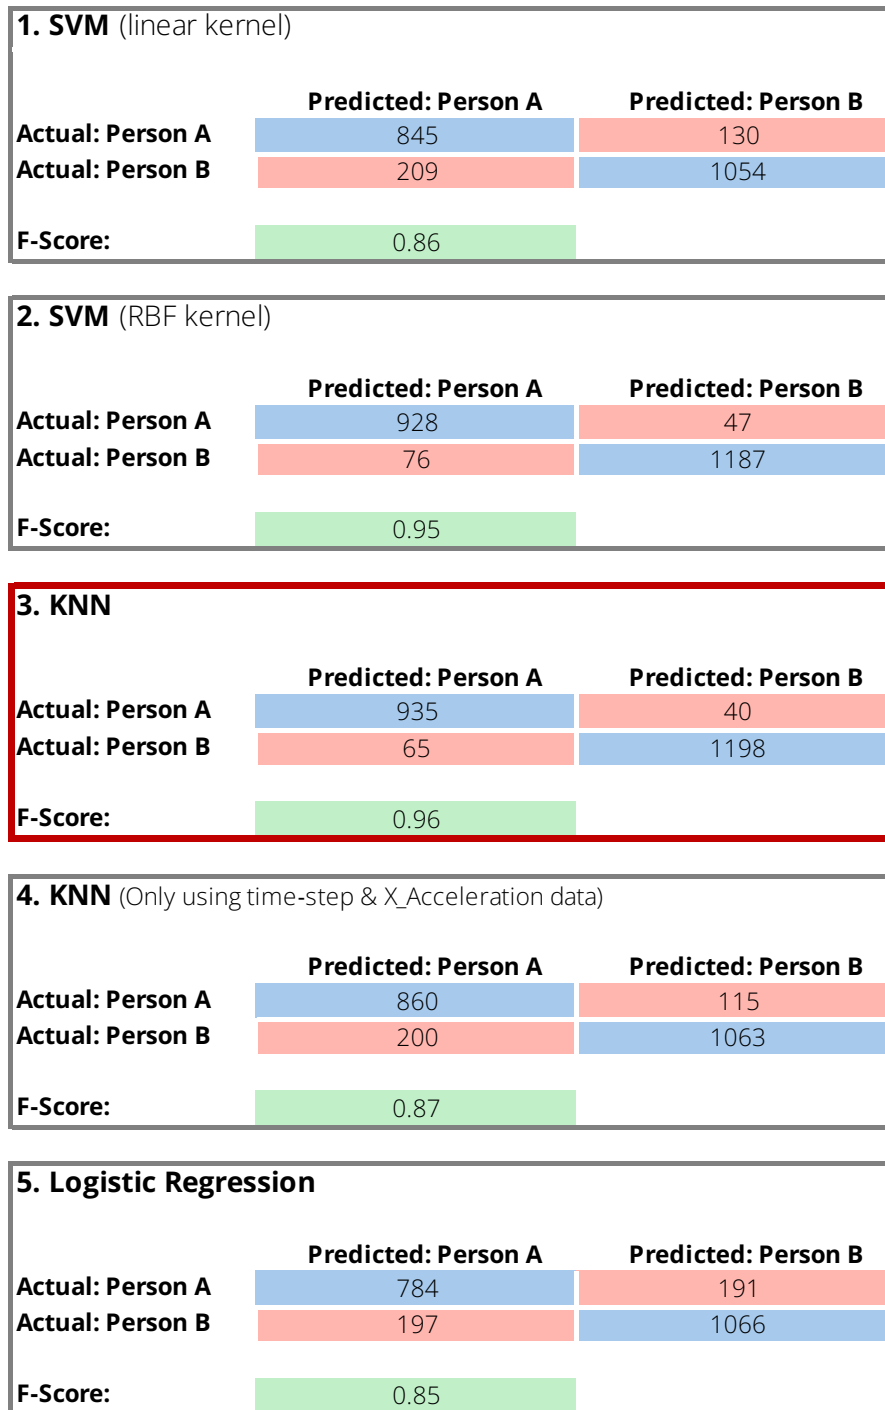


Figure 2 - Comparison of confusion Matrixes showing the results for different ML approaches. Approach 3 (KNN) achieved the best results.

4.3.1. Deliberation on the level of ML-information to include in the patterns

Through this experiment, I aimed to determine whether detailed algorithmic procedures – such as confusion-matrix interpretation and parameter-tuning specifics – would benefit non-expert UX designers in a pattern-based methodology.

Multiple classification trials (SVM with linear and RBF kernels, Logistic Regression, and KNN) confirmed that gait recognition is technically feasible, but the performance details had negligible bearing on core UX concerns. Although I observed that KNN achieved 95% accuracy, the nuances of kernel functions or numeric thresholds – i.e. the type of information that could potentially end up in a section of the pattern documentation – did not map directly onto practical UX design decisions.

Moreover, the fleeting relevance of any algorithm-specific guidance presented a significant challenge. The field of machine learning evolves so rapidly that embedding detailed parameter advice in formal design patterns could risk obsolescence within months. Consequently, the findings suggest that the expectation expressed in Sub-Proposition 2.1 – that richer ML detail would be beneficial to non-expert designers – does not hold in practice. On the contrary, an abundance of technical depth might distract design teams from pressing user-centric questions – such as informed consent for data collection, mitigating user confusion, or addressing privacy needs. A higher-level, principle-based approach therefore appears more conducive to long-term, user-focused patterns.

In light of these considerations, I advocate for selectively incorporating ML information that emphasises ethical usage, feasibility, and broad experiential requirements rather than delving into parameter tuning or confusion matrix breakdowns for example. This approach preserves the patterns' usability for a largely UX-oriented audience, adapts to a continually evolving ML field, and upholds the user-centric ethos that underpins this entire research project.

4.4. Automated Labelling of Material Design UI Screens

This section is guided by the exploratory Sub-Proposition 1.1 (SP 1.1) – **The Rico dataset, combined with computer vision techniques, may offer a viable approach for identifying and filtering UI screens of Consumer IoT apps, potentially facilitating the discovery of common UI patterns and providing a starting point for exploring machine learning-enhanced versions of these patterns.**

It is important to highlight that the patterns discussed in this section – UI Patterns – are not the same type patterns ultimately documented through this research project. The main distinction being that UI Patterns address problems associated with layout structures, visual consistency and behaviour at the screen level in a given Graphical User Interface (GUI), whereas UX patterns focus on wider user flows, multi-step processes, feedback loops beyond a single point of interaction, often at a more holistic level and taking into consideration users emotional and behavioural factors.

Given that user interactions with GUI cannot be easily facilitated without user interface elements, and that the scope of this research focuses on smartphone apps, in the initial phase of this research project I devised a methodology aimed at identifying ideal UI patterns that could be augmented with machine learning.

The variety of UI designs among apps makes it difficult to standardise interaction patterns. In the context of Consumer IoT, where physical device form factors frequently constrain standard three-dimensional input modalities, mobile app interface design is even more important.

User interface (UI) components are critical in enabling user interactions and experiences in programs. I thus conducted experiments to see whether identifying the most frequently used UI components in Consumer IoT apps could provide insights into potential opportunities for machine learning enhancement. To that end, I employed a method that relied on convolutional neural networks (CNNs) to automate the detection of user interfaces that employed Material Design.

Material Design (MD) is a design language created by Google to ensure a consistent user experience across platforms and device sizes (Google, 2014). Previous study by Doosti et al. (2018) found that Android applications that used Material Design were more likely to provide better user experiences.

Doosti et al. utilised text mining and computer vision techniques to identify apps using specific Material Design components and correlated this with user reviews and found that apps using MD components received more positive reviews. Guided by their approach, I sought to replicate and adapt their methodology within the constraints of resources available to me. My goal was to filter the most used UI components in Consumer IoT apps, with the plan to identify the patterns they were executing and then research ML-augmented versions of such patterns. Figure 3 summarises the process and objectives.




Figure 3 - The approach for identifying good candidate Consumer IoT UI patterns to be enhanced with ML.

Challenges

There were several challenges in this endeavour:

1. **Dataset limitations** include a shortage of annotated datasets for UI components in consumer IoT apps. Public datasets like as Rico contain a large number of UI screens but lacked categorisation for IoT specific applications.
2. **Resource Constraints** – Limited computing resources and manual labour for picture annotation limited the experiment's scope. Training deep learning models such as CNNs takes significant processing power, which was not easily accessible to me at the time of this experiment (2018).
3. **UI Design Diversity** – Developers' varying implementations and customisations make automatic identification hard. Designers and developers can modify common components or create bespoke ones that mimic standard MD components.

4.4.1. Data Sources

I selected the Rico dataset (Deka et al., 2017) as it was one of the largest publicly available repositories of mobile app UI screens at the time. It contains 72,200 UI screens from 9,700 Android apps, along with corresponding view hierarchies in JSON format.

4.4.2. Apparatus and Materials

I used the **IBM Cloud Annotations Tool** (which is no longer operational). The tool used a web-based interface that facilitated the training of a model to automatically annotate datasets. It was chosen due to its promising ability to simplify object annotation, user friendliness and being free. At the time, based on a contemporaneous informal review, there were no similar tools matching these criteria.

The automatic annotation happened by feeding the tool with a few samples of images which were manually labelled: Once uploaded, for each image that contained a FAB in the UI, I indicated its location by drawing a rectangular boundary around it - directly via the IBM Cloud Annotations interface.

IBM did not publish technical documentation about hardware or software architecture for the tool, other than stating that it was meant to 'complement their Watson AI platform'¹⁴. However, a GitHub repository suggests the use of

¹⁴ IBM Cloud Annotations JSON - The native JSON format of IBM's open-source Cloud Annotations labelling tool. <https://roboflow.com/formats/cloud-annotations-json>

'TensorFlow Object Detection API'¹⁵. I used this tool between 2019-2020 and the last commit shown in GitHub was on 21st April 2020.

4.4.3. Procedure

The planned approach based on the following phases:

1. Identification of the most common UI components used in Consumer IoT apps.

For that I used the IBM Cloud Annotations tool trained on a sample of the Rico dataset of Android apps UI screens. To achieve object detection¹⁶, the tool needed to be trained on at least 200 samples.

In this case the object being detected was the Floating Action Button (FAB) – a then very distinct UI component which was part of the Material Design library. By identifying UI screens that contained FABs, I was also identifying UI screens utilising Google's UI Design Language Material Design, which have been found to be linked to positive app ratings (Doosti et al., 2018). From that initial shortlist of UI screens using Material Design, I needed to filter even further to select those belonging to Consumer IoT apps, to only then identify the UI patterns in these screens.

2. Investigate the existence of the machine learning enhanced version of the same patterns.

Once I had the most common UI patterns used in Consumer IoT apps, the next step in the methodology would be to investigate the existence of the machine learning enhanced version of the same patterns being used in the real-world. These would have relationships to each other, potentially leading to the identification of a UI pattern language that encapsulates the use of machine learning.

A pattern language is collection of patterns organised in a coherent and interconnected way, addressing problems of a given field, which in this case is the field of UX Design of Consumer IoT apps.

¹⁵ <https://github.com/cloud-annotations/cloud-annotations> and <https://github.com/cloud-annotations/cloud-annotations/issues/49>

¹⁶ <https://developer.ibm.com/blogs/ibm-cloud-annotations-tool-eases-the-process-of-ai-data-labeling/>

Understanding the focus on the Floating Action Button (FAB)

Focusing on the Floating Action Button (FAB) was a strategic decision due to its distinctive appearance. The FAB is easily recognisable by its circular shape and its positioning visually 'above' the user interface – similar to the effect of using CSS's z-index property – often representing the primary action within a UI screen, which makes it particularly well-suited for object detection tasks. Additionally, its widespread use across many applications highlights its value as an important component for search.

Automated Identification of screens containing FABs

I developed a Python script to automate the identification of screens likely containing FABs:

1. Process Overview:
 - **Step 1:** Created a 'fab-screens' folder to store identified images.
 - **Step 2:** Parsed each JSON file recursively to search for nodes where the 'class' or 'resource-id' attributes contained the string 'fab'.
 - **Step 3:** If a match was found, copied the corresponding image file into the 'fab-screens' folder.
2. Implementation Details:
 - **JSON Parsing:** Utilised Python's json module to load and traverse the JSON structures.
 - **Recursive Search:** Implemented a function to recursively explore all nodes and sub-nodes in the view hierarchy, searching for the string 'fab'.
 - **File Handling:** Used the OS module to manage file paths and copying operations.
3. Outcome:
 - This approach significantly increased the proportion of images containing FABs shortlisted from the dataset, enhancing the efficiency of the annotation process.

Manual Annotation and Model Training

I uploaded the optimised subset of the dataset to the IBM Cloud Annotations and manually annotated the Floating Action Button (FAB) in each image by drawing bounding boxes around the FAB and labelling them accordingly. Additionally, potential false positives, such as circular buttons that were not FABs, were labelled as 'Non-FAB' to enhance the model's ability to accurately distinguish between actual FABs and similar-looking elements.

1. **The setup:** I utilised IBM Cloud Annotations for its capabilities in combining semi-automated annotations with model training in the cloud. This was

particularly useful given the limitations in computational resources available for free at the time, in 2019.




Figure 4 - Initial setup screen for IBM Cloud Annotations.

2. **Initial Labelling:** I began by annotating 20 images to test the tool and train an initial model. This was done by drawing a boundary around to objects in the image to be detected (in this case, FABs). As expected, the model performed poorly at identifying FABs due to the small sample size, yielding many false positives.




Figure 5 - Examples of Floating Action Button (FAB) selection for annotation.




Figure 6 - Due to the small amount of training data, the model did not perform as expected and led to false positives.

3. Identifying relevant images to better train the model:

- **First Approach:** I uploaded a random sample of 500 images but found no FABs to label. This suggested that FABs were relatively rare in random samples.
- **Second Approach:** I uploaded over 20,000 random images. After two weeks of manual labelling, only 40 images containing FABs were identified. The low yield confirmed the inefficiency of random sampling for this purpose.
- **Third Approach:** Optimised Data Collection via Metadata Parsing
In the Rico dataset each app screen file (JPEG or PNG) is accompanied by a heavily nested JSON file containing a hierarchical view of the UI components' coded structure (many of them are classes probably named by the developers). These files have consistent file names e.g., 64785.json matches 64785.jpg. I performed a few text searches in some of these JSON files and confirmed that some elements were named using conventions such as 'com.example.fab' or 'com.example_fab', where very often these elements represented a FAB button in the UI.

I considered that developers might use consistent naming conventions for FABs in the view hierarchy, such as including the string 'fab' in the class names, so I used the Python script described earlier to identify the screens that were most likely to contain FABs.

By recursively searching through the JSON files for nodes which had 'class' or 'resource-id' attributes with values containing the string 'fab', I was able to speed up the identification of relevant UI screens that contained the FAB component. These screens containing FABs were crucial for training the model to conduct automated labelling and therefore identifying the remaining UI screens that used Material Design.

A this point it is important to reiterate the purpose of this endeavour of identifying Consumer IoT app screens amongst the Rico Dataset's over 66,000 UI screens. These screens were from over 9,300 apps in 27 different categories, however there was no clear category for Consumer IoT, therefore this method was my attempt to identify apps matching this specific criterion. Figure 7 shows the distribution of the apps across all categories.




Figure 7 - Rico Android apps numbers across 27 categories¹⁷.

The plan was to identify apps using Material Design (MD), given previous evidence of them receiving better user ratings versus those Android apps not using MD (Doosti et al., 2018). From that pool of MD apps, I would then filter the screens of Consumer IoT apps and examine the UI patterns present in those screens.

After adopting the python script approach, I manually checked a small random sample of the files, and I noticed that it was far easier to find a screen using a FABs (meaning that the screen would be using Material Design). Unsurprisingly, there were false positives due to some JSON files having words like 'FABulous' for example somewhere in their structure. Still this approach proved to deliver a significantly better performance than the random selection of 20,000 images previously used.

The python script ran for approximately 5 hours scanning a folder with 66,000 images, and it identified just over 2,100 relevant images, which I uploaded to the IBM Cloud Annotations tool to once again, train the model by providing manual annotations and test for identification of relevant Consumer IoT app screens.

¹⁷ Graph source <http://www.interactionmining.org/rico.html>

I was then able to successfully label 279 Floating Action Buttons (FAB) and 21 'Not-FAB' items (which UI items likely be identified as false positives) in a space of 2 ½ hours. I also began to add the label 'nothing' to some screens, so these images were excluded from the non-labelled backlog (see Figure 8).




Figure 8 - Close up on IBM Cloud Annotations stats bar showing the thumbnails of screens that neither had FABs or Fab-like UI elements (not-FAB). In total, 279 FABs were annotated from a sample of 2,158 UI screens.

Typical examples of the UI screens containing FABs can be seen in Figure 9. These particular images were part of the selection of 279 images which had the FAB elements highlighted and labelled.




Figure 9 - Examples of UI Screens with a floating action button (FAB) located on the bottom right hand corner.

As a reference, using the previous approach of uploading a random sample of images, it took 14 hours to identify and label 40 FABs. The new approach, using the Python script to shortlist the screens with higher probability of having FABs proved to be over 38x more efficient.

Once the IBM Cloud Annotation model was trained, I fed into it images which did not belong to the training set to test the model's accuracy. The results showed that it still needed further tuning. Unfortunately, the tool did not allow access to its CNN's architecture to investigate any potential adjustments to parameters and potentially help with the model's performance issue.

Nearly three weeks were spent on this approach, where I manually labelled circa 2,000 images across different approaches expecting to lead to the identification of relevant Consumer IoT UI screens that used material design.

The final results for this experiment can be seen in the Table 2 below, where it summarises the number of screens identified to be using the Google Material Design (MD) language but most importantly, the number of Consumer IoT app screens encountered. Approach 1 involved uploading a random sample of 500 images to IBM Cloud Annotation to label FABs for training. Approach 2 was similar but uploading a much larger random sample of 20,000 images. Finally, approach 3 used a Python script to optimise the selection of images, resulting in 2158 images being uploaded to go through the same process.

	UI Screens using MD identified	Screens of Consumer IoT Apps using MD	Screens of Consumer IoT Apps (including those not using MD)
1 st Approach	0	0	0
2 nd Approach	40	0	0
3 rd Approach	279	0	0

Table 2 - Results of attempts to use a computer-vision approach (CNN) for the automatic identification of Consumer-IoT UI screens based on Material Design.

4.4.4. Results for The Automated Labelling Approach

This 'Automated Labelling' experiment did not produce the expected results – Despite successfully identifying hundreds of user interfaces using Material Design, there were no IoT (or Consumer IoT) apps amongst them – which was the purpose of this approach.

Considering that the ultimate goal was to identify common UI elements used in a Consumer IoT context to be augmented with ML, this method failed as a critical step towards this goal. This led me to pivot my approach to the one-to-one pattern mining study involving UX professionals described in chapter 5.

Initially, the one-to-one pattern mining workshops were conceived still to identify UI-patterns used in Consumer IoT which could be suitable for augmentation with machine learning. However, as I detail in Section 5.2.5, after the pilot study I shifted

the focus of the investigation from a UI patterns perspective towards a UX patterns one instead.

So crucially, although the automated labelling activity did not yield the results I anticipated, the fact that I had to shift towards a more 'analogue' and traditional manner of pattern mining led me to interact with and observe a UX professional during the study tasks.

That face-to-face contact enabled me to observe that UX designers would typically work with low fidelity user flows at that early, conceptual stage of the design process – even when UI patterns are available, and they were instructed to use them. That suggests that in the designer's mind, the concerns around the choice of components to assemble a UI happen much later, once the user flows and interaction problems have been addressed.

Limitations of the Automated Labelling UI Screens

General Limitations

Resource constraints had a significant impact on the scale of the experiments. Limited computational resources and time constraints influenced the approaches used. The reliance on publicly accessible datasets, which may not be precisely aligned with the unique context of Consumer IoT apps, presented issues, and the limited availability of specialist datasets restricted the breadth of study. Furthermore, since the experiments were carried out between 2018 and 2019, advances in machine learning and data annotation tools have developed, potentially now offering more efficient and effective methods.

Specific Limitations

1. **Dataset Bias:** The Rico dataset may not represent the latest design trends or the specific characteristics of Consumer IoT apps. Design practices evolve rapidly, and datasets can become outdated.
2. **Manual Annotation Challenges:** The quality and consistency of manual annotations can vary, impacting model training. Even though the purpose was to achieve automated annotation, the initial phase of manual annotation could be subject to annotation fatigue and therefore to errors.
3. **Generalisation of Models:** The trained models may not perform well on UI designs that differ significantly from those in the training dataset. Custom designs and stylistic variations can present challenges.
4. **Platform Limitation:** The focus on Android apps excludes UI components and therefore patterns prevalent in iOS or cross-platform applications. This may limit the applicability of findings.

5. **Correlation vs Causation:** Identifying the presence certain UI patterns does not necessarily reveal their impact on user experience. User satisfaction depends on various factors – if identified, these patterns would still need to go through a rigorous process to assess whether they were good patterns.

Methodological Constraints

1. **Technical Focus Over Design Principles:** Emphasis on technical implementation of the experiment could have overshadowed considerations of design principles and user needs, i.e. the UI problems these potential patterns were addressing. UX design is inherently user centred.

4.4.5. Relevance to the Central Argument

The results offer important insights regarding Sub-Proposition 1.1, which explored whether the Rico dataset, combined with computer vision, could serve as a viable approach for revealing Consumer IoT apps for further pattern analysis. Despite the initial success in identifying hundreds of Material Design screens, the approach did not help me locate actual Consumer IoT apps, indicating that an automated method alone was insufficient for streamlining the discovery of UI patterns in this domain. Rather than disproving RQ1 (the possibility of identifying an ML-enhanced pattern language for Consumer IoT), this unexpected outcome highlights a limitation in the dataset and method chosen.

Nonetheless, as will become clear in the following chapters, pivoting to more designer-led pattern mining ultimately confirmed the broader claim that a relevant pattern language can be identified, however the automated approach was unsuccessful for that, reaffirming how human insight (a key aspect the Pattern Theory) was integral for discovering truly relevant examples for ML-based IoT UX solutions.

4.5. Discussion and Implications for this Research

The experiments presented in this chapter were instrumental in addressing two exploratory sub-propositions linked to the overarching research questions. In the first experiment (addressing Sub-Proposition 1.1), I applied computer vision techniques to the Rico dataset to identify Material Design UI screens. Although the method efficiently identified Material Design UI screens, it did not isolate Consumer IoT examples to serve as a stand-alone data-driven approach. This outcome exposes the limitations of relying solely on automated techniques for discovering relevant UI patterns in a complex, multi-category dataset.

In the second experiment (addressing Sub-Proposition 2.1), I conducted a **user identification** task using gait recognition, with the aim of exposing a realistic ML

workflow to document. While the KNN classifier achieved 95% accuracy – suggesting the technical feasibility of such an approach – the fine-grained details of algorithmic configuration (such as kernel selection and parameter tuning) was not the type of information that could be understood by non-ML-expert UX designers. Further, examining the results of this experiment presented form of confusion matrices, reinforced the notion of this level of information to be outside the average UX professional skillset – to date, knowledge of these concepts is not typically listed as a requirement in the community of practice ¹⁸.

In fact, discussions with practitioners later when conducting the validation study further confirmed this. There was a clear demand for design guidance that focuses on handling errors gracefully, ensuring privacy, and clearly explaining the learning process rather than details on technical implementation. Although one participant did mention that it could be useful to have more technical information for some parts of the interaction (which could help with technical implementation). However, as explained previously, I kept the scope of these patterns strictly associated to UX practitioners need to avoid adoption challenges.

As discussed in the previous section, these findings indicate that the expectation expressed in Sub-Proposition 2.1 does not hold in practice and, more broadly, examining the results led me to advocate for a higher-level and principle-based approach in the pattern documentation. In doing so, the experiments collectively highlight that while technical ML methods are useful for demonstrating feasibility, the incorporation of human insight (through designer-led methods) is useful for uncovering truly relevant UX patterns for Consumer IoT. This directly addresses **Research Question 1** by leading to an approach that discovers relevant patterns that capture the nuanced requirements of user-centred design.

The limitations observed in the automated, data-driven experiments informed the subsequent shift towards a more qualitative, designer-centric pattern mining process, as detailed in Chapter 5. In essence, the ML experiments have provided both empirical evidence and practical guidance that underpin the later stages of this research, reinforcing the overall aim of bridging advanced ML capabilities with accessible UX design practice.

¹⁸ NN/g Skill Mapping template, a tool used by Design Operations to map the skillset of design teams does not have any reference to machine learning <https://www.nngroup.com/articles/skill-mapping/>

5. Pattern Mining

5.1. Introduction¹⁹

In order to achieve the aim of discovering relevant user interface (UI) patterns for ML enhancement, I devised an effective pattern mining technique. Pattern mining is the first phase in the pattern writing process, and it involves identifying and extracting important information known as "seeds" that serve as the foundation for producing formal pattern documentation (Alexander, 1979; Meszaros and Doble, 1997). These seeds contain early ideas or core principles taken from best practices in a particular field, in this case, user experience (UX) design.

Pattern mining is not an activity that one conduct in isolation; it largely relies on the collective expertise of a community of practice (Gamma et al., 1995, pp. 11–14). This communal approach ensures that the patterns developed are robust, widely

¹⁹ Parts of this chapter were published in my 2019 paper, 'A pattern approach for identification of opportunities for personalisation and automation of user interactions for the IoT' <https://doi.org/10.1145/3361149.3361157>

applicable, and grounded in real-world experiences. To effectively mine patterns in UX design, especially in the context of consumer Internet of Things (IoT) applications enhanced by ML, I recognised the necessity of involving UX professionals actively practising in the industry. To that end, pattern mining can take place in many ways, for example interviews or workshops (Sasabe et al., 2016) with practitioners.

In this chapter I cover the Pattern Mining activities, the first step in the pattern authoring lifecycle. The overarching goal of this research project is to document a pattern language that enables UX professionals to recommend UX solutions that integrate machine learning (ML) into Consumer IoT applications without requiring specialist technical knowledge. **Research Question 1 (RQ1)** asks whether such a pattern language can, in fact, be identified.

Chapter 4 documented an initial, data-driven search for potential patterns in a large-scale UI dataset. Such approach was not successful in identifying relevant patterns. This chapter therefore addresses the resulting alternative: a designer-centred study intended to tackle **RQ1** from a different angle.

This chapter is guided by the exploratory Sub-Proposition 1.2 (SP 1.2):

A designer-centric pattern mining process, grounded in direct engagement with UX professionals, may yield a more contextually relevant and practically applicable set of ML-driven UX patterns than a purely automated approach relying on large-scale UI repositories.

To explore this sub-proposition, I ran a series of designer-led workshops described in detail in the following sections, and applied rigorous qualitative analysis to the resulting artefacts. The results presented in this chapter capture insight from the UX community and reinforce the important role of human insight in bridging the gap between technical potential and user-centred design solutions.

5.2. Methodology

5.2.1. Ethical Considerations

This research study followed the institutional ethics guidelines from the University of Hertfordshire particularly regarding participant privacy and informed consent. Participants in the pattern mining workshops were informed about the study aims, format, and their right to withdraw at any time. Each participant provided written or electronic consent prior to data collection.

Any recordings or transcripts used for analysis were stored on password-protected drives, accessible only to the researcher and supervision team. Similarly, any

references to identifiable user data were de-identified before being used in previous publications and in this document.

Compliance and Approval

Ethical clearance was obtained from the University of Hertfordshire Social Sciences, Arts and Humanities Ethics Committee with Delegated Authority. The UH protocol number is aCTA/PGR/UH/03834. The data collection involved minimal risk and no vulnerable participants.

5.2.2. Approach Design and Rationale

To align with the foundations set by Akado et al. (2015) and Iba & Yoder (2014), in this research project I adopt a workshop-based pattern mining strategy that:

- Engages UX Professionals: Collects expert knowledge on recurring design problems and potential ML-driven solutions.
- Follows a Delphi-Like Protocol: Iteratively refines patterns through participant feedback, akin to a simplified Delphi approach (Chuenjitwongsa, 2017).
- Validates Draft Patterns: Submits pattern drafts to a peer review process (in this case European Conference on Pattern Languages of Program - EuroPLoP) for rigorous critique and refinement.

Through this process, I tailored established pattern-mining frameworks (Akado et al., 2015; Hentrich et al., 2015; Iba and Yoder, 2014; Sasabe et al., 2016) to suit common ways UX professionals operate, devising an approach to capture both the practical realities of ML-integrated design and the theoretical rigour of pattern scholarship in these patterns.

Further, the pattern mining approach needed to address some critical challenges, particularly the potential lack of expertise in ML and IoT among UX professionals. Table 3 summarises these challenges and the strategies employed to overcome them.

By systematically mining patterns, I aimed to gather evidence that a cohesive set of ML-enabled UX solutions can be identified (**RQ1, GP1**).

Challenge	Solution
1. Collecting best practices from the community of practice	I ensured that all participants were UX professionals currently practising in the industry. Their real-world experiences across diverse UX contexts provided valuable insights. While their expertise may not be in ML or IoT, their collective knowledge in design best practices was instrumental in identifying core interaction patterns based on established design principles.
2. Knowledge of consumer IoT user interaction (from a UX perspective)	Focused on smart devices with familiar, non-connected versions (e.g., smart windows) to bridge the knowledge gap, given that this would enable participants to both think as a designer and empathise as a user. This approach allowed participants to draw on their expertise in traditional products and identify potential interaction challenges arising from the introduction of IoT. This method facilitated the extraction of valuable seeds for consumer IoT interaction patterns.
3. Knowledge of appropriate use of ML to improve user interaction	I utilised the "What Would a Human Do" (WWHD) exercise to explore the potential for enhancing user interactions with ML. By comparing the human-human interactions to their equivalent human-computer interactions I was able to identify key stages where ML could be effectively applied. The seeds gathered through the WWHD exercise formed a strong basis for understanding ways ML could improve user interaction.

Table 3 - Challenges and their solutions for effective pattern mining

The Resulting Approach

To address the challenges and implement the solutions listed in Table 3 in the previous section, I devised a detailed, multi-stage process to collect and analyse first-degree data to answer **RQ1 – on whether the novel pattern language could be identified**. The approach focused on reaching the most suitable and relevant participants, using classic digital communication targeting strategies to maximise the response in the recruitment phase. It also involved a semi-structured workshop with each participant where key data was generated. Finally, a meticulous data analysis process was devised to ensure the findings were reliable.

The process comprised of four stages:

1. Participant Recruitment
2. Surveys
3. One-to-One pattern mining workshops
4. Data Analysis

5.2.3. Participants

This section addresses **Stage 1 - Participant Recruitment**. The recruitment plans targeted UX design practitioners currently active in the industry. Figure 6 shows the original recruitment process map. Recognising the potential difficulty in reaching a

large sample of UX practitioners, I adopted a "test and learn" approach. Initially, I utilised lower-quality channels – platforms less frequented by the target audience – to test the effectiveness of advertisements in attracting participants. This strategy aimed to optimise the messaging and creative aspects of the ads while avoiding any risks of ad fatigue among the prime audience (Resnick and Albert, 2014).

I experimented with text alternatives that focused on IoT, ML, or both. The language and visuals were tailored to appeal to UX designers, using terminology like "human-centric design" and themes associated with usability and accessibility, key concerns in the UX community.

See Appendix C for examples of the creative executions showing variations of text with focus on IoT, Machine Learning and both IoT and Machine Learning that I tested.

In addition, I conducted an A/B test using two distinct creative approaches (Appendix D). One ad showed an elderly man, representing a group often affected by poor usability and accessibility in design. The other showed a baby's hand touching an elderly hand, with the headline "Add the human touch," suggesting the value of human centric interactions - a pertinent topic in discussions about automation and technology in UX design.

The A/B test was conducted using Facebook's testing engine to determine the most effective ad. Upon analysing the results, the winning creative was utilised in higher-quality channels, such as LinkedIn's special interest groups. This strategic deployment ensured the ads reached the right niche of professionals.

As a testament to the quality and potential of the channels I chose for reaching experienced UX design professionals, renowned design pioneer Don Norman commented on one of my LinkedIn recruitment posts (Figure 11). He questioned the use of the term "Designing for humans," expressing a preference for "Designing for people." Although Norman himself helped to popularise the term "human-centred design"²⁰, his evolving perspective now favours "humanity-centred design"²¹ " emphasising a broader consideration of global impact (Norman, 2023). For the purposes of this research, I maintained my original terminology, as it remained relevant to the context and objectives of the study.

²⁰ <https://www.sandiegoreader.com/news/1994/dec/01/cover-when-you-dont-know-how-to-turn-on-your-radi/>

²¹ <https://jnd.org/humanity-centered-versus-human-centered-design/>




Figure 10 - The participant recruitment process initially leveraged social media, networking events, and industry contacts. Higher-quality channels were introduced in later stages, informed by insights gained from early recruitment efforts. Some channels shown on the map (e.g., Twitter) were expected to perform less effectively, whereas others (e.g., LinkedIn and in-person contact) were anticipated to have a greater likelihood of attracting the right participants.

Rubem Barbosa-Hughes posted in UX Professionals

Rubem Barbosa-Hughes
UX strategist | Doctoral researcher: Interaction Patterns using Machine Learning...
2yr · 🌐

Enjoy designing for humans? Visit uxdesigners.co.uk and let's have a chat. The aim of this project is to promote a more human-centric approach to the use of machine learning in the UI design of the Internet of Things. You can help by completing the survey but even better if you could participate in a short interview (either via skype or I can come to your workplace).

Please spread the word!
#humans #ux #artificialintelligence #uxdesigners #machinelearning #design #designthinking #iot #uxdesign



1 · 2 comments

👍 Like 💬 Comment

 Add a comment... 🧐 🖼️

Most relevant ▾

 **Don Norman** · 1st
Director, The University of California, Design Lab
2y ...

Actually, I don't believe in designing for humans: I design for people. Eh?

Curious · 🗨️ 1 | Reply · 1 Reply

 **Rubem Barbosa-Hughes** Author
UX strategist | Doctoral researcher: Interaction Patterns using Machine Learning...
2y (edited) ...

Hi Don Norman, now that the study has been and gone, I can be more open regarding the reason for insisting in using 'humans': it was simply due to the easier association with term human-centred design, which is effortlessly accessible in most designers minds :)

Like | Reply

Figure 11 - The recruitment strategy successfully reached a high-quality audience. In the example above, Donald Norman questions my use of the term 'humans,' as he now appears to prefer 'people' instead.

Prospective participants were directed to a dedicated landing page (Appendix E), which provided detailed information about the study and access to a screening survey.

Adopting guidelines from human-computer interaction research, smaller participant pools can yield in-depth insights when the aim is to capture qualitative nuances (Nielsen, 2000). Given the exploratory nature of pattern mining, achieving more than 20 workshops allowed for adequate emergence of themes, balancing breadth of perspectives whilst maintaining logistical feasibility. Similarly, for the Validation Study (discussed in chapter 7), focusing on five experienced UX designers permitted a close examination on whether the documented patterns supported realistic design tasks.

Stage 2 - Surveys

The survey was designed to meet ethical requirements such as obtaining explicit consent, and to collect initial data pertinent to the research. It was divided into three logical sections: **Screening, Initial Survey and finally Workshop Arrangements**, where participants could indicate their preferences to attend the one-to-one workshops.

I implemented a screening process to ensure that only individuals who identified themselves as professional UX practitioners could participate in the study. This step was essential to maintain the focus on the relevant community of practice and to ensure the validity of the findings.

I then conducted an initial online survey with a total of 28 participants (P1 to P28) who are professionals in the UX design field. The questionnaire collected data on demographics, work experience, involvement in IoT and machine learning projects, and the use of design resources such as design systems and pattern libraries. I applied a logic flow within the survey to enhance its relevance (Appendix I); certain questions were skipped based on prior responses to avoid irrelevant or redundant queries. This approach allowed me to gather information regarding participants' potential exposure to machine learning, the Internet of Things, and pattern theory, while also providing them with the opportunity to consent to participate in a subsequent workshop.

5.2.4. Apparatus and Materials

I used the same standard MacBook Air using in previous experiments: 13-inch screen, Mid 2012 version, 256GB flash storage, processor 1.8GHz dual-core Intel Core i5, 4GB of 1600MHz DDR3L onboard memory with Intel HD Graphics 4000.

Participants were provided with a list of the following six IoT devices on A4 sheet:

- Smart curtain track
- Smart kettle
- Smart fridge
- Smart bed
- Smart window
- Smart lock

The selection aimed to include devices with which participants were likely familiar in their non-connected forms, thereby minimising the cognitive load and allowing focus on interaction design.

They were provided with a collection of monochromatic, wireframed versions of Material Design UI components (Appendix F).

I also used MaxQDA 2018 and Microsoft Excel (Mac version) for data analysis.

Interviews took place both in person and remotely. For remote sessions zoom was used and for in person sessions, the participants were also given the option to provide their responses using A4 sheets provided.

5.2.5. Procedure

Stage 3 - The Pattern Mining Workshops

Following the survey, I made arrangements for the workshops based on participant's preferences. This step ensured that the workshop sessions could be organised effectively, accommodating the participants' availability. Moreover, this separation enabled individuals to contribute with the survey without necessarily having to participate in the workshops.

To assist pattern mining applying the solutions proposed in Table 3, I used focused interviews as one-on-one workshops, both in person and remotely. This methodology is consistent with qualitative research approaches used for exploratory studies in UX design (Robson, 2011). Furthermore, at the conclusion of each session, participants completed a brief, semantic differential questionnaire (Osgood et al., 1957) to assess their experience in the study and their views of the relevance of this research to their practice.

The workshops were designed as an evolution of the initial UI pattern mining method described in Section 4.4, shifting focus from UI components to broader UX flows encompassing both interaction and UI aspects. This shift was informed by insights gained during the pilot study, recognising that UX patterns offer a more holistic understanding of user interactions in IoT contexts.

Due to health and safety requirements outlined in the ethics application, workshops were conducted either at the participants' workplaces or online, ensuring compliance with institutional guidelines. All artefacts were provided either in physical or digital form, as required.

Because RQ1 focuses on whether an ML-driven UX pattern language can be identified for Consumer IoT, this phase of pattern mining is critical for gathering empirical evidence to explore Guiding Proposition 1 (GP1).

The workshops were divided into two parts.

Part 1: Each participant received A4 sheets containing the list of Smart devices and the set of UI components from the Material Design (MD) library, presented as low-fidelity, monochromatic wireframes to reduce aesthetic bias. The components were not explicitly identified as part of the MD library to reduce any potential bias. Initially the plan was to restrict the designers to use only material design UI components provided in the sheets, however, they were free to employ other UI patterns or suggest modifications, provided they justified their choices.

I asked the participants to choose one device from the list provided and conceptualise three interactions that would occur through a mobile application

This approach aimed to observe any IoT-specific adaptations the designers might propose to the UI patterns whilst avoiding any perceived researcher preferences.

Nonetheless, the insights gained during pilot study demonstrated that the requirement of focusing on UI aspects was too limited and to benefit from ML, the design solutions would often need to consider entire interactions, shifting the focus of this research towards the broader UX flows.

Part 2: The "What Would a Human Do" (WWHD) thought experiment
Participants were then introduced to the WWHD thought experiment, a method designed to uncover natural human interactions that I used to help the identification of ML-enhanced UX patterns. I asked them to describe how they would interact with a human being to achieve the same goals as in the digital interactions they just conceived.

To facilitate this exercise, I encouraged the participants to disregard feasibility constraints and focus solely on detailing the theoretical human-human interaction. By comparing these narratives with their digital interaction equivalents, I was able to identify key stages where ML could enhance human-computer interaction. Figure 12 shows the high-level plan for the original workshop.

Objectives:

- Identify IoT-specific patterns that may emerge due to designer's previous experience.
- Identify pattern enhancements based on the WWHD (What Would Human Do) approach (Yampolskiy, 2013). These enhancements will be considered as UI problems for machine learning (ML) and ultimately lead to the initial drafts of ML enhanced patterns.


Interviews may take place at their homes, workplace or other suitable premises.

The interviews will be video and audio recorded.

The setup should take between 10-15 minutes

White boards, sketch pads post it notes will be provided if necessary.

The first part of the study should take a maximum of 25 minutes



- The participants will be asked to choose an IoT device from the list.
- They will also be given a set of A4 sheets with wireframes of Material Design (MD) UI Patterns. These will be monochromatic and not clearly labelled as MD.
- The participant's will then be asked to describe or sketch how he or she would envisage 3 interactions for the chosen device, using the patterns provided.
- They will be told that they are allowed to make these patterns function any way they see fit, as long as all non-standard behaviours are fully described.

IoT devices for the use case exercise

Smart curtain track
Smart kettle
Smart fridge
Smart bed
Smart window
Smart lock

The second part of the study should take a maximum of 25 minutes

Finally, the participants will then be asked to disregard their knowledge of current technical constraints and industrial feasibility and ideate on improvements of those three interactions using the WWHD (What Would Human Do) approach (Yampolskiy, 2013).

This will be the most critical task of the study. The previous activity would have served to set the participants' focus on the problems they are solving using their experience.

By asking them to ideate beyond the known constraints, I expect to uncover recurring enhancement suggestions which could be abstracted and turned into reusable pre-patterns.

This approach is inspired by the pattern 'Mining' process (Sasabe et al., 2016), and it will aid the documentation of the knowledge which will be captured in the IoT UI patterns enhanced by machine Learning.



References:

Takashi Iba, Pattern Mining Workshop (PURPLSOC2017).

Sasabe, A., Kaneko, T., Takahashi, K., Iba, T., n.d. A Search for the Seeds of Patterns 16.

Yampolskiy, R.V., 2013. What to Do with the Singularity Paradox?, in: Müller, V.C. (Ed.), Philosophy and Theory of Artificial Intelligence. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 397-413. https://doi.org/10.1007/978-3-642-31674-6_30

Figure 12 - The original workshop structure document. Initially focused on UI patterns, it evolved to more effectively identify UX patterns. After the pilot study, participants were no longer required to use the MD UI Patterns, though it remained available as an optional resource.

5.2.6. Data Collection and Analysis

Stage 4 - Data Analysis

The pattern mining study generated a comprehensive corpus of qualitative data, supplemented by insights from the online survey. This data was critical in uncovering UX patterns for Consumer IoT applications enhanced by machine learning (ML). To analyse this data, I employed grounded theory methodologies, following systematic procedures of open coding, axial coding, and selective coding (Corbin and Strauss, 2014). This approach facilitated the emergence of core themes, categories and seeds, crucial for my goal of developing the UX design patterns informed by human-human interactions, in line with Alexandrian pattern theory (Alexander et al., 1977).

Grounded theory is a qualitative research methodology aimed at generating theories grounded in systematically gathered and analysed data. It is particularly effective for exploring complex phenomena where existing theories are insufficient (Bryant and Charmaz, 2010). In this study, the complex phenomena were the traditional UX patterns in Consumer IoT which could be enhanced with the integration of machine learning.

Grounded theory is a qualitative methodology concerned with generating theory from systematically gathered and analysed data, and is particularly appropriate where existing theoretical accounts are insufficiently developed to explain a complex phenomenon (Bryant and Charmaz, 2010). It was selected for this study in preference to thematic analysis because the analytic objective extended beyond the identification and description of recurrent themes. Although thematic analysis is well suited to organising qualitative material into descriptive categories (Robson, 2011), grounded theory in the Straussian tradition offers a more explicitly theory-building analytic logic. Specifically, open coding supports the identification of concepts and categories within the data; axial coding develops those categories and specifies their relationships to subcategories; and selective coding integrates and refines them around a central explanatory concept within a larger theoretical scheme (Strauss and Corbin, 1998, pp. 102–103, 125–126, 143, 161). This analytic progression aligns naturally with the pattern-mining process itself, which moved from individual interaction concepts, through interconnected pattern seeds, towards the basis of a coherent pattern language. The phenomenon under investigation in this case was the set of UX patterns in Consumer IoT that could be enhanced through the integration of machine learning.

Open coding is the initial phase where the collected data is examined for emerging themes, leading to identification, naming, and categorisation of such themes. In my

study, the open coding process involved an analysis of the transcribed workshop discussions, annotations, and narratives from the "What Would a Human Do?" (WWHD) thought experiment (still as part of the same workshops). Using MAXQDA software, I systematically coded the data, resulting in a comprehensive coding structure that included codes from all stages of the coding process: open, axial, and selective.

Axial coding is the second phase of grounded theory analysis, where I reassembled the data by finding connections between the categories and subcategories created during open coding (Strauss and Corbin, 1998). This process involves linking categories to form a more coherent understanding of the data.

In this phase, I focused on exploring how different codes relate to each other by examining conditions, contexts, actions/interactions associated with each category. For example, I observed that codes like "Voice - NLP," "Context awareness," and "Proactivity" often co-occurred, indicating a relationship where voice interactions are enhanced by context-aware systems that proactively assist users. By identifying the codes interconnections, I developed broader categories that encapsulate the underlying patterns in the data. This led to the emergence of categories such as:

- **Automation:** Encompassing system functionalities that perform tasks autonomously, reduce user effort, and anticipate user needs.
- **Personalisation:** Involving the adaptation of system responses and functionalities to individual user preferences, behaviours, and contexts.
- **Control:** Relating to the user's ability to oversee, influence, and override system behaviours, addressing concerns around autonomy, privacy, and security.

I slightly adapted the selective coding phase, typical in grounded theory method, so that it served its original purpose for this research project: the identification of traditional UX patterns, in the context of Consumer Internet of Things which could be enhanced with machine learning. This central core category had already been established. The results of this study were not meant as the final research outcome, but instead they served as a rigorous UX pattern identification step, informed by practitioners in the industry. The outputs, which were pattern seeds, interactions or basic concept of patterns, paraphrased in a paragraph or a list of steps, would then go through another review process to assess whether they were early indications of relevant patterns, and only then would they be formally documented.

I conducted the analysis using MaxQDA, a specialised software designed for qualitative and mixed-method research (Kuckartz and Rädiker, 2019). MaxQDA was

selected due to its intuitive coding mechanisms, stable functionality, and compatibility with macOS, the operating system I utilise to analyse this study.

While exploring potential software options, I considered NVivo, another well-known qualitative data analysis tool. However, at the time of the study, NVivo's macOS version was reported to have limited functionality and produced errors when importing data²². To ensure an efficient analysis process, I opted for MaxQDA, which provided stable performance and a user-friendly interface to facilitate coding and data management.

In addition to MaxQDA, I employed Microsoft Excel to facilitate specific analytical tasks, particularly those involving tabular comparisons and data organisation. Excel's grid-based structure allowed for the visualisation of themes and categories, helping to identify seeds of patterns and filtering irrelevant information.

Comparison of Interaction Concepts in Excel

After completing the coding phases, I utilised Excel to compare individual interaction concepts extracted from the data. By organising the coded extracts in a tabular format, I could systematically analyse similarities and differences across various interactions. This comparison allowed me to identify recurring interaction patterns and potential areas where ML could enhance some of those traditional user interactions.

For example, interactions involving remote operation of devices were compared across different smart appliances, highlighting common user needs for convenience and control. Patterns emerged in how users interacted with smart kettles, fridges, and beds, revealing opportunities for ML to introduce predictive functionalities and personalised experiences. However, not all these early pattern seeds would end up being formally documented as other criteria would become relevant, e.g. not needing specialised hardware.

Paraphrasing and Organisation in Excel

In Excel, the data was distributed across columns capturing specific attributes of the interactions. It included:

- **Device:** The specific IoT device related to the interaction.
- **Interaction 1, 2, 3, etc.:** Verbatim or paraphrase of the initial interaction concepts generated by participants.
- **Interaction 1 Desc, 2 Desc, etc.:** More detailed descriptions of each interaction.

²² <https://www.researchgate.net/post/Does-anyone-use-qualitative-analysis-software-on-a-Mac>

- Interaction 1 - Documented, Interaction 2 - Documented etc.: Interaction flows step-by-step, numbered in sequence.
- **WWHD 1, 2, 3, etc.:** Notes from the "What Would a Human Do" (WWHD) approach, transcript or paraphrased natural human interactions corresponding to the concept.
- **Final Interaction 1, 2, etc.:** Refined versions of the interactions, incorporating the WWHD adapted into the interaction.
- **PATTERN 1, 2, 3, etc.:** Tentative pattern names or themes emerging from the interactions (seeds of patterns) Figure 13 shows an excerpt from the Excel file that illustrates the structure:

Device	Interaction 1	Interaction 1 Desc	Interaction 1 - documented	WWHD 1	Final interaction 1	PATTERN 1
Smart window	The user can calibrate the amount of light coming into the room through the smart window, adjusting the level of opacity or darkness via the app.	The smart window's app interface allows the user to adjust the level of light entering the room using a slider. The slider provides both visual and text input options for accessibility.	1. Open app. 2. Select window control. 3. Use slider to adjust light level. 4. Confirm desired opacity level.	The user would ask a person to adjust the blinds or curtains manually to control the light level in the room, specifying the desired amount of light.	1. User asks someone to adjust the blinds or uses app to control light level. 2. Adjusts opacity via slider or gives specific instructions to achieve desired lighting.	Light Calibration: Provide users with an adjustable slider to control light opacity and brightness, offering both visual and text input for accessibility.

Figure 13 - Example of an excel line organising the notes from interaction concept to initial pattern seed.

This structured organisation allowed for efficient sorting, filtering, and comparison of related interactions, enabling the identification of recurring themes and unique concepts. Given the semi-structured flow of the workshop, I used paraphrasing to distil the essence of each interaction. This setup facilitated the identification of pattern seeds which represented foundational ideas for potential UX patterns within consumer IoT applications.

Following that, each pattern seed was then considered based on:

- **Relevance:** The degree to which the concept fits the context of this research (Consumer-IoT interactions enhanced by ML) and is likely to arise in everyday use.
- **Usefulness:** The potential of the concept improving the user experience by reducing effort, error or frustration.
- **Machine Learning Integration:** Whether the concept leveraged machine learning to provide adaptive or intelligent behaviours.

These criteria were adopted because they map directly on to the goals of the study. **Relevance** tests whether the proposed interaction would matter to ordinary users in the Consumer IoT context, avoiding off topic seeds (e.g. general mobile-app patterns, industrial-IoT workflows). **Usefulness** assesses if the concept actually addresses a recognised pain-points in current Consumer-IoT UX, ensuring the solutions have practical value, rather than just being wishful thinking. Finally, **Machine Learning Integration** immediately excludes seeds that are unable to leverage ML to personalise or automate relevant aspects of the user interaction – which is central to the RQ1.

Using the evaluation criteria described above, I conducted a deeper review into the pattern seeds, leading to the identification of pattern candidates which featured the adaptive, intelligent behaviours enabled by machine learning.

I then conducted an additional review of the shortlisted candidate patterns to assess whether the potential patterns would be generalisable enough to be useful (e.g. not depending on specialised hardware for example) or if they were in fact executions of a pattern – a typical confusion in the field. E.g. a particular layout of a login screen could just be a variation of the **Sign in/New account** pattern (Duyne et al., 2007). Wider generalisability is critical to help with the longevity of the patterns, ensuring they are not reliant on niche hardware or software to exist. The Early distinction between Patterns and Executions ensures that each resulting documentation is at the right level of abstraction, maximising the reusability, i.e. a real pattern.

Final selection of patterns for documentation

After the layers of reviews described above, I shortlisted the candidate patterns for formal documentation using the approaches discussed in Section 2.3, ensuring that solutions, problems and contexts are well understood (Iba, 2014; Wellhausen and Fießer, 2012) and further reviews and validations through the PLoP conferences.

By systematically documenting and gaining formal approval of these candidate ML-enhanced UX patterns, I aimed to address **Guiding Proposition 1 (GP1)**, which anticipates that a cohesive pattern language encapsulating ML-driven interactions for Consumer IoT can be identified. By formally documenting each pattern in a standard format, GP1 could be explored through the creation of explicit artefacts before moving on to the validation study, which addresses **Guiding Proposition 2 (GP2)**.

Triangulation of Data Sources

To reinforce the credibility of my qualitative analysis, I applied triangulation across multiple data sources obtained from the pattern mining workshops. In addition to participants' verbal accounts captured in the workshop transcripts, I considered a diverse set of design artefacts, observational notes, and (where relevant) short survey responses. By examining this variety of materials, I aimed at conducting a more thorough analysis reducing the impacts of potential personal bias.

Triangulating evidence from multiple data streams improves the validity of qualitative findings by enabling verification of emergent themes from different sources, and therefore mitigating the risk of researcher bias (Denzin, 2017). Although not without limitations, given that triangulation may risk introducing

'discrepancies and disagreements between the different sources' (Robson, 2011 pp. 158), this approach was well suited for this use case as it provided an additional layer of objectivity when capturing the insights.




Figure 14 - Example of material produced by a designer, which was also considered in the analysis process.

Specifically, after open coding each phase, I consulted the annotations I had made during the workshops themselves. In some instances, I found that participants' spoken reflections aligned closely with the content of their design sketches (see Figure 14) or diagrams; in other instances, the designed artefacts revealed nuances that the participants did not explicitly voice. Observing such discrepancies allowed me to refine and capture the full breadth of insights.

This practice of cross-verification ensured that the final patterns emerged from a holistic reading of the data, reducing the likelihood that I overlooked subtler or conflicting details and ultimately helped to improve the rigour of the overall pattern mining process.

5.2.7. Limitations of the Methodology

While the methodology was rigorous, several limitations were acknowledged:

Subjectivity in Coding

Qualitative analysis is inherently interpretive, and coding decisions are influenced by the researcher's perspectives (Mays and Pope, 2000). In this study, I acknowledge that subjectivity may have influenced the coding process. Future research could benefit from incorporating strategies like collaborative coding to enhance the trustworthiness of the results.

Data Volume Management

The extensive volume of data posed challenges. High-frequency codes might overshadow less frequent but significant codes. I addressed this by carefully reviewing all codes, regardless of frequency, to ensure that less common but important themes were not overlooked. However, the potential for missing subtle insights remains a limitation.

Reliance on Participant Articulations

The depth of insights depends on participants' ability to articulate experiences. Factors like social desirability bias²³ may have influenced responses. Ensuring a more comfortable and open environment during data collection could help mitigate this issue, but limitations in data richness may still happen.

Technological Limitations

Software constraints may have affected the analysis depth. Advanced features like ML-assisted coding was not available at the time, which could have enhanced pattern identification and analysis efficiency.

5.3. Recruitment Results

The Participant Recruitment activity involved the use of social media posts and assessment of the performance of creative execution (i.e. images and copy assets used in these posts) in attracting participants. I conducted an A/B test to assist with the selection of the most effective combination of creative assets where both adverts ran at the same time, sharing budget and schedule, ensuring that any observed differences derived from the creative. Table 4 shows the A/B test results and Figure 15 shows the assets **tested** along with the respective results.

Variant A had the image featuring the baby's hands and Variant B had the profile of an elderly gentleman.

²³ <https://www.sciencedirect.com/topics/psychology/social-desirability-bias>

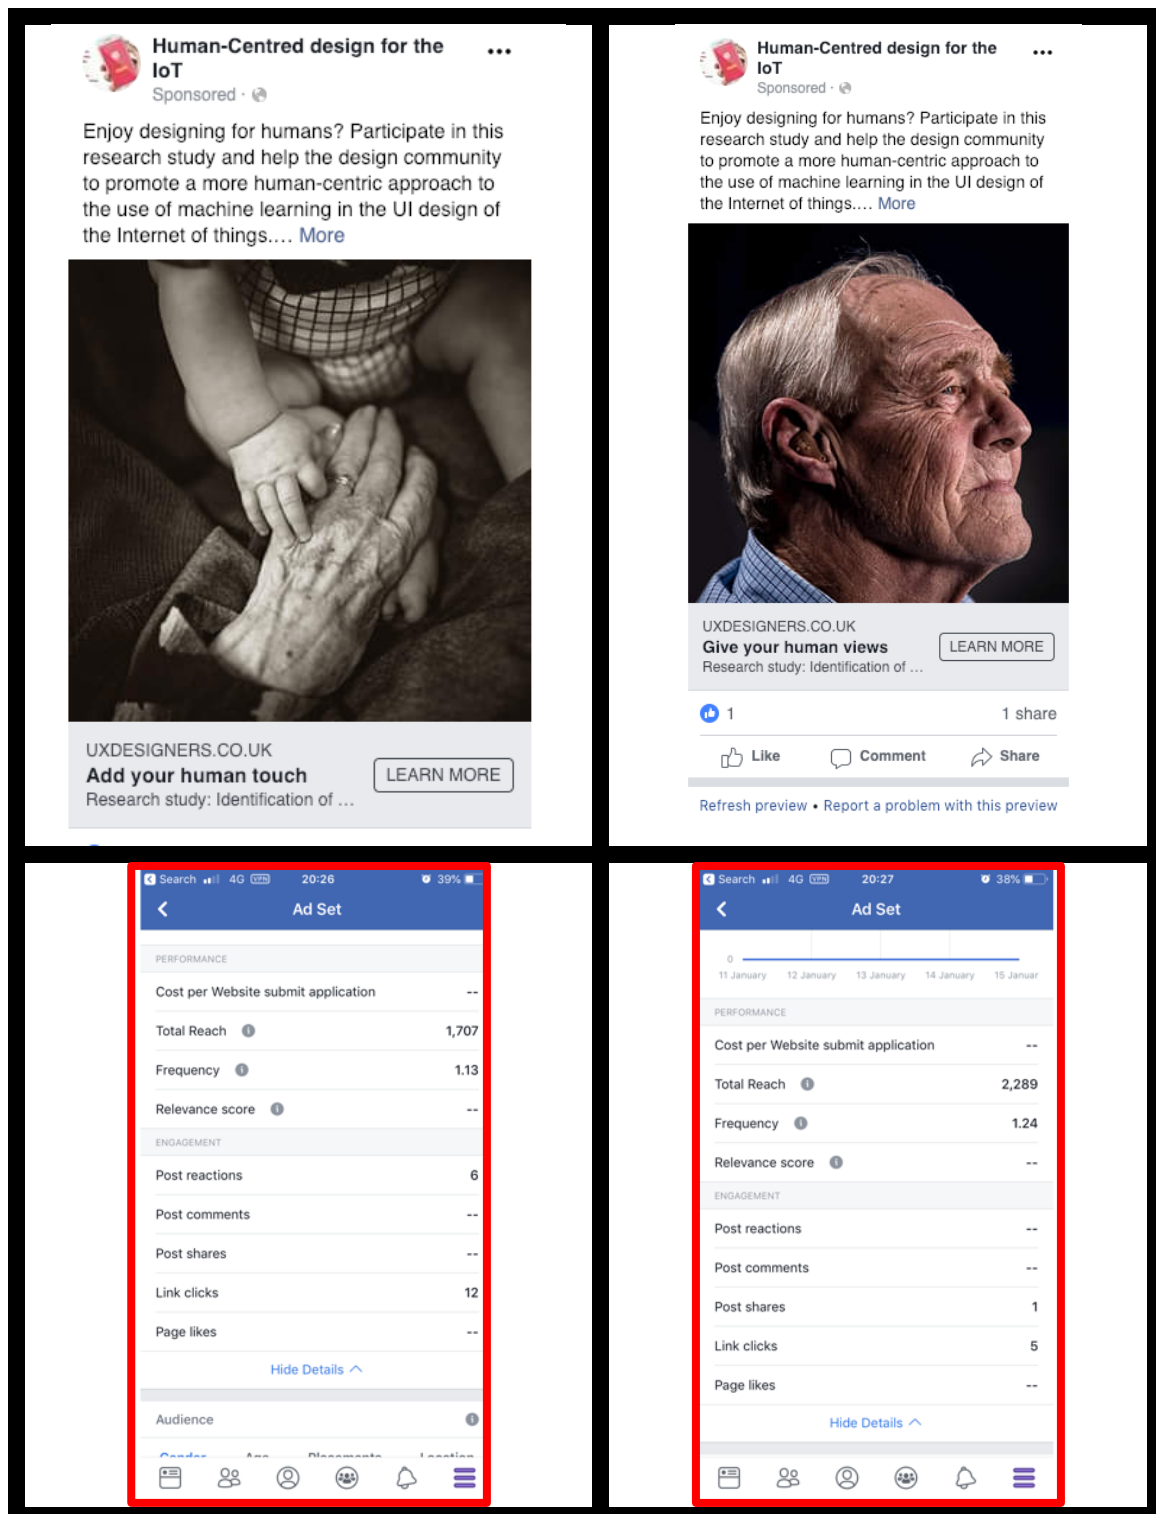


Figure 15 - A/B test results for two creative executions for the participant recruitment ad. The Ad featuring the baby's hand had greater engagement. The performance and engagement metrics for each variant are shown in the lower section.

Metric	Variant A	Variant B	Relative change (B vs A)
Reach (unique people)	1 707	2 289	+34 %
Mean frequency	1.13	1.24	+9.7 %
Link clicks	12	5	-58 %
Post reactions	6	0	-100 %
Post shares	0	1	+1
Click-through rate (CTR)	0.70 %	0.22 %	-68 %
Reaction rate	0.35 %	0.00 %	-100 %

Table 4 - Results of the A/B test assessing the creative assets for the participant recruitment posts.

Interpretation of the A/B test results

Greater exposure, lower engagement – Variant B achieved a 34% larger audience and a marginally higher exposure frequency, yet its engagement metrics were very poor. The CTR fell by roughly two-thirds compared to Variant A and no reactions were recorded.

Isolated positive signal – Variant B had a single post share, which was not recorded for Variant A. Given the low count, this is not considered an indicative of superior performance.

Variant A shows best performance – On all available engagement metrics Variant A outperforms Variant B; its CTR is more than three times higher.

Variant A delivered stronger user engagement and therefore I selected those creative assets to use in the higher quality channels: LinkedIn's general feed and special interest groups.

5.3.1. Results for the Recruitment Survey

I conducted an initial survey to gather demographic information from user experience (UX) designers and assess their exposure to machine learning (ML) and the Internet of Things (IoT). The purpose was to get a better understanding of the then state of UX designers' involvement with ML and IoT and recruit them to participate in the one-to-one workshops.

The final numbers were 28 survey responses. From those, 26 participants willing to participate in the workshops.

This section presents the survey results and highlights key findings and insights.

Demographic Information

Age Distribution: Respondents' ages ranged from 18 to over 54 years old. The largest age groups were 30–34 years (32.1%, n=9) and 40–44 years (25%, n=7).




Figure 16 - Age Distribution of Respondents.

Years of Experience in the UX Design practice: Experience levels varied, with the majority having between 5–10 years (50%, n=14) of experience.

The results revealed that only 21.4% (n=6) of participants had been involved in the UX design of an IoT product or application. This suggests that, despite the growing adoption of IoT technologies, a significant majority of UX designers in this sample have not worked on IoT related projects. This limited involvement indicates a potential disconnect between the rapid growth in numbers of IoT devices and the participation of UX professionals in designing the experiences with these devices.

Only one reported involvement in IoT projects that utilised ML to enhance user experience. This finding underscores a substantial gap in the integration of ML within IoT UX design. The minimal exposure could suggest that ML is not commonly employed in the context of Consumer IoT UX, possibly due to complexities associated with ML technologies or a lack of accessible design frameworks for UX practitioners.




Figure 17 - Years of Experience in UX Design.

In contrast, 28.6% (n=8) of participants had experience in the UX design of applications that used ML to assist user interaction outside the IoT context. This higher level of engagement indicates that ML integration is more prevalent in general app design. However, it is important to highlight that the presence of ML in these projects does not necessarily mean it was a result of UX design solutions. It is likely that ML was brought as part of a digital product feature.

Among those involved in ML-related UX design, a significant majority 75% (n=6 out of 8), were actively engaged in planning or designing the ML aspects of their projects. This involvement demonstrates that when ML is part of a project, UX designers tend to play a substantial role in informing how ML features are integrated into the user experience. Their contributions included cognitive task analysis and user research activities to better define ML requirements and collaborating on concept development and experience mapping. This is a significant finding – bearing in mind the small sample size of 6 out of 8 designers – as it contradicts previous views in the literature (Dove et al., 2017; Yang, 2018), that indicated that designers often were only involved later, after the role of machine learning had been defined. Clearly these numbers are not generalisable but poses a focus for future research.

The survey highlighted varied use of design resources among participants:

Open-Source Design Systems

Half of the respondents (50%, n=14) utilised open-source design systems, such as Google Material Design Guidelines²⁴, the GOV.UK Design System²⁵, and Bootstrap²⁶. This widespread use reflects the value placed on established design frameworks that provide consistency and efficiency in UI development.

Custom Design Systems

A majority (64.3%, n=18) reported using custom design systems developed for their companies or specific projects. This high usage of bespoke systems underscores the importance of tailoring design resources to meet unique organisational needs or project requirements. Many respondents who were freelancers or consultants indicated that it depends on the client.

Pattern Languages

Only 17.9% (n=5) indicated that they used a pattern language. However, it became apparent that many participants were unfamiliar with the term 'pattern language' as defined in the context of Alexandrian pattern theory. Some referred to UI libraries or design languages like Material Design instead. This suggests a lack of awareness or understanding of formal pattern languages within the UX community.

UI Pattern Libraries

Nearly half (46.4%, n=13) used UI pattern libraries, which aid in creating consistent and user-friendly interfaces.

Ideation Card Decks

Only one reported using ideation card decks, indicating that these tools are not widely adopted among the respondents.

Work environment

Most respondents worked as part of a larger studio, design or product team (39.3%, n= 11), while 21.4% (n=6) worked in UX specific team. However, many other roles associated with UX design were described, as well as shared responsibilities with other functions.

These survey findings support the overarching aims of the research by confirming that many UX designers have limited or no practical exposure to UX projects that Integrate ML. As such, these results align with **Research Question 1 (RQ1)** – the feasibility of identifying a pattern language that enable ML-driven UX design

²⁴ <https://m2.material.io/design/guidelines-overview>

²⁵ <https://design-system.service.gov.uk/>

²⁶ <https://getbootstrap.com/>

solutions for the Consumer IoT – and **Research Question 2 (RQ2)** – aimed at determining whether non-expert designers can effectively utilise such patterns. By revealing a significant gap in ML usage among practising UX professionals (only 28.6% had some relevant experience), the survey data supports the importance of creating accessible, well-structured pattern documentation for that user group.

Furthermore, the demographic profile captured here shows a broad pool of participants, enhancing the credibility of the patterns that subsequently emerge during the workshops. A fundamental aspect of patterns is that they must emerge from the community of practice.

5.4. The pattern mining workshops results

Building on GP1, which anticipates that a collection of ML-oriented UX patterns for Consumer IoT can indeed be identified, I put forward Sub-Proposition 1.2:

SH 1.2: A designer-centric pattern mining process, grounded in direct engagement with UX professionals, may yield a more contextually relevant and practically applicable set of ML-driven UX patterns than a purely automated approach relying on large-scale UI repositories.

This sub-proposition emerges from my earlier shift away from an automated labelling strategy and reflects the expectation that human insight may be critical for uncovering patterns that resonate with real-world design challenges in Consumer IoT. It also represents another key change in direction from the more limited scope of UI patterns towards the broader scope of UX patterns.

From the 26 respondents who agreed to participate in the workshops, three ultimately were unable to attend, making the total of workshop participants 23. Four of these sessions were face-to-face, in different locations in the UK and 19 were remote via video conference, with participants based in many different countries.

Following the completion of the 23 pattern mining workshops, the qualitative data consisted of transcripts capturing participants' insights and conceptualisations of user interactions with various smart devices. These transcripts were imported into MAXQDA for systematic qualitative analysis.

The coding process involved assigning codes to significant excerpts within the transcripts, isolating concepts such as:

- **User Concerns:** Wellbeing, safety, security, privacy.
- **Interaction Types:** Voice interaction, gesture-based control, remote operation, automation.

- **Machine Learning Aspects:** Proactivity, context awareness, user identification, adaptive automation.
- **Device:** Specific smart devices discussed during the workshops, such as smart locks, fridges, curtains, kettles, and beds.

The MAXQDA code system resulted in 322 coded segments, some of which were then exported into excel for the subsequent phase of analysis, resulting on **80 patterns seeds**, or early pattern concepts.

5.4.1. Shortlisted Good Pattern Candidates

In the pattern mining workshops, I gathered qualitative data from 23 UX professionals as they engaged in realistic Consumer IoT design scenarios. By applying the coding procedures described in Chapter 3, I initially extracted 80 distinct pattern seeds. These seeds emerged from a comprehensive analysis of participants' discussions and design artefacts, reflecting their real-world concerns, behaviours, and expectations regarding the integration of machine learning into UX design.

Importantly, this initial set of 80 seeds is the beginning of my direct response to Research Question 1 (RQ1) and Guiding Proposition 1 (GP1). I suggest that these seeds are not isolated ideas but are likely to be interconnected – forming the basis of a cohesive pattern language that integrates machine learning into adaptive, user-centred Consumer IoT experiences. To refine this broad set, I applied specific reduction criteria, which includes:

- **Relevance:** The pattern seed must address a recurring real-world UX challenge in Consumer IoT.
- **Generalisability:** The solution should be applicable across multiple contexts without necessitating specialised hardware.
- **Adaptive Capability:** The seed must suggest how machine learning can enhance interactions to create intelligent, adaptive experiences.

Using these criteria, I reduced the initial 80 seeds to 12 potential pattern candidates that genuinely leveraged machine learning to offer adaptive user experiences.

The 12 candidate seeds were given provisional names that captured their underlying solutions:

1. **Smart Inventory Management:** Integrates machine learning capabilities into smart fridges to monitor and manage food items by predicting restocking needs and reducing waste.

2. **Spoiled Food Detection:** Utilises sensors and ML algorithms within smart fridges to detect spoiled or near-expiry food and alert users.
3. **Gesture-Based Curtain Control:** Employs ML for gesture recognition to enable intuitive control of curtains or blinds.
4. **Voice-Activated Access Control:** Implements voice recognition supported by ML to control smart door locks.
5. **Adaptive Window Opacity Control:** Uses ML to automatically adjust window opacity based on environmental conditions.
6. **Adaptive Automation Based on User Habits:** Learns from user behaviours to automate settings and schedules, personalising the smart home environment.
7. **Sleep Pattern Analysis for Smart Beds:** Incorporates sensors and ML in smart beds to analyse sleep patterns and offer personalised recommendations.
8. **Predictive Boiling for Smart Kettles:** Utilises ML to predict when users need hot water, initiating boiling automatically.
9. **Proactive Energy Consumption Reporting:** Analyses household energy usage using ML to provide personalised energy-saving suggestions.
10. **Geo-Fencing-Based Automations:** Uses geo-fencing combined with ML to trigger actions based on a user's location.
11. **Intelligent Access Sharing:** Enables temporary access to Consumer IoT devices by using ML to suggest appropriate permissions based on past interactions.
12. **Smart Shortcuts:** Allows users to automate frequent, multi-step tasks by creating personalised shortcuts that can be triggered via smart assistants.

The fact that these 80 seeds could be reduced and interconnected into a smaller set of candidate patterns suggests that there is an underlying relationship with these adaptive ideas and indicates the potential for a pattern language that addresses the adaptive UX challenges in Consumer IoT. These findings further suggest that the decision to pivot towards a designer-driven pattern-mining approach was an effective decision.

5.4.2. Selected Patterns for Documentation

I further refined the 12 candidate pattern seeds by balancing specificity with generalisability. In this second stage, I ensured that the candidate patterns not only addressed the adaptive UX challenges of Consumer IoT but were also broadly applicable across diverse contexts without reliance on specialised hardware.

To achieve this, I considered the following dimensions critical for practical application:

- **Technical Feasibility:** The candidate pattern must be implementable with widely available technology.
- **Broad Applicability:** The pattern must be generic enough to be applied in multiple Consumer IoT settings.

By applying these dimensions, I further reduced the candidate list from 12 to a core set of four candidate patterns. In addition, during the investigation of implementations for the "**Adaptive Automation Based on User Habits**" and "**Geo-Fencing-Based Automations**" patterns, I identified a common, naturally coexisting design solution – which I documented as the '**Stop Learning Preferences**' pattern. This additional pattern did not emerge from the initial shortlisting exercise; rather, it arose as a complementary framework of best practices that supports user control and privacy, addressing a recurring need observed in conjunction with the first two patterns, i.e., an interconnection.

Thus, the five patterns I have selected for formal documentation are:

1. **Adaptive Automation Based on User Habits** (later renamed **Preference Learner**):

This pattern captures the ability of a system to learn from user behaviour and automate routine tasks. Its generalisable nature means that it can be applied across a wide range of Consumer IoT devices, making it a key candidate for supporting adaptive interactions.

2. **Geo-Fencing-Based Automations** (later renamed **Geographic Trigger**):

By leveraging geo-fencing technology, this pattern uses location data to trigger actions, offering a solution that leverages standard, widely available hardware and is applicable in multiple settings.

3. **Smart Shortcuts:**

This pattern enables users to automate frequent, multi-step tasks through integration with smart assistants. Its generality and support for 'ease of use' make it a valuable addition to the pattern set.

4. **Intelligent Access Sharing:**

Although I refer to this as a proto-pattern – owing to the limited real-world examples that meet the strict criteria of repeated implementation* – it presents a practical solution for temporarily granting access to Consumer IoT devices. It was included due to its support for adaptability and user control.

5. Stop Learning Preferences:

Identified as a natural relationship during my investigation of implementations of the above patterns, this pattern addresses the need for users to control or pause the learning process. It emerged as a critical component to ensure that the benefits of adaptive automation are balanced with ethical considerations and user autonomy.

* I was unable to find real-world examples for the adaptive version of the **Intelligent Access Sharing** pattern, i.e., examples that shows the integration of machine learning to augment any aspects of the setup and configuration. Therefore, this pattern does not adhere to the convention adopted by the pattern authoring community which determines that pattern documentations must include real-world examples as evidence that the solution have been re-applied multiple times (Alexander et al., 1977; Borchers, 2000a; Gamma et al., 1995; Harrison, 2003). For that reason, I use the term **Proto pattern** when referring to **Intelligent Access Sharing**. Proto patterns (Schümmer and Matschke, 2013) are initial ideas or solutions addressing recurring problems within specific contexts and they require further refinement and validation through repeated use – which forms the empirical evidence of broader applicability.

In addition to these five patterns that emerged from the candidate seeds refinement process, I later identified a sixth pattern – **Personal Spark**. It emerged during the **European Pattern Languages of Program** (EuroPLoP) pattern authoring review process when repeated reviewer enquiries led me to recognise that my methodological approach – initially intended solely to guide pattern discovery (or mining) – constitutes a reproducible design method in its own right, i.e. a solution for a recurring problem, or a pattern. This pattern encapsulates the human-centred approach I employed during the pattern mining study, particularly using the “What Would a Human Do” thought experiment, previously described. This led me to formally document **Personal Spark** as a separate pattern, which I classify under the Methodology Patterns category in chapter 6.

6. Personal Spark:

By making use of the “What Would a Human Do” thought experiment, this pattern enables designers to uncover adaptive opportunities by reimagining digital interactions through the lens of natural human-human interaction. It provides a structured method for uncovering subtle adaptive opportunities and serves as a mechanism for connecting advanced ML techniques with user-centred design practice.

Shortlisting these six patterns for formal documentation is a critical step in my research. This process marks the transition from exploratory pattern mining to the documentation of an interconnected pattern language. By focusing on patterns that satisfy the shortlisting criteria, I demonstrate that the seeds uncovered in the workshops (which, unbeknown to me at the time, had been generated using the Personal Spark pattern) are theoretically robust and practically useful. This result directly aligns with **Research Question 1** and **Guiding Proposition 1 (GP1)** by providing evidence of interconnected, adaptable design solutions that were about to form **the foundation of a UX pattern language for Consumer IoT that integrates machine learning**.

5.4.3. Contribution to the Research Project

The pattern mining process provided me with a related set of candidate pattern seeds, which serve as the initial step in directly addressing Research Question 1 and aligning with Guiding Proposition 1 (GP1). By extracting 80 seeds from the workshops and refining them down to 12 – and ultimately to five core patterns (with Personal Spark emerging later during the EuroPLoP review) – I have demonstrated that engaging UX professionals to simulate realistic design tasks can yield a more nuanced perspective on adaptive challenges in Consumer IoT.

These findings are significant for several reasons. First, they validate the effectiveness of a designer-centric pattern mining approach in uncovering the underlying concepts of adaptive solutions; the seeds are not isolated and show a natural relationship that suggests the potential for a coherent pattern language. Second, while the detailed interconnections among these seeds will be further elaborated during the formal documentation phase, these preliminary results already indicate the identification of a balanced, ML-enhanced UX pattern language is feasible. Such a language is critical to bridging the gap between advanced machine learning capabilities and user-centred design practice.

The current results lay a solid foundation for my subsequent step. They confirm that the methodological approach I adopted, based on qualitative analysis, is capable of uncovering a relevant set of UX design patterns. This stage directly addresses RQ1 by naturally leading to the formal documentation of patterns, which will also serve as practical resources for UX designers working in the Consumer IoT domain.

6. Formal Pattern Documentation

6.1. Introduction

Unlike the preceding study chapters, which follow a 'Participants → Apparatus → Procedure → Results' structure, this chapter is organised around the patterns themselves. The methodology through which the pattern seeds were identified is reported in Chapter 5; here, the focus is on the formal documentation of those patterns and the review process they underwent.

Building upon the pattern mining outcomes reported in Chapter 5, this chapter presents the formal documentation of interaction patterns that leverage machine learning to enhance human-computer interaction, inspired by human-human interaction. The aim was never to mimic human-human interaction but to abstract suitable aspects to encapsulate within these novel patterns. Influenced by the concept of Natural User Interfaces (Wigdor and Wixon, 2011), this approach aspired

to achieve more natural user interactions within the context of Consumer IoT devices.

After identifying high-quality interaction patterns²⁷ with feasible integration of machine learning to solve interaction problems more efficiently, I searched for examples of real-world implementations of these patterns. I evaluated these examples against established user experience principles and machine learning implementation guidelines. Specifically, Norman's interaction principles (Norman, 2013) provided a foundation for assessing usability aspects, such as affordance, feedback, and consistency. Additionally, Amershi et al.'s (2019) human-AI interaction guidelines offered a framework for objectively evaluating the potential of the patterns in enabling good user experiences.

The pattern mining approach described in Chapter 5 resulted in a substantial set of candidate pattern seeds. I converted these exploratory findings into a formally documented collection of six patterns. This phase is critical, as it transforms my initial, qualitative insights into the interconnected pattern language that directly addresses Research Question 1 and aligns with Guiding Proposition 1 (GP1). In doing so, I aim to bridge a gap between advanced machine learning capabilities and user-centred design practice within the Consumer IoT domain.

The successful patterns were then documented using formal approaches for pattern authoring (Harrison, 2003, 1999; Iba, 2014; Meszaros and Doble, 1997). Four of these patterns underwent a rigorous peer review process conducted at the premier European Conference on Pattern Languages of Program, EuroPLoP. The authoring and review process is described in the following section.

²⁷ Throughout this document I use the terms 'Interaction Patterns' and 'UX patterns' interchangeably, given that UX encompasses the wider scope ranging from user interactions and journeys to the whole experience. When these patterns address user interaction problems, they are inherently addressing user experience problems.

	AI Design Guidelines		Example Applications of Guidelines
Initially	G1	Make clear what the system can do. Help the user understand what the AI system is capable of doing.	[Activity Trackers, Product #1] "Displays all the metrics that it tracks and explains how. Metrics include movement metrics such as steps, distance traveled, length of time exercised, and all-day calorie burn, for a day."
	G2	Make clear how well the system can do what it can do. Help the user understand how often the AI system may make mistakes.	[Music Recommenders, Product #1] "A little bit of hedging language: 'we think you'll like'."
During interaction	G3	Time services based on context. Time when to act or interrupt based on the user's current task and environment.	[Navigation, Product #1] "In my experience using the app, it seems to provide timely route guidance. Because the map updates regularly with your actual location, the guidance is timely."
	G4	Show contextually relevant information. Display information relevant to the user's current task and environment.	[Web Search, Product #2] "Searching a movie title returns show times in near my location for today's date"
	G5	Match relevant social norms. Ensure the experience is delivered in a way that users would expect, given their social and cultural context.	[Voice Assistants, Product #1] "[The assistant] uses a semi-formal voice to talk to you - spells out 'okay' and asks further questions."
	G6	Mitigate social biases. Ensure the AI system's language and behaviors do not reinforce undesirable and unfair stereotypes and biases.	[Autocomplete, Product #2] "The autocomplete feature clearly suggests both genders [him, her] without any bias while suggesting the text to complete."
When wrong	G7	Support efficient invocation. Make it easy to invoke or request the AI system's services when needed.	[Voice Assistants, Product #1] "I can say [wake command] to initiate."
	G8	Support efficient dismissal. Make it easy to dismiss or ignore undesired AI system services.	[E-commerce, Product #2] "Feature is unobtrusive, below the fold, and easy to scroll past...Easy to ignore."
	G9	Support efficient correction. Make it easy to edit, refine, or recover when the AI system is wrong.	[Voice Assistants, Product #2] "Once my request for a reminder was processed I saw the ability to edit my reminder in the UI that was displayed. Small text underneath stated 'Tap to Edit' with a chevron indicating something would happen if I selected this text."
	G10	Scope services when in doubt. Engage in disambiguation or gracefully degrade the AI system's services when uncertain about a user's goals.	[Autocomplete, Product #1] "It usually provides 3-4 suggestions instead of directly auto completing it for you"
	G11	Make clear why the system did what it did. Enable the user to access an explanation of why the AI system behaved as it did.	[Navigation, Product #2] "The route chosen by the app was made based on the Fastest Route, which is shown in the subtext."
Over time	G12	Remember recent interactions. Maintain short term memory and allow the user to make efficient references to that memory.	[Web Search, Product #1] "[The search engine] remembers the context of certain queries, with certain phrasing, so that it can continue the thread of the search (e.g., 'who is he married to' after a search that surfaces Benjamin Bratt)"
	G13	Learn from user behavior. Personalize the user's experience by learning from their actions over time.	[Music Recommenders, Product #2] "I think this is applied because every action to add a song to the list triggers new recommendations."
	G14	Update and adapt cautiously. Limit disruptive changes when updating and adapting the AI system's behaviors.	[Music Recommenders, Product #2] "Once we select a song they update the immediate song list below but keeps the above one constant."
	G15	Encourage granular feedback. Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.	[Email, Product #1] "The user can directly mark something as important, when the AI hadn't marked it as that previously."
	G16	Convey the consequences of user actions. Immediately update or convey how user actions will impact future behaviors of the AI system.	[Social Networks, Product #2] "[The product] communicates that hiding an Ad will adjust the relevance of future ads."
	G17	Provide global controls. Allow the user to globally customize what the AI system monitors and how it behaves.	[Photo Organizers, Product #1] "[The product] allows users to turn on your location history so the AI can group photos by where you have been."
	G18	Notify users about changes. Inform the user when the AI system adds or updates its capabilities.	[Navigation, Product #2] "[The product] does provide small in-app teaching callouts for important new features. New features that require my explicit attention are pop-ups."

Table 5 - The 18 human-AI interaction design guidelines (Amershi et al., 2019)

6.1.1. Pattern Authoring and Review Process

I documented the patterns in a structured format, which typically includes the following elements (Harrison, 1999):

- **Context:** The circumstances under which the problem occurs, suggesting that solutions may not be applicable outside this context.
- **Problem:** A detailed description of the problem, highlighting the conflicting forces and challenges.

- **Solution:** The explanation of the proposed solution, including its general structure and examples.
- **Consequences:** Account of the implications of applying the pattern, encompassing both benefits and potential drawbacks.

This structured approach ensures clarity and facilitates the communication of design knowledge to the pattern users.

Peer Review at EuroPloP

Upon submission to the European Conference on Pattern Languages of Programs (EuroPloP), the patterns underwent a specialised reviewing process known as "Shepherding." In this process, an experienced pattern author, referred to as a "shepherd," provided iterative feedback and guidance. The shepherd, possessing expertise in the relevant field, conducted multiple rounds of reviews (**Figure 18**) allowing for substantial refinement of the patterns prior to resubmission.

Following acceptance, the patterns were subjected to further scrutiny during the conference workshops. These workshops comprised small groups of pattern authors from similar or closely aligned fields. All group members receive a copy of the papers they will be discussing in advance and are encouraged to read it before the workshop.²⁸ The collaborative environment encouraged participants to offer detailed feedback on each section of the documented patterns. This collective critique aimed to address specific issues and enhance the overall quality of the patterns.

- I would be curious to see a scientific reference that underlines the following statement: "The advent of the IoT is leading us to a new era [...]".
That is my interpretation of the views of Mike Kuniavsky in his book 'Smart Things: Ubiquitous Computing User Experience Design'. IoT now allows everyday objects to embed computation and data communication in our environments at larger proportions than we had before.

- IoT is not just about the everyday things in the consumer life. It is also about production planes, machine parks, and trade.
The scope of this paper is home/domestic IoT. I define that in section 3.1

Figure 18 - Extract of an exchange between myself (in purple) and reviewer during the pattern writing review process. It shows typical peer review aspects covering arguments presented throughout the entire paper, not restricted solely to the patterns.

²⁸ EuroPloP writers' workshop <https://www.europlop.net/conference/>

The multi-layered review process at EuroPloP plays a crucial role in ensuring the high quality of published patterns. Despite the rigorous Shepherding and workshop reviews, patterns remain subject to a final recommendation regarding their inclusion in the conference proceedings. Papers that do not meet the required standards may be excluded from the final publication.

This extensive review process guarantees that the patterns published through EuroPloP represent exemplary contributions to the field of pattern theory. Moreover, accepted papers that undergo further development are eligible for submission to the Springer journal LNCS Transactions on Pattern Languages of Programming (TPloP), a pathway I intend to pursue in the future.

I organised the six formally documented patterns into four categories:

- **Methodology Patterns:** Patterns that encapsulate methods and techniques which assist designers in the adaptive UX design process. In this category, I include the **Personal Spark** pattern.
- **Automation Patterns:** Patterns that enable systems to automate routine tasks, and therefore reducing cognitive load. These include patterns such as **Geographic Trigger** (previously named Geo-Fencing-Based Automations) and **Smart Shortcuts**.
- **Personalisation Patterns:** Patterns that facilitate the customisation of user experiences based on past interactions and user feedback. The **Preference Learner** pattern (formerly referred to as Adaptive Automation Based on User Habits) is a key example.
- **Privacy and Control Patterns:** Patterns that enable users to maintain control over their data and privacy. This category includes the **Stop Learning Preferences** pattern and the proto-pattern **Intelligent Access Sharing**.

Of these six patterns, four have undergone the more rigorous pattern authoring process, including multiple rounds of review and subsequent publication by the Patterns Language of Programs (PloP) conferences. The acceptance of these patterns by the pattern authoring community is in itself a testament to their quality. Due to time constraints, the remaining two have not yet been through the same review process.

This formal documentation phase is a crucial step in my research. By converting the conceptual pattern seeds into well-structured and formal design artifacts, I produced practical resources that demonstrate the feasibility of an ML-enhanced UX pattern language. This language not only advances theoretical knowledge in adaptive UX design but also offers practical guidance for addressing real-world

challenges in Consumer IoT. The sections that follow present each pattern within its respective category.

6.2. Methodology patterns

In this category, I consider patterns that encapsulate proven methods or techniques which can be directly integrated into the design process. These patterns provide designers with structured approaches to uncover adaptive opportunities and translate advanced machine learning capabilities into user-centred design solutions.

6.2.1. The Personal Spark Pattern – Identification and Naming ²⁹

Personal Spark is a methodology pattern based on the approach that I have developed to conduct the pattern mining study. This pattern emerged as a significant milestone during the iterative peer review process at the EuroPLoP conference in 2019, where reviewer enquiries prompted me to recognise that the methodology I had employed – especially my reliance on the “What Would a Human Do” (WWHD) thought experiment – was not merely a means to discover pattern seeds but a reproducible design strategy in its own right.

The Personal Spark pattern can be used to identify relevant fragments of human–human interaction as an initial step to inform requirements for adaptive user interfaces and data processing systems, particularly in the context of Consumer IoT. It is most applicable in the early stages of the interaction design process when the interaction requirements are still emerging. To design effectively for human behaviour, one must understand both the reactionary and anticipatory nature of the natural dialogue between two parties, in this case a person and a system – a dynamic that goes beyond mere usability (Kolko, 2010). Kolko also asserts that the most successful products and services often incorporate subjective qualities that make them stand out, and **Personal Spark** can assist in capturing these nuances.

At its core, **Personal Spark** is built upon the adoption of the WWHD thought experiment. This approach invites designers to envision how a helpful human might interact in a natural, spontaneous manner to complete a given task, instead of using digital means. By adopting this perspective, I was able to systematically explore opportunities where machine learning can be harnessed to imbue systems with adaptability and anticipatory capabilities. It is important to note that **Personal Spark** is not intended to produce outputs that mimic human–human interactions; rather,

²⁹ Parts of this section, including the ‘Personal Spark’ pattern were published in my 2019 paper ‘A pattern approach for identification of opportunities for personalisation and automation of user interactions for the IoT’ <https://doi.org/10.1145/3361149.3361157>

it offers a framework for isolating the ideal use cases for machine learning to suggest solutions for user interaction problems in the early stages of design.

The name **Personal Spark** was inspired by Christopher Alexander's concept of 'self-maintaining fire' present in 'living patterns', which posits that the integration of multiple interrelated patterns can bring an environment to life. According to Alexander (1979), when numerous living patterns are present, the environment 'comes to life as an entirety', exhibiting a self-sustaining vitality that is difficult to articulate yet universally recognised. In a similar vein, **Personal Spark** seeks to inject 'vitality' into design solutions by providing a structured, human-centred method for identifying opportunities for adaptive approaches. This pattern is aimed at helping UX designers navigate the challenges and opportunities of using data and machine learning as design resources by reimagining digital interaction problems in a human-to-human context.

Just as Alexander's living patterns enhance physical environments by addressing important human needs; by embracing traditional interaction design and Human-AI interaction principles in UX design, we can create user experiences that are not only effective and efficient, but also evoke a sense of satisfaction in users. This approach emphasises the creation of digital experiences that feel more natural, with the potential of ultimately contributing to a more positive relationship between users and technology.

'The more living patterns there are in a place (...) the more it comes to life as an entirety, the more it glows, the more it has that self-maintaining fire which is the quality without a name.'
(Alexander, 1979:X)

6.2.2. Personal Spark – Formal Documentation³⁰

Pattern name: Personal Spark

Context: Within an iterative design environment, the UX designer needs to find ways of improving the user experience design by identifying relevant opportunities for personalisation and/or automation based on what the system knows about the user and/or the context of use at a given time. Since these findings may lead to new

³⁰ As published in my 2019 paper 'A pattern approach for identification of opportunities for personalisation and automation of user interactions for the IoT' <https://doi.org/10.1145/3361149.3361157>

requirements, earlier identification is critical to ensure that appropriate resources are secured.

Problem: At the early concept stage, UX professionals need to identify indicative, relevant adaptation variables and provide a contribution to early discussions about suitable inference mechanisms necessary to design adaptive user interactions.

Forces:

- Personalisation can take place in many ways, and it may become overwhelming in the early stages to identify the ideal hierarchy of elements to personalise.
- User testing sessions are extremely useful if designed correctly. However, the User Centred Design process outside academia is often less rigorous due to time and budget constraints. The availability of user testing in the project lifecycle is often reduced to certain critical stages such as testing prototypes or evaluation between design solutions.
- More complex requirements such as Machine Learning or device-device communication are less likely to be implemented if their need is only identified at later stages in the project lifecycle.
- Modern product design approaches such as Design Thinking and GV Design Sprints favour informal and straightforward methods before committing to formal and more rigorous ones.
- Despite the small appetite formal methods at the early stages of the product design, most designers still need to structure their design approach with some form of validation and prioritisation.
- Designers are familiar with the idea of working from a high-level concept towards a more constrained and better scoped solution.

Solution:

Designers can initially conceptualise the user interaction based on their own expertise, applying the available heuristics, and using well established Human-Computer Interaction approaches. Once that initial concept is completed, they can then ask the collaboration of a 'Proxy' to use the What Would a Human Do (WWHD) approach to identify relevant opportunities for personalisation and/or automation, which will be influenced by the natural human interaction perspective. The Proxy can be a colleague or any other person familiar with the problem the design concept is trying to solve. This person is called Proxy because they will envisage how the whole human to human interaction would happen, therefore representing both themselves and the other human during the WWHD thought experiment.

The WWHD approach can be performed as follows:

If the concept covers more than one task, the designer should isolate each task to allow a more relevant analysis.

It is not unusual to find two or more tasks to merge into a single and more complex interaction during the WWHD part. However, it is more valuable for the design that such convergence occurs naturally during this process, rather than the designer imposing them in the original concept.

1. The designer shares the newly conceived user interaction with the Proxy.
2. The designer then asks the Proxy to perform a very simple thought experiment: to imagine that, instead of the UI as a point of interaction, there is a human being performing the role of the interface.

So, in effect the designer is lifting the constraints imposed by the user interface whilst creating the opportunity for documenting what a human-human interaction could look like in the narrow context of that task completion.

3. The designer then focuses on documenting information, asking the Proxy to explain any unclear step, a technique very similar to moderating user testing sessions. The main objective of this step is to identify the elements of human-human interaction which could in turn inspire a more natural interaction.

Using the WWHD once should already highlight relevant opportunities for personalisation, but more complex tasks will benefit from more than one session. However, we found that beyond two or three sessions the novelty of the findings decreased dramatically.

Consequences

Benefits:

The designer now has the original interaction concept and the corresponding human-human interaction, allowing them to extract the most relevant and feasible fragments from the later to enrich the user interaction concept. As a result, the technology recommendations necessary to support the personalisation and/or automation are driven by the user interaction needs, rather than resource availability, technical convenience or any other non-user-centric motivation.

The WWHD approach fits naturally within the cyclic process of the iterative design methodology. Its concept anticipates the need for further refinement and improvement, inherent to this design technique.

Liabilities:

Certain tasks cannot be performed by a human and therefore the WWHD approach would fail to achieve any useful results. In certain scenarios, a graphic user interface, for example, offers a better solution than a human-human interaction (e.g., scheduling a repetitive task or reviewing long activity logs). Any improvement in this case would frequently fall under well documented usability heuristics or HCI methods, rather than warrant uniquely tailored interaction solutions. Further, there must be special considerations for certain interactions supporting the input and output of information: people do not naturally expect high precision in a casual human interaction, but we are accustomed to retrieving and providing accurate values when interacting with a computer. Moreover, the information an average human would naturally communicate to another human is not only limited in volume, but also far less accurate than that which a computer system can easily provide. In some cases, users may favour the accuracy over a humanised interaction e.g., entering the number of degrees instead of simply instructing the system to make it warmer or cooler.

Example 1:

We used Personal Spark on nine IoT user interaction concepts for a smart kettle which led us to quickly identify that:

- 'A degree of autonomy' was the most common feature expected for an improved interaction. It highlighted a need for an autonomous verification functionality, which would be responsible for checking if all resources necessary to perform the task are available. This should happen without direct user instruction and be informed by the user's patterns of usage (e.g., check water levels, set maximum temperature or pre-empt daily interactions). Potential use case for time series analysis.
- Other improvements in the interaction could be achieved by simply applying a few of well-known Nielsen's usability heuristics: 'Match between system and the real world', 'Error prevention', 'Recognition rather than recall' and 'Visibility of system status'.

Example 2:

The pattern is used to design contact forms (Meyer and Wachter-Boettcher, 2016). It helps to assess whether the form language is clear enough and identify any possible questions could arise during its use. It also assists in exposing more

relevant flows. Although this example concentrates on aspects of design and it does not necessarily show any clear use case for machine learning, the potential still exists.

Example 3:

The use of the Personal Spark pattern has been recommended when creating the knowledge base to inform the behaviour of intelligent agents for game engines. The aim is to use the pattern to achieve frequent adaptation of the knowledge and improve the realism of the agents as the pattern always tries to extract 'What Would a Human Do' (van Lent and Laird, 1999).

Example 4:

In the earlier stages of a Design Sprint for conceptualising a solution to interact with a smart window, the Personal Spark pattern can be used to help to identify suggestions of more natural interactions to be tested in a user interface prototype.

Example 5:

I used Personal Spark with 23 participants as a method for pattern mining. All were UX design practitioners, which led to the identification of the other five patterns presented in the next sections. The resulting interaction concepts contained executions of UX patterns that had real world implementations. Although I used the pattern to identify Novel, undocumented UX patterns incorporating machine learning in the context of Consumer IoT, a total of 80 interactions were identified, many of which could be used to inform design decisions in the right circumstances.

6.3. Automation Patterns

Automation patterns are designed to reduce cognitive load by enabling systems to perform routine tasks automatically. In my formal documentation, I include two patterns in this category:

- **Geographic Trigger** (previously presented as Geo-Fencing-Based Automations): This pattern leverages location data to trigger actions, providing a broadly applicable solution that uses widely available hardware.
- **Smart Shortcuts**: This pattern allows users to automate multi-step tasks through integration with smart assistants, therefore streamlining frequent interactions.

The **Geographic Trigger** pattern has undergone expert review at the European Conference on Pattern Languages of Programs and has been published. In addition, this pattern was tested in the validation study described in Chapter 7.

6.3.1. Geographic trigger pattern – overview³¹

The Geographic Trigger Pattern focuses on designs that rely on GPS and other location technologies to use location as a form of input. The primary use case is that users can set a connected device to perform certain operations when they enter or leave a specific location.

6.3.2. Geographic trigger – formal documentation

Pattern name: Geographic Trigger

Context:

During their daily routine, a user finds him/herself moving between different locations but constantly carrying a mobile phone.

Problem:

The user wants a connected device to perform a particular action when he/she enters or leaves a specific location to alleviate the burden of remembering certain mundane, but critical tasks linked to a geographic location.

Forces:

- In certain circumstances, simply scheduling tasks to happen at a specific date/time is not an ideal solution as the user may not know when he/she needs the task performed.
- There are situations where the task execution might only be applicable if the user enters or leaves certain locations.
- When the users are travelling, they may be satisfied that their location is a good indication of estimated time of arrival, and they may want to use that as a task trigger.
- Certain notifications from the device are more relevant if the users receive them in certain locations.
- Leaving or arriving at certain locations are effective events to indicate progression on user's daily routine.
- Most modern mobile devices have built-in GPS.
- The task in question should happen as part of a sequence of events, rather than the typical chronological date and time specification.

Solution:

³¹ Parts of this section, including the 'Geographic Trigger' pattern were published in my 2020 paper '**Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT**' <https://doi.org/10.1145/3424771.3424777>

Allow the user the option for his/her mobile phone location to be used as a passive mobile application UI input mechanism to prompt the connected device to execute predetermined tasks associated to certain geographic locations.

The user should be able to mark geographic areas and associate what tasks would be performed by the connected device when he/she approaches, enters or leaves the demarked area. For example, the user could instruct a smart kettle to boil the water when he/she is approaching a certain distance from home, or a smart lock could be set to lock itself if the user crosses a certain distance from home.

To ensure effectiveness, clearly communicate with the user the need to grant access the mobile device's location information to allow the app to collect the user's current location appropriately.

Likewise, for greater efficiency the application's UI should display the following options in the same categorical context:

1. A method to easily select the task to be performed by the connected device.
2. A visual method to select locations in an interactive map and associate it with the desired task.
3. A method to allow the user to choose whether the task should be triggered if he/she enters, leaves or reaches a specific distance from the location.
4. A list of previously set up tasks associated to geographic locations, to avoid unintentional repetition or conflicts.

In addition, the UI should immediately confirm whether the setup was successful, with appropriate error handling.

Primary Principles of Interaction:

- Mapping: The UI exhibits correspondence or representation of the desired outcome, or it behaves in a manner that does so.
- Feedback: Communicates the results of an action.
- Conceptual Models: Subjective, high-level explanation of how something works.

Consequences

Benefits:

- No impact on the user's lifestyle. The automation happens while the user goes about their daily routine as usual.

- The connected device can be instructed to perform tasks which are applicable when the user approaches, enters or leaves certain geographic locations.
- The user can instruct the connected device to issue notifications if they enter a specified geographic location.
- Opportunity for sequences of different locations the user travels through, rather than just a single location, to be used as distinct triggers for tasks.
- Less cognitive demand and less stress caused by oversights. He/she can rely on the connected device's task being executed automatically based on his/her geographic location, rather than having to inform precise dates/times.

Liabilities:

- Unplanned changes in itinerary may have adverse effect in this kind of automation. E.g., if the user changes their normal route due to traffic or floods and does not pass through a previously targeted location, this could cause the automated task to be ignored. That might not be what the user expected.
- Areas with poor mobile signal may impact on the functionality.
- The solution does not work if the user does not grant access the mobile device's location information.
- Finally, it may be less effective for usage within a short range of distance as issues associated with GPS precision, WiFi, Cell Towers and latency could impair the functionality.

Known uses

The controller app for the smart lock Ultraloq U-Bolt Pro allows the users to setup the auto lock functionality.

It sets a geofence of around 300 meters around the lock and if the users (i.e., their mobile phones) leave the device's position for at least 300 meters, it locks itself automatically.

The August Smart Lock controller app also uses the same pattern to allow similar functionality.

The IFTTT mobile app uses the pattern to allow users to set up location-based services to control a range of connected devices.

Finally, the automobile manufacturer BMW uses a variation of the pattern to enable the firm's proprietary eDrive Zones function. In this case it automatically switches

newer hybrid vehicles engines to fully electric mode if they reach a geofenced emission-free zone.

Additional comments

Although not documented in the first publication of this pattern, Geographic Trigger may leverage ML to provide more accurate automations. For example, by processing road traffic data, it can more accurately trigger automations that are meant to have started when user arrives at their destination, e.g. heat a room to a specific temperature or having a bath ready.

6.3.3. Smart shortcut pattern – overview

The **Smart Shortcut** pattern enables users to automate frequently performed, multi-step tasks in a more convenient way. By allowing users to create personalised shortcuts, this pattern leverages the integration with smart assistants (e.g. Siri, Alexa, or Google Assistant). After a user completes a task, the application prompts them to save such sequence of steps as a shortcut, which they can name for easy recall. These shortcuts can then be triggered through voice commands, reducing the need to open the app each time.

6.3.4. Smart shortcut – Formal Documentation

Context:

In the consumer Internet of Things (IoT) domain, users often perform frequent, multi-step tasks using mobile applications to interact with connected devices. With the proliferation of smart assistants (e.g., Siri, Alexa, Google Assistant), there is an opportunity to enhance user experience by making use of the integration with these assistants to simplify task execution.

Problem:

The user needs to automate frequently performed, multi-step tasks with Consumer IoT devices, in a more intuitive way, allowing them to initiate these tasks conveniently even when the outcomes may vary slightly each time.

Forces:

- The advent of smart assistants promoted a now mature interaction model, primarily centred around voice user interfaces with advanced natural language processing capabilities. It is sensible for app developers to leverage these assistants rather than developing proprietary functionalities.

- To adopt automation, methods need to be straightforward without complex configurations. Previous approaches such as macros were not widely adopted by non-expert users.
- Integration with smart assistants is a common practice in the industry and the resulting experience is familiar to existing users.
- Normally users have the ability to pause, modify, or stop smart assistants' integrations at any time, making it a non-permanent, lower risk integration.
- Some users place their time-saving gains above the concerns about data sharing and how their interactions are monitored, especially with third-party assistants.
- Some businesses have a financial incentive to simplify how users complete certain multi-step tasks.
- In some Consumer IoT environments, smart assistants are used as form of hub to control multiple smart devices. It is logical integrating with these existing entry points of interaction to incentivise the use of key features.
- Making the integration available would not interfere with users who prefer manual control or are uncomfortable with smart assistants.

Solution:

Provide users with the option to create personalised 'Smart Shortcuts' within the Consumer IoT device's application for certain key tasks they perform frequently. After completing a multi-step task, the app prompts the user to save this sequence as a shortcut, allowing them to name it for easy recall. These shortcuts integrate with smart assistants, enabling users to trigger them through voice commands without opening the app.

Key Implementation Aspects:

1. **In-App Prompting:** After task completion, prompt the user to create a shortcut, ensuring the scope and context are clear.
2. **Naming Shortcuts:** Allow users to assign custom names to shortcuts, aiding recall when using voice commands.
3. **Integration with Smart Assistants:** Seamlessly provide mechanisms to connect the new Smart Shortcuts with the user's preferred smart assistant (e.g., Siri Shortcuts, Alexa Skills, Google Assistant Routines).
4. **User Control Panel:** Provide an interface within the app to manage shortcuts, including options to pause, modify, or delete them.
5. **Feedback and Confirmation:** Offer clear feedback during shortcut creation and usage, confirming actions and outcomes.
6. **Privacy Settings:** Include settings that inform users about data usage and allow them to control what information is shared with smart assistants, specially what is essential and optional to make the automation possible.

Strive for transparency – make use of the **Stop Learning my Preferences** pattern.

7. **Account to interruptions:** Enable users to easily recover if the setup is interrupted – by either clearly signposting to continue from where they stopped or explicitly constraining them to start over.

Primary Principles of Interaction (Norman):

- **Signifiers: Communicate what actions are possible.** The app should clearly indicate the option to create a shortcut after task completion, perhaps through a prominently placed button or icon.
- **Conceptual Model: Help users understand how the system works.** Offer explanations or tutorials on how shortcuts function and how they can be triggered via smart assistants.
- **Constraints: Limit possible actions to prevent user errors.** Ensure that only appropriate tasks can be saved as shortcuts to avoid unintended consequences.

Primary Human-AI Interaction Design Guidelines (Amershi et al.):

Initially

- **G1: Make clear what the system can do.** Ensure that users understand the system's capabilities, specifically that they can create and use voice-activated shortcuts via integration with smart assistants. This information should be available in the app, ideally presented when users first encounter the shortcut feature.

During interaction

- **G4: Show contextually relevant information.** Display the option to create a shortcut immediately after the user completes a relevant multi-step task. This contextual timing leverages user attention at a moment when the feature will appear most relevant and useful.
- **G7: Support efficient invocation:** Enable the users to easily call the shortcut using voice commands.

When Wrong

- **G8: Support efficient dismissal.** Allow users to dismiss the shortcut creation prompt easily if they are not interested, ensuring that this action is simple and non-intrusive. Users should feel that they can opt out without disrupting their task flow.

Over Time

- **G16: Convey the consequences of user actions.** Clearly explain how creating a shortcut will impact future interactions, such as enabling voice activation for the specified task. This guidance ensures users understand the persistent nature of the shortcut and its effects on their routine.
- **G17: Provide global controls.** Offer settings to manage all created shortcuts, allowing users to pause, modify, or delete them. This centralised control supports user autonomy by allowing them to adjust how shortcuts operate over time according to their changing needs or preferences.

Consequences:

Benefits:

- **Efficiency:** Users save time by reducing the number of steps required for frequent tasks.
- **Convenience:** Enables hands-free operation via voice commands using smart assistants.
- **Personalisation:** Users create shortcuts to fit their routines and preferences.
- **Non-Essential Nature:** Does not impede users who prefer manual control, as the feature is optional and non-intrusive. Those unfamiliar with the use of smart-assistants and those who decide not to use them, will be able to interact with the app without the automation.

Liabilities:

- **Privacy Concerns:** Users may hesitate due to concerns about data sharing with smart assistants.
- **Dependency on Third-Party Services:** Relies on external platforms, which may introduce compatibility or reliability issues.
- **Complexity for Some Users:** Despite the pattern incentive for simplicity, the general concept might be intimidating for less tech-savvy individuals.
- **Limited Adoption:** Users unfamiliar with or distrustful of smart assistants may not utilise the feature.
- **Risk of Errors:** Without proper safeguards, automation might lead to unintended actions if misused.
- **Need for high confidence on which are the users' top tasks:** In order to simplify the action of saving relevant tasks, designers will have to identify the top candidate tasks to be automated and carefully design the 'moment of interruption' of the interaction to prompt the user to automate such tasks.
- **Privacy compliant:** The pattern should not be used as a covert mechanism for collecting user data for any other use (e.g. advertisement), unless it is clearly stated.

Avoid this pattern when:

- Tasks involve sensitive operations where errors could have serious consequences (e.g., financial transactions).
- The target audience is unlikely to adopt voice-activated features due to accessibility issues, have significant privacy concerns or prefer not to use smart assistants. (As usual in UX design, knowing who the primary users are is key)
- Dependency on third-party platforms introduces unacceptable risks or complexities.

Known uses:

Philips Hue smart lighting enables users to create scenes activated by voice commands through smart assistants³².


Alexa Routines let users automate sequences of actions with a single voice command or at scheduled times³³ and The Tile app allows users to locate lost items using personalised voice commands³⁴.

Finally, eBay's iOS app (Figure 19) enables users to access their most recent order through a Siri phrase. Although not a direct example of use in a Consumer IoT context, this implementation is a good example of a relevant moment in the interaction to invite users to automate the task.

³² <https://www.philips-hue.com/en-gb/explore-hue/works-with/amazon-alexa/alexa-commands-for-hue-lights>

³³ <https://www.amazon.co.uk/gp/help/customer/display.html?nodeId=G2PYLKJN3XVZ55EQ>

³⁴ <https://www.philips-hue.com/en-gb/explore-hue/works-with/amazon-alexa/alexa-commands-for-hue-lights>



Screenshot of eBay's iOS app – this implementation of Smart shortcut serves to streamline the multi-step process of accessing the most recent order.

Figure 19 – An implementation of the Smart Shortcut pattern used in the eBay iOS app.

6.4. Personalisation Patterns³⁵

Personalisation patterns focus on tailoring user experiences to individual behaviours and preferences. This category of patterns facilitates the implementation of personalised user experiences primarily by relying on data derived from past interactions with user-defined preferences and feedback.

The **Preference Learner** pattern (initially referred as Adaptive Automation Based on User Habits) is example in this category. In a nutshell, it leads to solutions can learn from a user's interactions over time, thereby automating routine tasks in a way that is customised to the user's needs. This pattern is essential in reducing the burden of manual configuration while ensuring that the adaptive features remain aligned with user expectations.

Personalisation proves to be especially advantageous in cases where users engage with a system repeatedly, and when the variables influencing the interactions are diverse and substantial in volume. Such complexity invites more sophisticated solutions beyond the mere consideration of temporal factors, such as date and time.

³⁵ Parts of this section, including the formal documentation of the 'Preference Learner' pattern were published in my 2020 paper 'Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT' <https://doi.org/10.1145/3424771.3424777>

6.4.1. Preference Learner Pattern – Overview

The **Preference Learner** pattern enables smart devices apps to learn and adapt to individual preferences and routines over time. Instead of requiring users to manually input extensive settings and schedules, this pattern produce solutions that instruct applications to observe interactions during a designated learning period. This observation is actually a data collection process that, when combined with the use of machine learning, leads to automation of tasks in a personalised way - aligned with the user's habits. Users can guide the learning process by providing feedback that work as corrections, ensuring the applicaton's behaviour reflects their preferences. This approach improves the user experience by reducing the need for complex, manual configurations, offering a more intuitive interaction with their connected devices.

6.4.2. Preference Learner – Formal Documentation

The user can instruct the connected device to learn specific characteristics of automation by watching their interactions with the application for a period of time.

Context:

The user has a number of preferences for the operation of the smart device which typically form part of his/her routine.

Problem:

The user wants to set a smart device to perform certain tasks autonomously in a personalised fashion, however it is very time consuming to enter every single preference, scheduling and other data values to achieve that.

Forces:

- Machine learning enables the use of vast amounts of data as input to make inference of desired outputs.
- There are long sequences of events and data points that would need to be fed into to the system to achieve the desirable level of personalisation and it would be an arduous task for any user to input that manually.
- Modern mobile phone devices have a number of sensors which can be used to better inform the user's routine and accelerate the data entry.
- The mobile device and the connected device have limited capabilities and resources.

Solution:

The smart device's UI allows the user to enable a learning mode through which the application can learn the user's preferences and usage routine by monitoring their interactions for a period of time.

To ensure effectiveness, during the 'learning period', the user should be given opportunities to correct and nudge the application towards their preferred behaviour.

In addition, when technically viable, provide an option to allow the application to continue to monitor the user's interactions even after the initial 'learning period' has ended, so that the user can benefit from interactions getting continuously further aligned with their needs. In that case, the user should be made aware of how their interactions with the device may influence the application's learning.

This pattern increases the user interaction efficiency by transferring the burden of the complex data entry to the application, which translates in greater user satisfaction.

Primary Principles of Interaction:

Signifier: Communicates what action is possible.

Feedback: Communicates the results of an action.

Conceptual Models: Subjective, high-level explanation of how something works.

Primary Human-AI interaction design guidelines:

- During Interaction:
 - G5 Match relevant social norms: Ensure the experience is delivered in a way that users would expect, given their social and cultural context.
- When wrong:
 - G9 Support efficient correction: Make it easy to edit, refine, or recover when the AI system is wrong.
- Over time:
 - G12 Remember recent interactions: Maintain short term memory and allow the user to make efficient references to that memory.
 - G13 Learn from user behaviour: Personalize the user's experience by learning from their actions over time.
 - G14 Update and adapt cautiously: Limit disruptive changes when updating and adapting the AI system's behaviours.

- G15 Encourage granular feedback: Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.

Consequences

Benefits:

- The user can continue to go about his/her daily business without interruption and have that routine used as a learning period to automate a number of mundane tasks.
- For convenience, the mobile device's own sensors can be used as inputs (e.g., GPS and accelerometer).
- The user can replace a number of laborious data input activities by just normal operation of the UI to 'train' the application on his/her preferences.
- Once the application learns the user's preferences, the task delegation can be automated in a more personalised manner.
- Even though the user will be performing the tasks for a while before the tasks can be performed automatically, the knowledge that learning is in progress serves as feedback which might put the user at ease, potentially due to the anticipation of the benefits of having a higher degree of personalisation.

Liabilities:

- This pattern works well if the user performs a daily routine. If the user changes his/her habits of interaction with the connected device, the automation settings may need manual adjusts.
- In certain cases, the user will still have to assist the learning algorithms by tagging certain events manually. Without such interventions the learning process would take too long or even be unfeasible for consumer level products.
- For simpler applications, some machine learning models can be trained locally in the user's mobile phone. This may significantly drain the devices' performance.

Known uses

The Google Hello doorbell allows the user to teach the camera which faces are familiar by tagging unfamiliar faces as 'familiar'. Overtime (approximately two weeks), the device begins to automatically recognise familiar faces and when they ring the doorbell, a notification with their names appears on the mobile screen or announced through the smart speaker. Other similar examples are both Google Cam Outdoor and Google IQ Outdoor which offer similar familiar faces feature, however instead of a doorbell interaction, these are internet connected security cameras.

The Nest thermostat features the Auto-Schedule functionality which serves to instruct the system to learn from the user's initial interactions. The learning period lasts for about a week and after that the learning functionality remains on to continue to adjust the user's routine but attributing less weight to the impact of each change the user makes.

The controller app of Amazon Alexa virtual assistant allows users to instruct the device to learn their voice so that they can have personalised interactions for news, shopping, messaging and calling. The learning is achieved by having the user to read out loud certain sentences and eventually a voice profile is created.

Related patterns:

Action-Reaction (Vega-Barbas et al., 2018) – Vega-Barbas et al. describes the Action-Reaction pattern in the context of a Smart Space where 'points of interaction between the service and the user are transparently immersed in the environment through the use of everyday objects that have been augmented with computational capabilities'. The pattern uses the user's actions and tasks to model the Smart Space's interaction.

However, in Preference learner the user is in explicit control of the interaction i.e., the user has to enable the functionality. Further, the pattern addresses an interaction problem with a connected device, and not with an environment. Ubiquitous computing is not a pre-requisite for the Preference learner pattern, as it is the case for the Action-Reaction pattern.

Stop Learning Preferences – This pattern is presented in the next section. It solves interaction problems that can only occur once the Preference learner pattern is applied.

6.5. Privacy and Control Patterns³⁶

Privacy and Control Patterns address the critical need of enabling users to manage how their data is used and to maintain control over adaptive systems. In this category, I have documented:

³⁶ Parts of this section, including the formal documentation of the 'Stop Learning Preferences' pattern were published in my 2020 paper '**Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT**' <https://doi.org/10.1145/3424771.3424777>

- **Stop Learning Preferences:** This pattern addresses the recurring user need of pausing or stopping the learning process triggered by patterns such as **Preference Learner**, ensuring that user control and privacy are maintained.
- **Intelligent Access Sharing:** Although designated as a proto-pattern due to limited real-world examples meeting the standards to be considered a repeated implementation, this pattern presents an innovative framework for temporarily granting access to Consumer IoT devices. It promotes the use of adaptive and context-sensitive access management while safeguarding user control.

6.5.1. Stop Learning Preferences Pattern – Overview

Recognising that continuous data collection may not always align with user preferences, the **Stop Learning Preferences** pattern enables users to pause or permanently discontinue the learning processes initiated by mechanisms like the **Preference Learner** pattern. While machine learning enhances personalisation by adapting to individual behaviours, there are times when users may wish to halt this data gathering – perhaps due to privacy concerns, changes in routine, or impacts on device performance. This pattern offers clear options for users to disable learning functionalities and delete any stored data, ensuring they maintain control over their personal information.

6.5.2. Stop Learning Preferences – Formal Documentation

At some point the user may decide that he/she does not want the smart device, nor its controller app tracking their interactions.

Context:

User had previously activated a learning functionality of a connected device to increase the personalisation of the automation routine, as described in the **Preference learner** pattern.

Problem:

The user who previously set the connected device's UI to learn from his/her routines now wants the learning process to pause or stop permanently.

Forces:

- At some point during the learning process the user starts having privacy concerns.
- The user decided to operate the device manually for a period of time or permanently.
- The learning process causes performance degradation of the user's mobile device.

- There is a noticeable increase in resource consumption of both the user's mobile device and the connected device during the learning process.

Solution:

The interface allows the user to pause the learning for as long as he/she wants or to completely reset any learned behaviour and delete any existing history. The option should be placed in the same spatial context where the controls for activation of the learning functionality are located.

Primary Principles of Interaction:

- Signifier: Communicates what action is possible
- Feedback: Communicates the results of an action
- Conceptual Models: Subjective, high-level explanation of how something works

Primary Human-AI interaction design guidelines:

- During interaction:
 - G4 Show contextually relevant information: Display information relevant to the user's current task and environment.
- Over time:
 - G15 Encourage granular feedback: Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.
 - G17 Provide global controls: Allow the user to globally customize what the AI system monitors and how it behaves.

Consequences

Benefits:

- The user has the sense of control over their data and privacy. He/she can decide to pause or stop the data collection.
- The user can permanently delete the data and has no worries about misuse of any previously stored data.
- The user's mobile device performance improves.
- No resources are used in the user's mobile device or in connected device

Liabilities:

- The user may not completely comprehend the impact of disabling the functionality and erasing all the learnt history as he/she might have been accustomed with the device's personalised behaviour for a long time.
- Only once the historical data is lost, the user might appreciate how much the lack of the finely tuned personalisation impacts his/her daily routine.

Known uses:

The Google Hello doorbell gives the user the option to turn familiar face detection off and to delete a familiar face profile.

Both Google Cam Outdoor and Google IQ Outdoor once more offer very similar ways of stopping the face detection functionality, delete a previously identified person's profile and to delete the entire library.

In the Nest thermostat mobile app, the user can turn the Auto-schedule off as well as to erase the entire learnt history.

Related patterns:

Farewell (Vega-Barbas et al., 2018) – The Farewell pattern also addresses an interaction problem associated with control and privacy, but with an entire ubiquitous computing environment instead of a single connected device. The pattern also anticipates that the Smart Space initiates the interaction autonomously, in contrast Stop Learning Preferences anticipates the user's intent as the starting point of the interaction. The purpose of the collected data is not explicit in the Farewell pattern, but the solution refers to a dialogue between a pervasive service and the user on whether to delete or store information when he/she stops using a service or leaves the Smart Space. On the other hand, Stop Learning Preferences focuses on giving the user control over an ongoing data collection process, which may or may not continue in the future.

Preference learner – The Stop Learning Preferences pattern is directly linked to the Preference learner pattern, solving interaction problems associated with the use of the latter. The pattern allows the user to pause or stop whatever was initiated by the Preference learner pattern.

6.5.3. Intelligent Access Sharing Proto-Pattern – Overview

The Intelligent Access Sharing proto-pattern provides a user-friendly method for granting temporary rights to operate Consumer IoT devices or resources, such as smart locks for example. With the integration of machine learning it can provide contextual access predictions and adaptive permission recommendations by analysing data such as date/time, location, calendar events, routine behaviours and past access-sharing instances. It may, for example, suggest appropriate permission levels and durations for each contact or situation, thus automating the process and reducing effort.

This pattern simplifies access sharing, reduces cognitive load, and maintains user control and privacy, all without imposing on users the complexities of traditional user management systems.

6.5.4. Intelligent Access Sharing – Formal Documentation

Context:

In the consumer Internet of Things (IoT) context, users sometimes need to grant temporary access to their smart devices or specific features to different people.

Problem:

The user needs an efficient, effective and flexible way of enabling different levels of temporary operation access to their smart devices or resources – such as allowing guests to enter their home using a smart lock or enabling a service provider to adjust settings on a connected appliance. This task needs to be accomplished without requiring the users to navigate complex user management systems.

Forces:

- Admin and member management systems require the user to add the new person and ‘remember’ to reset any changes that are meant to be temporary.
- Users must feel confident that sharing access is safe and that they maintain control over permissions and duration.
- Users may encounter situations with limited internet connectivity, necessitating alternative methods for granting access.
- Users normally appreciate assistance that simplifies tasks, provided it respects their autonomy.
- Traditional access management systems can be overly complex and cumbersome for non-expert users, and the complexity and sophistication of typical administration panels based on access control lists is not usually necessary.
- Users may be cautious about how their data is utilised, preferring transparent systems that respect their privacy.
- Smartphone devices are pervasive, with most individuals in developed countries likely to own one.

Solution:

Provide an intuitive mechanism within the Consumer IoT application that enables the process of granting temporary access. Similar to sharing access token approach, commonly used between applications to determine the identity and levels of privilege.

Enhance the user experience by integrating machine learning (ML) to provide intelligent assistance in two ways:

1. **Contextual Access Predictions:**

- Employ ML to analyse contextual data (e.g. time of day, location, calendar events, and routine behaviours etc) to predict when the user might need to grant access, and proactively offer prompts.

2. **Adaptive Permission Recommendations:**

- Use ML to analyse past access-sharing instances to recommend appropriate permission levels and access durations for different contacts or situations. If the user preferences change, the ML model should adapt, providing up-to-date recommendations that reduce effort and improve consistency.

Key Implementation Aspects:

1. Design a step-by-step process that is intuitive, reducing the effort required to share access.
2. Use plain language to describe access levels, avoiding technical jargon.
3. Present ML-suggested permission levels and durations based on user's past behaviour, allowing quick selection or easy adjustment.
4. Allow user to set specific durations for access, with options for immediate revocation.
5. Intelligent Behaviours:
 - Offer helpful prompts when the system predicts a likely need to grant access – for example, suggesting access sharing when a calendar event indicates a guest is arriving or if routine behaviours imply an upcoming service visit.
 - Ensure suggestions are timely, relevant, and presented without being intrusive.
6. Provide options to share access via SMS or other non-internet-dependent methods when necessary, accommodating limited connectivity situations.
7. Employ the **Stop Learning Preferences** pattern (or a relevant variation) so that the user is offered settings to customise privacy preferences, allowing control over the extent of ML analysis, including options to disable ML features.

Primary Principles of Design:

- **Affordance and Signifiers:** Clear use of buttons, iconography and labels to help the user through the access-sharing process and what options are available.

- **User Control and Freedom:** User can easily adjust permissions, revoke access, or dismiss suggestions, maintaining a sense of control.
- **Consistency and Standards:** The interface integrates known UI design patterns, reducing the learning curve.
- **Visibility of System Status:** User is kept informed about the system's status through timely notifications.

Primary Human-AI interaction design guidelines (Amershi et al.):

Initially

- **G1: Make clear what the system can do:** Clearly communicate the capabilities of the system, including how machine learning enhances access sharing with intelligent suggestions.

During interaction

- **G3: Time services based on context:** Offer suggestions at appropriate moments, enhancing relevance and without disrupting the user.
- **G4: Show contextually relevant information:** Provide access-sharing options and ML-driven suggestions when and where they are most helpful.

When wrong

- **G8: Support efficient dismissal:** Allow users to easily dismiss suggestions, ensuring that ML assistance enhances rather than hinders the experience.
- **G13: Learn from user behaviour:** Continuously adapt to the user's actions and preferences, refining predictions to be more personalised.

Over time

- **G16: Convey the consequences of user actions:** Clearly explain what will happen when the user grants access, including permissions and duration.
- **G17: Provide global controls:** Offer settings to adjust how the system uses data and to opt in or out of ML-driven features, respecting user autonomy. (Consider using the **Stop Learning Preferences** pattern)

Consequences

Benefits:

- Users can grant access with minimal steps, aided by intelligent prompts that anticipate their needs.
- Adaptive permission recommendations simplify setting permissions, tailoring the process to individual preferences.
- ML assistance reduces the mental effort required to remember when to grant access and what permissions to set.
- Time-limited and easily revocable access helps to maintain user control.

Liabilities:

- Users may have concerns about data usage; transparent communication and control options are essential to build trust.
- ML suggestions may not always align with user needs; the system should allow easy dismissal and learn from user feedback.
- To be effective, ML depends on the quality of the data available; limited data may reduce accuracy, leading to repeat instances of not useful recommendations and potential user frustration and loss of trust.
- The adaptive aspect is not the core benefit of the pattern, the core UX problem can still be reasonably well addressed without ML enhancements.

Known uses³⁷

The August smart lock controller app allows users to provide temporary access to guests, within desired access levels. https://support.august.com/how-do-i-invite-guests-r1BRlIJA_M

LastPass users can share passwords and decide whether to allow the recipient to view the password or to solely use it without viewing it <https://www.lastpass.com>

The Apple's iOS Share Wi-Fi functionality is a variation of this pattern which allows a user to share the password with visitors, however it lacks the levels of control the other implementations often feature.

Apple Home app enables users to share controlling privileges of smart home resources with someone, as well as to revoke such privileges. <https://support.apple.com/en-gb/102386>

6.6. Discussion

The formal documentation of these six patterns represents a significant contribution to UX design within the Consumer IoT domain. By developing the pattern seeds into the set of patterns that formed the foundation of a pattern language, I have directly addressed **Research Question 1** and found strong alignment with **Guiding Proposition 1 (GP1)** – demonstrating that a designer-centric, ML-enhanced pattern language could be identified.

³⁷ By the time of writing, I was able to identify only non-adaptive real-world implementations of this pattern. In these instances, it is common to refer to it as a 'proto pattern' until at least three implementations of the pattern can be found.

Throughout the documentation and peer-review process, I was attentive to potential conflicts or tensions between patterns, an aspect that the Alexandrian tradition may affect any pattern language, where patterns may exert competing forces. In this language, the most prominent tension currently exists between Preference Learner, which progressively automates aspects of the interaction based on inferred user preferences, and Stop Learning Preferences, which protects user autonomy by allowing the suspension of adaptive behaviour. Rather than treating this as a contradiction, I documented the two patterns as complementary: Stop Learning Preferences is explicitly noted as a required companion whenever Preference Learner is deployed, a dependency captured in the directed acyclic graph (DAG). Another example, though less pronounced, tension was identified between the convenience offered by Intelligent Access Sharing and the privacy considerations it introduces; I addressed that by embedding the use of Stop Learning Preferences explicitly as part of the solution and including transparency recommendations in the Consequences section. Naturally, future tensions may emerge as the language expands; the most appropriate strategies to mitigate them should be explored as they arise.

However, it is clear that the patterns documented by this research do not represent the totality of the Pattern Language. Nonetheless, one of my contributions can still be leveraged to further expand this pattern collection – interested pattern authors may leverage the Personal Spark pattern to mine more patterns for this language.

Each of the six documented patterns occupies a specific role in the pattern language. In the formal documentation, I have included “Related patterns” sections where relevant to explicitly delineate how they interconnect – in some cases there are mention of relationships with patterns that are outside this pattern language. However, this is limited to the identified patterns at the time of writing the documentation. The most effective way of showing the interconnections of this set of patterns forming a language is through a directed acyclic graph (DAG), as introduced by Borchers (2000b).

Figure 20 shows a proposed representation of this pattern language as a DAG. In this representation, I illustrate how the patterns – documented in this Chapter 6 – are interrelated. The graph is based on my documented categories and reflects the following structure:

- **Personal Spark** (Methodology Pattern) is the foundational node that underpins the past, and potentially the future, discovery process of new patterns.
- From Personal Spark, all other patterns emerge:

- Preference Learner
- Geographic Trigger
- Smart Shortcuts
- Stop Learning Preferences
- Intelligent Access Sharing (proto-pattern)
- In order to deliver better user experience, **Stop Learning Preferences** need to be applied whenever **Preference Learner** and **Intelligent Access Sharing** are present.
- The presence of nodes called “Future Pattern” represent the ability of this pattern language to evolved through the use of **Personal Spark** as a pattern mining mechanism.




Figure 20 - The resulting directed acyclic graph (DAG) representing this novel pattern language.

This directed acyclic graph succinctly encapsulates the structure of the ML-enhanced UX pattern language. It demonstrates that the patterns are not isolated but are interconnected – each contributing to a larger framework that bridges advanced machine learning capabilities with user-centred design practice.

Because the documented patterns occupy complementary rather than competing roles (spanning methodology, automation, personalisation, and privacy/control) no direct conflicts between them were identified during documentation or peer review. Where potential tension exists between automated personalisation and user privacy, the pattern language resolves it structurally: the Stop Learning Preferences

pattern for example, is prescribed whenever Preference Learner or Intelligent Access Sharing is present, ensuring that privacy safeguards are embedded in the language itself rather than left to ad-hoc designer judgement. This structural conflict-resolution mechanism is also visible in the DAG above, where the directed edge from Stop Learning Preferences to the personalisation patterns encodes this mandatory co-occurrence. This representation further addresses **Research Question 1** and aligns with **Guiding Proposition 1 (GP1)**.

The formal documentation phase allowed me to synthesise diverse insights into a unique framework – a Novel Pattern Language. The explicit mapping of related patterns reinforces the interconnected nature of these design solutions and suggests that a ML-enhanced UX pattern language addressing Consumer IoT context is feasible. This work advances the theoretical knowledge and offers tangible design resources, ultimately making a significant contribution to the field.

7. Validation Study

7.1. Introduction

The validation study presented in this chapter is designed to address **Research Question 2 (RQ2)**: "If identified, can such a pattern language support the effective design of user interaction concepts for Consumer IoT scenarios by designers with no machine learning expertise?".

This study also examines **Guiding Proposition 2 (GP2)**, which anticipates that "When employed by designers with no deep expertise in machine learning, this pattern language may support the creation of feasible Consumer IoT interaction concepts."

To explore this proposition, I conducted a proof-of-concept study with a sample of experienced UX professionals, using two formally documented UX design patterns: **Preference Learner** and **Geographic Trigger**. My aim was to evaluate whether these patterns provided sufficient design guidance for participants to create technically

feasible and user-centred IoT UX flows without requiring deep knowledge of machine learning.

The validation study involved one-to-one workshops with experienced UX designers. During these sessions, I provided them with structured design tasks via pre-prepared materials on Miro boards. They were required to generate UX flows using one of the patterns. After that, these flows were evaluated using a comprehensive rubric that measured criteria such as comprehensiveness, flexibility and customisation, user control and privacy, ease of use, and the potential of facilitating communication with engineers. In addition, as detailed in Section 7.5.3, an independent IoT expert assessed the technical feasibility of the design solutions.

This validation was intended to explore whether the documented pattern language could be used in a plausible professional UX design setting, particularly during early-stage concept development in the Consumer IoT context. More specifically, the study examined whether structured patterns might serve as a useful design aid when designers are asked to incorporate ML-related capabilities without needing deep technical expertise.

Taken together, the qualitative analysis, descriptive ratings, and expert feasibility assessment provide an exploratory basis for considering the practical usefulness of the pattern language. This chapter presents an initial proof-of-concept account of how the patterns were used, how the resulting artefacts were assessed, and what these findings may suggest for future research.

7.2. Methodology

7.2.1. Ethical Considerations

This study followed the institutional ethics guidelines from the University of Hertfordshire particularly regarding participant privacy and informed consent. Participants in the validation study were informed about the study aims, format, and their right to withdraw at any time. Each participant provided written or electronic consent prior to data collection.

To maintain confidentiality, participants were assigned unique codes (e.g., D1-D5). Any recordings or transcripts used for analysis were stored on password-protected drives, accessible only to the researcher and supervision team. Similarly, any references to identifiable user data were de-identified before integration into the dissertation (with exception of the IoT Expert who was happy to be identified in this work).

Compliance and Approval

Ethical clearance was obtained from the University of Hertfordshire Social Sciences, Arts and Humanities Ethics Committee with Delegated Authority. The UH protocol number is aCTA/PGR/UH/05871. The data collection involved minimal risk and no vulnerable participants.

7.2.2. Research Design

The theoretical underpinning for this part of the research methodology is based on scenario-based design and the use of design patterns in HCI. Scenario-based design focuses on creating and analysing narratives that describe user interactions with systems in specific contexts (Rosson and Carroll, 2002). Design patterns, as conceptualised by Alexander et al. (1977), provide reusable solutions to common design problems, facilitating communication and efficiency in the design process. By combining these frameworks, I evaluated how UX patterns can be effectively employed by designers to recommend complex technological integrations in the context of the Consumer IoT without requiring deep technical expertise.

This section addresses **RQ2** – whether the newly documented patterns are practical and usable for designers lacking in-depth ML skills – and thus directly explores **Guiding Proposition 2 (GP2)**, which anticipates that non-expert designers can successfully use the pattern language to produce adaptive, user-centred IoT concepts.

I adopted a qualitative, exploratory research design grounded in HCI principles. Given the novel nature of UX patterns integrating machine learning in the context of IoT applications, an exploratory approach was appropriate to gain in-depth insights into designers' experiences and perceptions when using it. As the use of UX patterns in Consumer IoT is relatively less explored (when compared to other domains), this approach is suitable for identifying key themes and insights that may inform future research and practice. In addition, the scenario-based method aligns with the study's focus on practical application, enabling the assessment of how designers interact with the patterns in a simulated but realistic context. Finally, incorporating an expert assessment adds an extra layer of validation by evaluating the technical feasibility of the designers' outputs, ensuring that the findings are based, as much as possible, on practical implementation considerations.

I organised this Validation Study into four primary stages:

1. Participant Recruitment: Identification and selection of UX designers who met specific inclusion criteria to participate in the study.

2. Workshops with UX Designers: One-to-one workshops where designers were introduced to two novel patterns and asked to design solutions using these patterns.
3. Assessment of the solutions produced: Using a specific rubric, I conducted a scenario-based evaluation of the solutions produced by the UX designers in the stage two above.
4. Assessment by an IoT Expert: Engaging an IoT expert to evaluate the solutions provided by the designers to achieve a thorough and objective assessment of the impact of using the patterns.

This structure allowed for a comprehensive examination of the patterns' utility in a UX design context.

7.3. Participants

Participant Recruitment

To ensure the relevance and applicability of the findings, I established specific inclusion criteria for participant selection:

- Age: Participants had to be 18 years old or older to comply with ethical research standards involving adults.
- Professional Experience: Participants were required to have at least five years of experience in UX design or a related field. This criterion ensured that participants possessed sufficient professional expertise to engage with the design tasks³⁸.
- Availability: Willingness and availability to participate in a one-hour online workshop were required.

Exclusion criteria included individuals without professional experience in UX design, those under 18 years of age, and those unable to participate in the workshop.

Recruitment approach

I utilised the professional network LinkedIn to recruit participants, aiming to achieve five participants, which I successfully did. To mitigate potential biases associated with recruiting from personal networks, I made efforts to open the invitation to a diverse group of professionals in terms of age and years of experience.

Given that I had to rely on participants' self-reported years of experience, having previously met four of the participants in settings associated with professional UX

³⁸ As a reference, the **Chartered Society of Designers**, among other requirements, requires applicants **seeking Chartered Designer Status** to have practised in the field of design full-time for at least five years. The 'Chartered status is the ultimate international recognition of professional practice' in a design specialism, the Society's website states.
<https://www.csd.org.uk/membership/chartered-designer/>

practice gave me greater confidence that the study was being attended by individuals who genuinely possessed the claimed levels of expertise.

This approach balanced the practical considerations of recruiting experienced UX designers currently practising in the industry within a limited timeframe while maintaining the integrity and validity of the study.

The target of five participants was determined by both the proof-of-concept scope of this study and the practical realities of recruiting experienced UX designers willing to volunteer their time for a research study. As established in human-computer interaction research, small samples of experienced practitioners can yield in-depth insights when the aim is to surface the principal strengths and weaknesses of a design artefact rather than to establish population-level generalisability (Nielsen, 2000, 1993). Recruiting senior professionals who are actively practising in industry presented challenges due to their limited availability, competing professional commitments, and the absence of direct incentives for participation. Because this study sought to determine whether the pattern language could, in principle, enable designers with little ML expertise to produce feasible IoT interaction concepts, depth of engagement with each participant took precedence over breadth of sample. Each one-to-one workshop generated detailed qualitative data that complemented the rubric-based ratings, and the subsequent independent expert assessment provided an additional layer of external validation.

Study Information Page

Participants were directed to the study information page (see appendix G). Given the temporary nature of these studies and the fact that I owned and had full hosting control of the uxdesigners.co.uk domain, I reused it for this second study. In addition, for consistency, I used the same page template used for the previous study webpage. The information page provided all the necessary information about the purpose of the research study, format and the voluntary nature of participation.

From the information page, participants could access a **Screening Survey** hosted on Jisc's Online Surveys platform.

The Screening Survey aimed to:

- a) Capture Informed Consent: Obtain participants' consent to partake in the study.
- b) Screen for Eligibility: Ensure participants were 18 years old or older and met the experience requirement.
- c) Gather Demographic and Professional Background Information: Collect data on the designers' age group, years of experience in UX design, involvement in IoT and ML projects, and their usage of design systems.

Throughout the recruitment process participants were assured of the confidentiality of their responses and the anonymity of their identities in any publications resulting from the study. I ensured that they understood their right to withdraw from the study at any time without the need for explanations. The recruitment materials avoided any coercive language, emphasising the voluntary nature of participation.

7.4. Apparatus and Materials

The Proto Persona

To simulate a real-world design scenario, I developed a proto persona named Alex Smarty – a summarised user profile. This approach aligns with common UX practices where designers are provided with user personas to guide their design decisions. The inclusion of a proto persona helped to enable a setup closer to real-world scenarios where the designer is asked to design for a particular user group and often has some information about that user to provide greater context around the user's needs.

An illustrative image accompanied the persona to enhance engagement and empathy. (See Figure 21).

Some brief information about our user

Alex Smarty



Goals:

- Streamline everyday tasks and save time.
- Ensure home appliances and systems provide comfort and efficiency.
- Stay connected and updated with personal preferences even with a busy lifestyle.

Key pains and constraints

1. Concerned about the security and privacy of smart devices, especially when granting temporary access.
1. Wants to ensure devices don't become overly intrusive in personal preferences and habits.
1. Looks for intuitive user interfaces in smart devices; doesn't want to spend too much time adjusting settings.

Behaviours and Preferences:

- Enjoys experimenting with device settings initially but appreciates when devices learn and adapt over time.
- Occasionally reads tech blogs and watches review videos to stay updated on the latest smart devices.
- Values voice commands and integrations with smart assistants for a unified home automation experience.
-

Background

- Lives in a 3-bedroom house in the countryside.
- Tech-savvy and always eager to integrate the latest smart devices into daily life.
- Prefers convenience and automation in everyday tasks.

Figure 21 - Alex Smarty, the proto persona used in the workshop. Photo by Caique Nascimento via unsplash.com.

Adapted Pattern Documentation

I provided documentation of each pattern divided into three main parts to facilitate understanding and application by the designers (see Figure 22 and Figure 23)

1. Summarised Version of the Formal Documentation: Using a less formal written register. The summarised description aimed to be accessible, avoiding technical jargon and overly scholarly language. The focus was on facilitating practical application by the users (UX designers in this case).
2. Diagrammatic Version of the Pattern: I created a visual representation in form of the typical user flows that would result when the pattern is applied.
3. A wireflow³⁹ with a real-world example of the pattern using actual app screenshots or mock-ups. These wireflows demonstrated:
 - User Interfaces: How the pattern appears in the interface, including layout and design elements.
 - Interactions: The user's actions and the system's responses.
 - Contextual Integration: How the pattern fits within the broader application or system.

This resource provided concrete guidance to aid designers in understanding how to integrate the patterns into their designs.




Figure 22 - Miro board with the pattern documentation for the Preference Learner pattern.

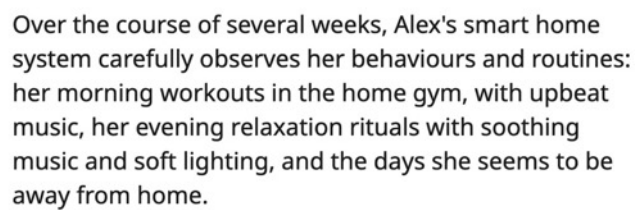
³⁹ Combination of wireframes and flowcharts <https://www.nngroup.com/articles/wireflows/>



Figure 23 - Miro board with the pattern documentation for the Geographic Trigger pattern.

Development of Scenarios and Tasks

I crafted two scenarios, each associated with one of the patterns, to contextualise the design task. These scenarios were introduced as a desired outcome, followed by a task that asked the designers to use the pattern documentation to design a UX flow enabling the user (Alex) to make the necessary configurations to achieve the desired outcome (see Figure 24 and Figure 25).



This scenario is the resulting outcome of an earlier setup process or user interaction. Using the Preference Learner pattern as your reference, please document the user interaction flow for the initial configuration that could enable the experience Alex had.

129



After a long day at work, Alex is driving home. As she approaches her neighbourhood, her phone's GPS data is relayed to her smart home system.

Sensing she's about 10 minutes away, the system activates the driveway lights. The the indoor temperature had also been adjusted to her preferred evening setting. By the time she pulls into her driveway, her home is well-lit, welcoming, and comfortably warm, ready for her arrival.

This seamless integration ensures that Alex's transition from outside to indoors is smooth, personalised, and energy-efficient.

Using the pattern below as your reference, please document the user interaction flow for the initial configuration that could enable the experience Alex had.

Figure 25 - Scenario and task associated with the Geographic Trigger pattern.

Tools Used

To prepare and conduct the validation study, I utilised the following tools:

Miro Virtual Whiteboards

Miro was used to design and store the pattern documentation and assets, and to enable the UX designers to document their solutions. Its collaborative features facilitated real-time interaction during the workshops and enabled the participants' design artefacts being captured in situ, automatically saving to the cloud and therefore reducing the chances of loss of such visually rich data. The virtual whiteboards allowed participants to:

- Interact with Materials: View and navigate through the proto-persona, pattern documentation, and scenarios.
- Document Solutions: Create user flow diagrams using shapes, connectors, and text elements.
- Collaborate: Engage in discussions with me and view any comments I would add to board during the guided discussion.

Additionally, Miro's video chat facility was used on two occasions when audio problems occurred with Microsoft Teams.

Midjourney AI Image Generation

I employed Midjourney to create the scenario and pattern header illustrations. This AI tool was chosen given its ability to generate bespoke visuals based on natural language prompts in matter of minutes. This allowed me to easily generate customised visuals that enhanced the materials' appeal and contextual relevance.

The images helped in:

- Visual Engagement: making the materials more engaging and easier to relate to.
- Increase Contextual Understanding: Providing visual context to the scenarios and patterns.

Microsoft Teams

Microsoft Teams was used to conduct the remote sessions. It provided a secure, GDPR-compliant communication method which was freely available through the University's office 365 subscription. It has video conferencing capabilities, screen sharing for collaborative review of materials, and chat functions for more immediate communication. Features utilised included:

- Video Calls: Facilitated remote interaction, enhancing communication.
- Screen Sharing: Allowed me to guide participants through the materials and observe their work on the Miro board.

- Chat Functionality: Enabled the exchange of messages and links when needed.

Spyder - The Scientific Python Development Environment

Used for statistical analyses, calculation of Spearman's correlation coefficient and statistical significance testing.

Microsoft Excel

Used for basic formatting given its efficiency in handling tabled data.

7.5. Procedure

7.5.1. The Workshops

Prior to the workshops, I scheduled sessions at mutually convenient times for each participant. I sent reminder emails before the scheduled session, including instructions for accessing the Miro board and Microsoft Teams meeting.

I also prepared individual Miro boards for each participant, pre-loaded with all the necessary materials, including the proto-persona, pattern documentation, scenarios, and spaces for the designers to create their user flows.

Conducting the Workshops

Each one-to-one workshop lasted approximately one hour and followed a semi-structured format, aiming at having as much consistency across sessions as possible, while aiding deeper discussions. The one-hour long sessions offered a pragmatic compromise between the data collection needs and the practical constraints of recruiting busy practitioners who were making themselves available voluntarily. The remote format enabled flexibility for time and location, while also simplifying the data collection by using tools like Miro, as explained previously.

In addition, dividing the sessions into three segments – Introduction, Design Task, and Discussion – helped to ensure that the required data were produced while minimising participants' cognitive fatigue.

This is how each segment of these workshops was structured:

Introduction (15–20 minutes)

The session began with a warm welcome and an overview of the study's objectives. I then introduced the materials:

- Proto-Persona: I presented Alex Smarty's profile, highlighting her as a typical user relevant to the design task.

- Pattern Documentation: I guided the participant through the three components of the pattern documentation, ensuring they understood the pattern's purpose and application.
- Scenarios and Tasks: I read aloud the scenario corresponding to the assigned pattern and discussed it to clarify any questions.

Participants were encouraged to ask questions throughout this segment to ensure clarity and understanding.

Design Task (20 minutes)

Participants were given the task instruction:

"The scenario above is the resulting outcome of an earlier setup process or user interaction. Using the [Pattern Name] pattern as your reference, please document the user interaction flow for the initial configuration that could enable the experience Alex had."

I explained to the participants that the experience described was the outcome, the result of a UX flow – what could that UX flow have looked like?

Designers were then instructed to use Miro's tools to create a user flow diagram, and I remained available for questions but refrained from providing guidance that could influence the design decisions.

Guided Discussion and Feedback Collection (20–25 minutes)

Following the design task, I conducted a semi-structured interview to gather qualitative feedback. The discussion aimed to explore:

- Experience Using the Pattern: How the designer found the process of applying the pattern.
- Challenges Faced: Any difficulties or obstacles encountered during the task.
- Perceptions of the Pattern's Utility: The perceived effectiveness of the pattern in facilitating the design.
- Suggestions for Improvement: Ideas for enhancing the pattern documentation or the pattern itself.

Specific questions included:

1. Did the UX patterns effectively address the needs of the users?
2. How would you describe the process of implementing the patterns in the design?
3. If not mentioned above, did any challenges arise while using the patterns?
4. Was the presentation of the pattern clear and understandable?
5. Would you suggest utilising these patterns in relevant upcoming projects?
6. Has the pattern impacted either positively or negatively the efficiency of the design process?

7. What improvements would you suggest for the documentation of these patterns to make them easier to use?

However, the interview was conversational and depending on comments already provided by the designer during the session, I would skip some the questions above as I already had an answer. This allowed for follow-up questions based on the participant's responses and to delve deeper into relevant points.

Ethical Considerations During the Workshops

I ensured that participants were comfortable throughout the session. I reiterated that there were no right or wrong answers and that the focus was on understanding their experiences. If any participant showed signs of discomfort or stress, I was prepared to pause or terminate the session.

7.5.2. Assessment of Designers Output using Scenario-Based Evaluation

I codenamed the UX designers D1, D2, D3, D4 and D5 in order to anonymise their output. Three designers created user experience flows using the Preference Learner pattern and two designers created user experience flows using the Geographic Trigger pattern.

I conducted the evaluation of proposed UX flows using the rubric detailed in section 7.5.2. I also included a dedicated category to capture the effectiveness of each flow in aiding communication between UX designers and engineers – an important role for these novel patterns in the modern Agile product development environments.

Alternative Methods Overview

In developing the research methodology for this study, I considered several established evaluation methods within the field of Human-Computer Interaction (HCI), including Heuristic Evaluation, Cognitive Walkthrough, Task Analysis, and typical Usability Guidelines Compliance. However, upon careful examination, I determined that scenario-based evaluation was the most appropriate approach for assessing the effectiveness of the UX flows provided by the designers. This decision was significantly influenced by the nature of the design artefacts produced during the study.

The UX flows were not traditional user interface designs; instead, they were presented as diagrams or documented in sequences of Miro sticky notes. This unique format had a considerable impact on the suitability of the evaluation methods, as it required an approach capable of assessing abstract, high-level conceptual designs rather than concrete interface elements.

Heuristic Evaluation, a method where experts assess an interface against established usability principles (Nielsen and Molich, 1990), was considered but ultimately deemed unsuitable. This method is primarily designed for evaluating designs with detailed user interface, focusing on elements such as visual consistency, error prevention, and user control within tangible interface components. Given that the UX flows in this study were abstract representations lacking specific interface elements like buttons, menus, or visual layouts, applying heuristic principles would be difficult. I would have no concrete interface components to examine, making it challenging to identify usability issues as defined by traditional heuristics. Therefore, the heuristic evaluation's reliance on detailed UI elements made it incompatible with the abstract nature of the designers' outputs.

Similarly, **Cognitive Walkthrough** was considered but found inappropriate for this context. This method simulates users' cognitive processes when interacting with an interface, focusing on learnability and ease of use for new users (Nielsen and Mack, 1994). Once again, it requires a more detailed interface through which the evaluator can step, considering the user's goals, actions, and system responses at each stage. However, the UX flows produced were high-level diagrams outlining sequences of actions and decisions without concrete user interface interactions. There were no specific steps or interface elements for an evaluator to walk through cognitively. As such, it would be impractical to apply this method to abstract diagrams, as it relies on the presence of tangible interface components to assess potential user difficulties and learning challenges.

Task Analysis, which involves breaking down tasks into their component actions to understand user goals and the interactions required to achieve them (Kirwan and Ainsworth, 1992) was also considered. While task analysis is valuable for understanding the complexity of tasks and identifying potential issues in task completion, it is most effective when applied to detailed processes within existing systems or prototypes. The UX flows in this study were conceptual representations, lacking the granularity required for a traditional task analysis. They did not specify detailed user actions or aimed to capture comprehensive system responses at the interface level but instead mapped out overarching flows and decision points. Consequently, task analysis would not be a good fit to assess the effectiveness of these high-level designs, as it could not adequately dissect the abstract sequences presented in the sticky note diagrams for example.

Lastly, **Usability Guidelines Compliance** involves evaluating a design against established standards and best practices to ensure usability and accessibility. This method requires specific interface elements and interaction designs to compare

against guidelines, such as those relating to visual contrast, navigation structures, or input controls. The absence of concrete interface components in the designers' outputs meant there were no specific elements to assess for compliance. Once more, the abstract diagrams did not contain sufficient detail to evaluate adherence to usability guidelines, which are typically geared towards fully realised interface designs. As a result, this method was not suitable for assessing the effectiveness of the UX flows provided in this study.

Considering the limitations listed above, **scenario-based evaluation** emerged as an appropriate method for this research. This approach focuses on assessing how well a design addresses the specific needs and goals outlined in a user scenario, making it more suitable for evaluating conceptual designs. By using a rubric tailored to the study's objectives, the designers' output could systematically be evaluated based on criteria relevant to the scenarios provided. The rubric included dimensions such as comprehensiveness, flexibility, user control, privacy considerations, ease of use, and the ability to communicate effectively with engineers. This allowed for a structured assessment of the abstract designs, focusing on how effectively they translated the UX patterns into solutions that met the scenario's requirements.

Evaluation Rubric Categories and Respective Rationales

In developing an objective approach to evaluate the UX flows created in this study; it was essential for me to consider critical aspects of user experience and pattern effectiveness. The evaluation categories I chose were grounded in established UX evaluation dimensions from the literature, particularly drawing on the theoretical frameworks proposed by Hassenzahl (Hassenzahl, 2004) and Hussain et. Al. (Hussain et al., 2021)

Hassenzahl's model distinguishes between **pragmatic** and **hedonic** qualities in user experience. Pragmatic qualities pertain to the product's ability to support the achievement of practical goals, focusing on usability, functionality, and task efficiency (Hassenzahl, 2004). Hedonic qualities, on the other hand, relate to the user's psychological well-being and emotional fulfilment, encompassing aspects such as stimulation and identification.

Building upon these concepts, Hussain et al. (2021) proposed a model for the design and evaluation of interactive products, which includes an additional dimension of **trust**. This dimension addresses factors such as privacy, security, credibility, dependability, and transparency.

For the purposes of this study, three dimensions were most relevant for evaluating the UX flows:

- **Pragmatics:** This dimension encompasses ergonomic and task-oriented aspects of the user experience that meet users' functional goals (Hassenzahl, 2004; Hussain et al., 2021). Under this dimension, I included the categories of **Comprehensiveness** and **Ease of Use**, assessing how effectively the UX flows enable users to accomplish their tasks efficiently.
- **Hedonics:** Focusing on emotional and psychological aspects that contribute to user satisfaction and pleasure (Hassenzahl, 2004; Hussain et al., 2021). The category of **Flexibility and Customisation** falls under this dimension, as allowing users to tailor their experience can enhance enjoyment and fulfilment.
- **Trust:** A critical dimension that involves users' perceptions of privacy, security, and control when interacting with systems (Hussain et al., 2021). The category of **User Control and Privacy** aligns with this dimension, reflecting the importance of users feeling secure and in control within adaptive systems.

In addition to these dimensions, I included a fifth category for evaluation: **Communication with Engineers**. This category is vital because one of the key roles of the novel UX patterns is to facilitate effective communication between designers and engineers (Borchers, 2001). By assessing how well the UX flows support this communication, we can evaluate the patterns' effectiveness in bridging the gap between design and technical implementation.

Rubric for Evaluating UX Flows

To ensure a systematic and objective evaluation of the UX flows, I developed a detailed rubric. This rubric provides guidance for rating the workflows across the five criteria introduced in the previous section:

1. Comprehensiveness of Setup
2. Flexibility and Personalisation
3. User Control and Privacy
4. Ease of Use
5. Communication with Engineers

Each category is graded on a scale from 1 to 5, with specific standards defining each level of performance. This structured approach allows for a nuanced assessment of the designs, grounded in established principles of usability and human-computer interaction.

Comprehensiveness of Setup

Definition: Comprehensiveness assesses the thoroughness of the setup process within the UX flows, ensuring that all necessary steps are included to facilitate a smooth and error-free user experience during initial configuration.

Reasoning: A comprehensive setup process is crucial for system usability and functionality. Nielsen (1993) emphasises that a well-designed, comprehensive design minimises user errors and reduces cognitive load. By ensuring that all essential steps are present, users can configure complex systems, such as smart home automation devices, correctly and efficiently. This thoroughness is particularly important in preventing misconfigurations that could lead to user frustration or system malfunctions.

Rating Standards:

- **5:** All essential steps are included, with thorough explanations, enabling users to complete the setup without external assistance. The process anticipates potential issues and offers solutions.
- **4:** Most steps are included with good documentation. Some minor details may be missing but do not significantly hinder the setup process.
- **3:** The setup is largely comprehensive, with a few missing steps or details. Users can complete the setup without major difficulties.
- **2:** Only the basic steps are present, lacking sufficient detail, which may cause users to require external help at times.
- **1:** The setup is oversimplified, omitting many necessary details, potentially leading to user confusion and errors.

Flexibility and Personalisation

Definition: Flexibility and Customisation gauge the extent to which the system allows users to personalise settings and features to align with their individual needs and preferences.

Reasoning: Providing flexibility and opportunities for customisation is critical to user satisfaction and system adaptability. Norman (2013) discusses the importance of accommodating different user needs through flexible design, highlighting that personalisation can enhance user engagement and satisfaction. In adaptive user experiences, enabling users to tailor their experience is vital for creating meaningful and personalised interactions.

Rating Standards:

- **5:** Extensive customisation options are available, allowing users to completely adjust the system to their specific needs. Advanced settings and preferences are easily accessible.
- **4:** High level of customisation options similar to the top rating but missing a few advanced settings.
- **3:** A decent range of customisation options is provided, enabling users to personalise the main aspects of the system.
- **2:** Limited customisation options are available, offering basic flexibility but missing several advanced features.

- 1: Minimal customisation choices; users can change some fundamental settings but cannot significantly personalise the system beyond that.

User Control and Privacy

Definition: User Control and Privacy assess the degree of control users have over their personal data and the transparency regarding privacy measures implemented within the system.

Reasoning: Ensuring user control over personal data and maintaining transparency about privacy practices are fundamental to ethical design and fostering user trust. Cavoukian (2010) introduces the concept of "Privacy by Design," advocating for systems that support users to manage their own data and understand how it is used. In the context of smart home automation, where personal data may be collected and processed, prioritising user control and privacy is essential to build confidence and reduce barriers for user adoption.

Rating Standards:

- 5: Extensive user control and privacy features are provided, with full transparency and control over data usage. The system communicates clearly about how data is used and offers comprehensive privacy protections.
- 4: Great user control and privacy features, with strong transparency and data handling practices, though slight improvements could be made.
- 3: Good level of user control and privacy options, although communication about data usage could be clearer.
- 2: Basic user control and privacy features are available but lack depth and detailed communication.
- 1: Minimal user control and privacy, with significant transparency and control issues.

Ease of Use

Definition: Ease of Use represents the system's user-friendliness and intuitiveness, focusing on simplicity and the provision of clear instructions.

Reasoning: Ease of use is a fundamental principle in usability engineering and well documented in decades of HCI literature. Nielsen (1993) asserts that systems should be easy to learn and operate, reducing the time and effort required for users to achieve their goals. The design of intuitive interactions is particularly critical in smart home automation experiences where users may interact with more complex technologies. Despite the designers output not including user interfaces, by prioritising aspects that support ease of use in the UX flows, designers can enhance user satisfaction and promote efficient task completion.

Rating Standards:

- 5: Exceptionally user-friendly flow, with clearly marked steps, indicating that users can navigate the setup without any external assistance.

- **4:** Highly user-friendly and straightforward flow. Minor improvements could make the process even more intuitive.
- **3:** User-friendly flow with clear instructions, though some steps might be complex for non-technical users.
- **2:** Somewhat user-friendly but with enough complexity in the flow that may confuse users.
- **1:** Difficult to use, with many confusing elements that require significant effort to understand.

Communication with Engineers

Definition: Communication with Engineers evaluates the clarity and effectiveness of the documentation provided by designers and the extent to which it facilitates collaboration between UX designers and engineers, ultimately supporting optimal design implementation.

Reasoning: Effective communication between UX designers and engineers is critical for successful design implementation. Collaborative design is critical in modern digital product development and the need for clear documentation between design and development teams is well known in the industry. Good communication helps to ensure that engineers understand the design intent and can accurately translate it into functional code.

Rating Standards:

- **5:** Exceptional documentation that fully communicates the design specifications and intentions.
- **4:** Excellent documentation and communication, with minor gaps. However, one can easily comprehend what is required.
- **3:** Good documentation, although some details might need clarification. Communication is effective but not optimal.
- **2:** Basic documentation that could lead to ambiguity when communicating with engineers.
- **1:** Minimal standard of documentation and communication, which could lead to frequent misunderstandings and the need for clarifications.

Note: Although no development engineers were directly engaged in this study, the assessment of 'Communication with Engineers' remains rigorous by evaluating the UX flows against established best practices for design documentation and clarity. By focusing on the quality and completeness of the designers' documentation – such as the inclusion of sufficient detail, clear annotations, and logical structuring – I can assess the communication potential for each UX flow.

7.5.3. IoT Expert Assessment

I engaged an IoT expert, Dr Massimiliano Dibitonto, to provide an independent assessment of the technical feasibility and practicality of the designers' outputs. His expertise offered valuable insights into whether the proposed UX flows could be implemented effectively using current IoT technologies.

In this proof-of-concept research project, one domain expert was deemed sufficient because the study sought to establish technical feasibility rather than provide a fine-grained comparison of alternative solutions.

At this early stage of validation, the key question was binary: can the proposed user-experience flows be implemented with off-the-shelf Consumer-IoT hardware and software that are widely available today? For such 'good enough' assessment, one expert is considered sufficient under ISO/IEC/IEEE norm 15288:2023 **'Systems and software engineering – System life cycle processes'**.

The ISO/IEC/IEEE 15288 (2023) recommends that organisations make use of one of more assessment methods, including expert judgement, and explicitly notes that an organisation 'can be as small as a single individual, if the individual is assigned responsibilities and authorities' ISO/IEC/IEEE 15288 (2023). Consequently, an organisation comprised of one expert may conduct a valid expert assessment of a given solution.

The practice is not without precedent: the use of a single expert in early-stage design has been documented in both research and industry contexts (Albeshier, 2023; Habib and Cranor, 2022; Harley, 2018; Weber et al., 2019).

A further practical and methodological consideration concerned the relative scarcity of experts whose knowledge sits specifically at the intersection of Consumer IoT and design methodology, rather than solely on the technical engineering side of IoT. This distinction mattered for this study because the artefacts under review were UX flows and design concepts rather than implemented systems. Dr Massimiliano Dibitonto's profile was particularly well matched to the evaluative task as it combines substantial IoT expertise with recognised experience in UX and the design process in the context of connected products.

Dr Dibitonto holds a PhD in Computer Science and is specialised in IoT and UX design. With over 15 years of experience, he has been involved in developing innovative IoT solutions, including energy management systems and electric vehicle charging infrastructure.

In addition to providing his feasibility assessment for each of the UX designers' output, Dr Dibitonto kindly completed a 13-question survey, aimed at providing a deeper understanding of his experience and expertise. The full questionnaire can be found in appendix H.

Assessment Procedure

The expert assessment was conducted via a scheduled Microsoft Teams meeting, lasting approximately one hour and took the following structure:

Introduction and Briefing

I provided an overview of the study, the objectives of the assessment, and the materials to be reviewed. I gave the expert access to the Miro boards containing the pattern documentation, scenarios, and designers' outputs.

Review of Designers' Outputs

For each participant's UX flow, the expert:

- Examined the proposed flow: Reviewed the user flow diagrams, considering the sequence of interactions.
- Considered Technical Aspects: Assessed the feasibility of implementing the proposed interactions using current IoT and ML technologies.
- Identified Potential Challenges: Noted any issues related to scalability, interoperability, user control, privacy, or complexity.

Feedback Collection

The expert provided verbal feedback for each design, which was recorded (with consent) for analysis.

I then examined the feedback and comments from the IoT expert, focusing on his assessment of the feasibility and potential impact of these patterns. These insights are discussed in this thesis, and some of them will inform future updates to the formal documentation of the patterns.

7.6. Data Collection

I collected data through multiple channels to ensure a rich dataset for analysis:

Design Artefacts

The user flow diagrams created by the participants were captured within the Miro boards. These artefacts included visual elements, annotations, and any textual descriptions added by the designers. I exported and saved the Miro boards for analysis.

Video and/or Audio Recordings and Transcriptions

With participants' consent, I recorded the workshops to capture the dialogue accurately.

Observational Notes

During the sessions, I took observational notes, documenting any technical difficulties encountered, and notable behaviours such as confusion or enthusiasm.

Survey Data

The surveys provided demographic and background information on the participants.

Data Management

All collected data were stored securely in the University's One Drive. I organised the data systematically, with each participant assigned a unique code (D1–D5) to anonymise their contributions.

Statistical Analysis methods

I used simple descriptive statistics together with an exploratory correlation analysis to summarise the rubric-based ratings of the five UX flows. These techniques were not intended to provide confirmatory statistical evidence, but rather to support interpretation of the small dataset by highlighting broad tendencies and possible relationships between criteria.

Data Collection

I conducted a systematic rating for each UX flow based on the rubric criteria to ensure uniformity and impartiality during my evaluation. The result was a small dataset with the ratings of all five criteria, for each one of the five participants' output.

Descriptive Statistics: Calculation of Central Tendency

Given the sample size and nature of the dataset, the median rating for each criterion was calculated. The median, representing the middle value in a sorted list of ratings, was chosen as a measure of central tendency.

Formula (in this case it is an odd number of ratings)

$$\text{median} = \left(\frac{n + 1}{2} \right)^{th} \text{ term}$$

This approach provided a more holistic view of the performance consistency in UX flows.

Correlation Analysis

Correlation analysis was used to clarify relationships between different evaluation criteria, aiming to understand whether variations in one criterion relate to others. This analysis highlights interdependencies that could warrant further investigation from an execution perspective.

Calculation of Correlation Coefficients

Spearman's rank correlation coefficient was calculated to measure relationships between criteria. This non-parametric measure assesses how well the rankings of one criterion correspond to the rankings of another, even if the relationship is not linear.

Spearman's rank correlation formula⁴⁰:

$$\rho = 1 - \frac{6\sum d_i^2}{n(n^2 - 1)}$$

ρ = Spearman's rank correlation coefficient

d_i = Difference between the two ranks of each observation

n = Number of observations

Statistical Significance

I also report the p-values associated with the observed correlations as a tentative aid to interpretation. Formula used for the statistic test:

$$t = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}}$$

t : This is the **t-statistic**, which is the value used to compare against a t-distribution to determine whether the correlation is statistically significant.

⁴⁰ Source: <https://www.simplilearn.com/tutorials/statistics-tutorial/spearmans-rank-correlation>

r : This is the **correlation coefficient** between two variables. It ranges from -1 to 1, where 0 indicates no correlation, 1 indicates perfect positive correlation, and -1 indicates perfect negative correlation.

n : This represents the **sample size**, or the number of data points.

Although inferential tests can highlight potential relationships among variables, given the very small sample size, these values are used only in an exploratory sense and are not treated as a basis for confirmatory inference. They are included to indicate which apparent associations in this dataset may be worth examining in future studies with larger samples.

7.7. Results

In this section, I present the validation study results, organised into four parts. First, I report the screening survey results, which provide an overview of the participating designers' professional backgrounds and experience. I then provide a general account of the results of my evaluation of the UX flows generated by the participants who used two novel UX patterns: **Preference Learner** and **Geographic Trigger**. This includes a summarised overview of my ratings across multiple evaluation criteria and a summary of the technical feasibility assessment conducted by an independent IoT expert. Next, I go into the detail of each individual UX flow produced by the designers, offering a thorough rationale for the ratings assigned to each criterion. Finally, I present the detailed qualitative feedback from the IoT expert, which substantiates the technical feasibility of the proposed solutions.

7.7.1. Screening Survey Results

To join the study, participants had to successfully complete a screening survey. The survey responses provided a comprehensive overview of the professional backgrounds and experiences of the participating UX designers.

The participants were given codenames D1 through D5. The results of the survey can be seen in Table 6.

The majority of the participants have extensive experience in UX design, with three designers having more than a decade of experience. This indicates a high expertise level among the participants, which is crucial for the main task of using and providing feedback on the novel UX patterns. It is relevant to note that despite their long experience, only two designers had prior involvement in IoT projects.

Only two out of five designers had experience with machine learning in a UX design project, and of those designers neither of them was involved in any stage of planning or designing the role of machine learning.

Usage of Design Systems

All participants indicated that they currently use a design system. This is relevant as Design Systems are one of the most recent embodiments of patterns and pattern theory in the industry. The use of custom-made design systems by most designers may suggest the importance of having tailored design systems to address organisations' specific needs and constraints in digital product design. In addition, it could potentially be associated with the common challenges with the adoption and customisation of open-source component libraries to fit into companies' existing technology stack. However, the sample is not large enough to draw confident conclusions.

Codename	Age group	Experience	Involvement in IoT projects	Involved in projects that used ML?	Involved in deciding the role of ML?	Currently using a design system?	Open source or Custom made?
D1	30-34	5-10 years	No	No	N/A	Yes	Custom made
D2	35-39	10+ years	No	No	N/A	Yes	Open source
D3	45-49	10+ years	Yes	No	N/A	Yes	Custom made
D4	45-49	10+ years	No	Yes	No	Yes	Custom made
D5	30-34	5-10 years	Yes	Yes	No	Yes	Custom made

Table 6 - Results of the validation study's screening survey

7.7.2. Quantitative Analysis of UX Flow Evaluations

In this subsection, I present descriptive quantitative summaries of the five UX flows. Using the evaluation rubric described earlier in this chapter, I report the ratings across the five evaluation criteria, together with median values and an exploratory Spearman correlation analysis. Within this proof-of-concept study, these quantitative results are not treated as confirmatory evidence; rather, they are used to complement the qualitative analysis by identifying broad tendencies in the ratings and possible relationships that may warrant investigation in larger future studies.

Three designers (D1, D2, and D3) created user experience flows using the **Preference Learner** pattern and two designers (D4 and D5) created user experience flows using the **Geographic Trigger** pattern. A rating of 1 signifies the lowest standard possible, minimally meeting the criteria whereas 5 represents excellent level and meeting all the expectations for a given criterion. The evaluation rubric is described in Section 7.5.2.

Table 7 below shows the overall ratings for each criterion of all UX flows produced by the designers:

	Using Preference Learner pattern			Using Geographic Trigger pattern	
	D1	D2	D3	D4	D5
Comprehensiveness	4	4	3	4	4
Flexibility and Customisation	5	5	5	5	3
User Control and Privacy	4	5	3	3	4
Ease of Use	3	4	4	4	4
Communication with Engineers	4	5	3	4	4

Table 7 - Summary of the ratings for the UX Flows created by the designers. Each coded between D1-D5 to maintain the anonymity of the participants.

The overall ratings overlaying all five designers' output can be seen in the radar chart below. None of the solutions achieved the highest score in all criteria, however none of them received the lowest score for any criterion.




Figure 26 - Overlaid ratings for all UX flows produced by the participants – none received the lowest rating in any criterion.

Descriptive Analysis of UX Flow Ratings Results

The following table summarises the median ratings for each evaluation criterion:

Criteria	Median
Comprehensiveness	4
Flexibility and Customisation	5
User Control and Privacy	4
Ease of Use	4
Communication with Engineers	4

Table 8 - Median values for each criterion.

Across this small sample, the median ratings cluster between 4 and 5. This indicates that, within the rubric used in this study, the five outputs were generally assessed positively across the evaluated criteria. However, these values should be read as descriptive summaries of this particular cohort and task setting, rather than as evidence of broader performance or generalisable effectiveness.

Correlation analysis

The correlation matrix shown in Table 9 quantifies the strength and direction of the relationships between the evaluation criteria using Spearman correlation coefficient. Values range from -1 to 1, with -1 indicating a perfect negative correlation, 0 indicating no correlation and 1 indicating a perfect positive correlation.

Criteria	Comprehensiveness	Flexibility and Customisation	User Control and Privacy	Ease of Use	Communication with Engineers
Comprehensiveness	1.00	-0.25	0.56	-0.25	0.79
Flexibility and Customisation	-0.25	1.00	-0.19	-0.25	0.00
User Control and Privacy	0.56	-0.19	1.00	-0.19	0.82
Ease of Use	-0.25	-0.25	-0.19	1.00	0.00
Communication with Engineers	0.79	0.00	0.82	0.00	1.00

Table 9 - Spearman Correlation Matrix for the evaluation criteria

Statistical Significance (P-Values)

The table below reports the p-values corresponding to the Spearman rank coefficients. Once again, it is critical to reiterate that given the small sample size, these values should be interpreted with considerable caution and only as rough exploratory indicators. A 10% threshold is shown here solely to flag possible tendencies within this limited dataset, not to support strong inferential claims.

Criteria	Comprehensiveness	Flexibility and Customisation	User Control and Privacy	Ease of Use	Communication with Engineers
Comprehensiveness	0.00	0.69	0.33	0.69	0.11
Flexibility and Customisation	0.69	0.00	0.76	0.69	1.00
User Control and Privacy	0.33	0.76	0.00	0.76	0.09
Ease of Use	0.69	0.69	0.76	0.00	1.00
Communication with Engineers	0.11	1.00	0.09	1.00	0.00

Table 10 - P-values indicating the statistical significance of the Spearman's rank correlation coefficients.

With only five cases, the coefficients are highly sensitive to small shifts in the ratings and should not be regarded as stable estimates. At most, the analysis points to possible associations within this dataset. In particular, the positive association between User Control and Privacy and Communication with Engineers may be of interest for future work, but even here the evidence remains tentative and exploratory. The remaining coefficients are best read as descriptive indications of how the ratings happened to vary together in this small proof-of-concept sample, rather than as evidence of reliable underlying relationships.

7.7.3. Qualitative Analysis: Individual UX Flow Case Studies

In this section, I present detailed evaluation of the UX flows produced by each designer using the novel UX patterns. For each designer (D1 through D5), I provide their proposed UX flow (with corresponding artefacts) followed by the detailed rationale of my evaluation.

D1's Proposed User Experience Flow – Using Preference Learner Pattern




Figure 27 - D1's proposed user experience flow.

D1's User Experience Flow Evaluation

Criteria: Comprehensiveness of Setup

The proposed flow demonstrates a well-structured setup process that covers essential aspects of device synchronisation, permissions, and configuration. The flow begins with device syncing (e.g., smartwatch, smart lights), followed by permissions and options for activating routines based on triggers, such as heart rate (HR) and time. This approach indicates a clear initial setup considerations, anticipating necessary configurations for successful operation. However, some steps, particularly around initial connectivity checks and permissions, could benefit from more detailed guidance to ensure completeness for a less tech-savvy user. The provision of specific error-handling instructions or troubleshooting tips could elevate this section to a full score.

Rating:

4

Criteria: Flexibility and Personalisation

The UX flow exhibits plenty of flexibility and personalisation options, allowing users to customise routines according to individual preferences, such as selecting workout types, relaxation routines, and setting goals (e.g., exercise frequency). The presence of "Opt-out" and "Partial opt-out" features enhances user autonomy, enabling individuals to fine-tune the learning period – this suggests a very adaptable experience, meeting different user needs while providing clear personalisation mechanisms.

Rating:

5

Criteria: User Control and Privacy

The design accounts for user control and privacy through various elements, such as explicit options for complete or partial opt-out from features. However, more detailed communication around data usage, specifically how data is collected and what it is used for would enhance transparency and user trust. An addition of more granular controls for data-sharing preferences could provide users with a better sense of privacy control, leading to a higher rating.

Rating:

4

Criteria: Ease of Use

While the flow includes user-friendly elements (like tips of what some options mean), the overall interaction structure may present complexity to those unfamiliar with smart device interactions. For example, interactions expecting user input on aspects such as Interoperability with other devices and/or applications and 'Learning' Progress review may confuse less tech-savvy users or those unfamiliar with terminology associated with automation and interconnected devices. It was not clear whether users would be provided with simpler descriptions for each step or if cues/tooltips or any form of help would be embedded – which could improve the intuitiveness of the flow.

Rating:

3

Criteria: Communication with Engineers

The documentation of the UX flow is comprehensive and clear, with logical structure and well-defined stages. This clarity is likely to facilitate effective communication with engineers responsible for implementation. Each phase of the user flow – from device syncing to preference configuration and notification management – is outlined with sufficient information for engineers to interpret the designer's intent. However, the flow could benefit from more explicit annotations indicating specific technical requirements or constraints, such as dependencies on third-party platforms.

Rating:

4

D2's Proposed User Experience Flow – Using Preference Learner Pattern




Figure 28 - D2's proposed user experience flow.

D2's User Experience Flow Evaluation

Criteria: Comprehensiveness of Setup

D2's UX flow presents a well-structured setup process that encompasses critical stages such as initial configuration, activation of learning modes, and information about the ongoing learning option. Like the flow presented in D1, it effectively covers essential elements. It also includes notifications to users about data tracking, establishing device connectivity, and emphasising privacy control options. These aspects collectively represent a solid example of requisites for a high-quality user experience.

However, like observations made in D1's evaluation, there is room to enhance the setup's comprehensiveness, by offering more detailed guidance for users who may require additional support. For example, after establishing device connectivity, the flow could include specific instructions or prompts to assist those unfamiliar with smart home automation systems. Incorporating troubleshooting tips or context-sensitive help at this stage would ensure a smoother onboarding experience, catering also to users with varying levels of technical proficiency without overwhelming them with technical jargon.

Rating:

4

Criteria: Flexibility and Personalisation

The UX flow offers extensive flexibility and personalisation options, enabling users to tailor the system to their preferences. The 'Learning Mode' allows users to decide when and how the system learns from their behaviour, giving them control over the personalisation process – an

Rating:

5

approach that matches the high level of adaptability noted in D1's design.

Additionally, the UX flow considers choice between a local network setup and cloud-based configuration and poses questions about the role of the user amongst other users in the household. The combination of these features demonstrates a commitment to user-centric design, particularly considering the Consumer IoT context.

Criteria: User Control and Privacy

User control and privacy are prominently emphasised throughout D2's UX flow. The design ensures users are informed about behaviour tracking and continuous learning capabilities, providing options to manage data handling – including options such as local or cloud storage. This contextual privacy information enhance transparency, fostering trust. These measures align with privacy by design principles, ensuring users retain autonomy over their data and understand its utilisation, demonstrating an ethical approach to UX design.

Rating:

5

Criteria: Ease of Use

The UX flow in general deliver a user-friendly experience, featuring clear interaction points like 'Learning Mode' activation and periodic privacy control reminders, reducing cognitive load by making system interactions straightforward. However, certain aspects could be refined to enhance usability for non-technical users. This flow has a particular emphasis on connectivity, but it does not clearly communicate the steps to break the process down to end users, nor it mentions error handling – critical in these types of tasks to enable user recovery and next steps.

Rating:

4

Criteria: Communication with Engineers

The documentation provides enough detail for effective communication with engineers, featuring clear steps and logical structure that cover key technical requirements like device connectivity, learning mode configurations, and data handling options. In addition, it clearly specifies the learning mode requirement and poses technical feasibility questions to explore the possibility of an ongoing learning functionality. Finally, it uses a more formal diagrammatic notation, making it easier for engineers.

Rating:

5

D3's Proposed User Experience Flow – Using Preference Learner Pattern

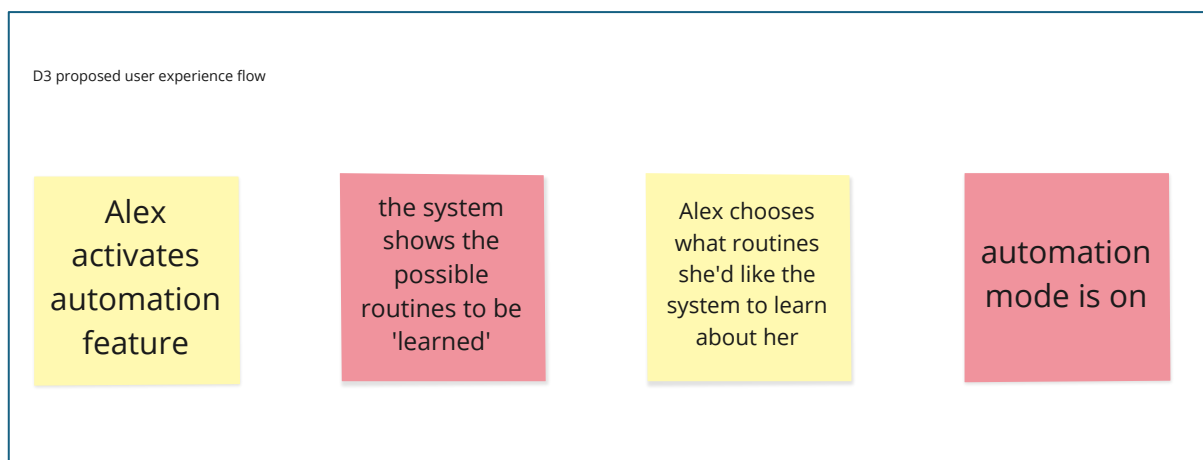


Figure 29 - D3's proposed user experience flow.

D3's User Experience Flow Evaluation

Criteria: Comprehensiveness of Setup

D3's UX flow presents all key steps in the logical sequence that would allow Alex to activate the automation feature and select routines for the system to "learn" (e.g., morning and evening routines). The sequence includes steps that are the basis for a smooth configuration, such as activating the feature and selecting specific routines, which align with the goals of the scenario provided for the task. However, details on feedback mechanisms and alternative interaction paths would make the setup flow more realistic and useful. Thus, while the setup would support many relevant discussions in the typical next steps of the design process, this UX flow could do with significant enhancement.

Rating:

3

Criteria: Flexibility and Personalisation

D3's UX flow allows Alex to select which routines the system should learn, providing a high level of personalisation, which aligns well with the Preference Learner pattern's adaptive goals. By letting Alex choose specific modes, D3's solution enables a flexible approach to automation, which is a key strength. The setup also considers the user's control over the system's learning scope and does not impose default or rigid structures. This is a good example of flexibility in UX design, meriting a score of 5.

Rating:

5

Criteria: User Control and Privacy

This UX flow does include basic user control by allowing Alex to select routines; however, it lacks explicit considerations for privacy and data transparency. For instance, the flow could inform Alex about what data is collected during the "learning" process or by offering options to modify or pause learning modes. Since privacy and transparency are increasingly central to user trust in systems using machine learning, D3's flow would benefit from making these controls visible. Therefore, while the flow addresses some user control, it falls short on privacy-related aspects.

Rating:

3

Criteria: Ease of Use

The UX flow is straightforward and intuitively structured, beginning with the activation of automation and progressing logically through routine selection and automation initiation. In theory, if executed through user interfaces, users with minimal experience could follow the sequence without major hurdles, as each step flows into the next naturally. However, incorporating contextual details – such as privacy options and configurations for ongoing learning for example – would create an explicit step which would improve the usability. Overall, the flow achieves an easy-to-follow design but could benefit from adjustments.

Rating:

4

Criteria: Communication with Engineers

The documentation provided by D3 for the UX flow is generally clear but there is room for improvement when it comes to technical specificity. While the flow represents the setup sequence well, it lacks some details on steps or notes covering potential user behaviour and options. For instance, including notes on data handling, triggers, and error recovery would facilitate better understanding of technical requirements. Consequently, while functional when presented by the UX designer, the documentation could be improved to support collaboration with engineers.

Rating:

3

D4's Proposed User Experience Flow – Geographic Trigger Pattern




Figure 30 - D4's proposed user experience flow.

D4's User Experience Flow Evaluation

Criteria: Comprehensiveness of Setup

D4's UX flow offers a fairly comprehensive setup process, guiding users through clear stages when creating new automations. Users begin by assigning a unique name to the automation, followed by specifying location-based triggers and selecting devices or actions, such as controlling lights or heating systems. The option to add additional trigger conditions and define specific outcomes enhances the depth of configuration, accommodating more complex user requirements. However, the flow lacks mention of an explicit confirmation or feedback mechanisms to communicate successful/failure and recover routes when relevant. Incorporating these aspect could improve the user experience further.

Rating:

4

Criteria: Flexibility and Personalisation

The UX flow communicates extensive flexibility and personalisation. Users can customise numerous aspects of the automation, from specific the geographic location as a trigger to specifying device actions and setting specific conditions. Additional control is introduced through the selection of multiple limiting conditions, including date, time, and specific aspects like temperature. This wide array of options enables users to tailor automations to their preferences, aligning with the principles of adaptive and context-aware interactions in Consumer IoT systems – a strong user-centred perspective.

Rating:

5

Criteria: User Control and Privacy

While D4's flow allows users to adjust triggers and specify conditions, it offers limited visibility of privacy and data control measures. There is a lack of clarity about how location data will be used, stored, or accessed, which may raise concerns for users sensitive about their personal information. The absence of explicit privacy settings or transparency elements could undermine user trust. Enhancing the flow by incorporating privacy options and clearly communicating data handling practices would improve user control and align the design with higher standards of privacy assurance.

Rating:

3

Criteria: Ease of Use

Rating:

4

The UX flow demonstrates ease of use through its step-by-step approach, guiding users logically through defining triggers, conditions, and actions. Each segment follows a sequential order, supporting users in completing the setup without complexity. The inclusion of features like “favourites” or frequently used actions simplifies decision-making and promotes efficient interaction. To stay consistent to the level of detail documented (e.g. ‘clicks + or add CTA’), this UX flow could be further enhanced by providing indications of inline assistance or tooltips during configuration steps - this would assist novice/first time users. In addition, no adaptive functionality was explicitly explored, although predictive models are likely to be used at some point to calculate travel times. Overall, the flow is usable and largely self-explanatory, which contributes positively to the user experience.

Criteria: Communication with Engineers

Rating:

4

D4's UX flow is overall effective in conveying the design to engineers, featuring distinct stages and a logical arrangement that clarifies user actions and system responses. However, the UX flow does not explicitly address critical error handling, such as what happens if the location data is temporarily unavailable. The flow also lacks considerations for feedback from the system, which combined with error handling becomes essential do design recovery interactions, which need to be supported technically. Enhancing the flow with these specifics would improve communication with the engineering team.

D5's Proposed User Experience Flow⁴¹ – Geographic Trigger Pattern




Figure 31 - D5's proposed user experience flow.

⁴¹ D5 initially confused the persona's name, calling her 'Alexa' instead of 'Alex'

D5's User Experience Flow Evaluation

Criteria: Comprehensiveness of Setup

D5's UX flow presents a solid process for users like Alex to configure their smart home systems effectively. The process begins logically with app installation and profile creation, followed by customisation of settings related to heating and lighting. These steps align well with the goal of enabling geographic triggers for home automation. However, could be enhanced by providing more explicit instructions on linking specific device settings to location triggers. Additionally, accounting for potential setup errors or issues would make the process more robust, but instead the initial part of the flow covers purchase decision and product evaluation.

Rating:

4

Criteria: Flexibility and Personalisation

While D5's flow allows for some customisation – such as adjusting time, lighting, and heating preferences based on location – it offers limited personalisation options. The ability to set more complex conditional triggers is not evident. Expanding the flow to include some aspect of adaptive behaviour using preferences and user data could significantly enhance personalisation for example. This expansion would enable users to tailor the system more closely to their lifestyles and needs, while reducing the amount of user input, therefore also improving the flexibility of the UX flow.

Rating:

3

Criteria: User Control and Privacy

The UX flow demonstrates a mindful approach to user control and privacy by prompting Alex to enable location services and explaining why it is needed – ensuring users are informed about data usage. However, the flow could further improve user control by offering additional privacy options, such as allowing users to set data retention durations or manage permissions at a more granular level. These enhancements would increase users' sense of control over their personal information.

Rating:

4

Criteria: Ease of Use

D5's UX flow is straightforward and user-friendly, guiding users through each step from app installation to setting up geographic triggers in a logical manner. This approach reduces cognitive load and facilitates ease of navigation without the need for external assistance. To further enhance intuitiveness, the flow could mention contextual tooltips or interactive walkthroughs, especially when setting up geographic triggers for multiple devices. Such features would support users with varying technical skills, making the system more accessible and user-friendly.

Rating:

4

Criteria: Communication with Engineers

The documentation of D5's UX flow provides a functional overview that allows engineers to interpret the basic setup requirements, including enabling location services and linking device actions to geographic triggers. However, the flow could benefit from more detailed annotations. For example, including edge cases and exception scenarios – what should happen if connection to smart devices is not available? Connectivity is a critical aspect of Consumer IoT experiences.

Rating:

4

7.7.4. The IoT Expert's feasibility Assessment

To evaluate the technical feasibility of the UX flows proposed by the designers, I sought the independent assessment of an IoT expert. He provided insightful feedback grounded in both technical understanding and user-centric considerations.

Including an independent feasibility assessment guard against the risk of designer or author blind spots and anchors the proposed UX flows in real-world engineering constraints. By inviting an external IoT specialist – who was neither involved in the UX practitioners' workshops nor invested in the research outcomes – I gained an impartial verdict on whether the concepts could be achieved with today's Consumer-IoT hardware and software and a view on any key implementation considerations. This independent perspective further strengthens the results credibility, ensuring that the pattern language is both theoretically robust and applicable in practice.

The expert affirmed that, from a technical standpoint, **all proposed UX concepts are feasible to implement**. He recommended Leveraging existing IoT ecosystems, such as those compatible with Amazon's Alexa for example, to mitigate interoperability issues by adhering to established standards and application programming interfaces (APIs). This approach enables the integration of different devices and services, simplifying the practicality of adoption of such systems in real-world scenarios.

About the expert

I was privileged to have Dr Massimiliano Dibitonto to evaluate the technical feasibility of the solutions proposed by the designers. In this document, Dr Dibitonto is also referred as 'The Expert'.

Dr Dibitonto has 15 years in UX Design/Research and related fields. Co-authored "IoT Design Method: A Method for the Co-Design of the UX of Connected Products"⁴², which led to the production of The IoT Design Deck – included in the 2023 ADI Design Index by the Industrial Design Association which catalogues excellence in Italian design.

⁴² <https://www.iotdesigndeck.com/>

The Feasibility Assessment

Below I provide the detailed account of the results for Dr Dibitonto's assessment of technical feasibility for the design solutions produced, including his feedback.

Solutions using the Geographic Trigger pattern:

The expert pointed out that this automation approach is practical to implement, and real-world examples can be found. The feedback provided to this group of solutions emphasised simplifying the initial setup, providing user control over automated actions, and considering multi-user environments.

He emphasised the need for feedback mechanisms that allow users to adjust or override automated actions, particularly when their plans change unexpectedly.

"I would specify the feedback and the possibility to change the condition (...) Maybe I'm approaching, and I receive a feedback notification that says, 'Okay, I'm starting to heat up your home.' But from that notification, I could change my mind. I could change the automation."

Careful balance needs to exist to avoid excessive notification, however. Critical stages in the user's routine need to be identified so that feedback can enable users to change their preferences. He recognises the challenges around identification of these critical stages to achieve the right balance, given that every notification demands attention from the user. He added that so far, he has not yet found a 'perfect pattern' addressing this problem.

Solutions using the Preference Learner pattern:

Although all solutions using Preference Learner were deemed technically feasible, the expert notes that they were more complex and that it brings subjective aspects to the user experience. For example, each user has unique routines and preferences. What one user finds convenient or desirable, another might find intrusive or unnecessary. Users may have different comfort levels with a system that learns from their behaviours and makes autonomous decisions.

The feedback he provided for this group of solutions focused on managing subjectivity in user preferences, ensuring transparency in data usage, offering adjustable automation levels, and simplifying system complexity through integration.

The D1, D2 and D3 designs need to consider the multi-users situations more explicitly. The key design challenge is to recognise when to activate automation for groups or individuals. These solutions need to add this variable, potentially relying on an algorithm that determines the prioritisation of the relevant needs.

He emphasised the significance of transparency regarding data collection and usage, especially since the **Preference Learner** pattern relies on monitoring user behaviours.

There are conspiracy theories, for example, about devices listening to conversations. Connecting multiple devices is an attractive idea to enable more personalised automation, however, user trust is needed. Transparency practices play a key role in this by explicitly stating which kind of data is being collected and providing the opportunity to opt out.

Overall comments and recommendations useful for both cases

Consider that there is a large probability that the system will provide the wrong guess for what you would prefer in the future. In that case, you need to provide opportunity to fix that incorrect prediction.

The expert highlighted the importance of integrating devices into some form of hub to reduce complexity, suggesting that a single gateway device could manage multiple connected devices. Such a gateway device, such as Amazon Alexa for example could be used to facilitate the communication of the preferred levels of automation. This could save the users from informing these preferences in every vendor specific app.

"I imagine having a single device that is able to control all other devices, or a gateway. And this, should be the gateway that I use to control the level of automation."

He proposed allowing users to select their desired level of automation in a sliding scale, ranging from manual control to full automation, enhancing user comfort and trust.

In Consumer IoT, a key challenge is the UX and usability of the initial setup. Considerable effort goes into the initial configuration and understanding of the system. That is the case because one must 'forecast the future' to feed that initial information to the system (e.g. a heating schedule for winter). He mentioned the concept of Human-centred AI (Shneiderman, 2022), which promotes the balance

between automation and control. Good defaults, he points out, can represent the difference between appliance and tool. Tool demands effort from the user. If the system initial configuration is the most useful to the user, users are unlikely to change it.

Dr Dibitonto suggested enhancing the user experience by utilising machine learning to propose the most probable choices, thereby reducing the initial configuration effort.

"I think that an add-on here could be using machine learning to see what other users have chosen and have the most probable choices proposed by the system (...) So the system helps you to make the configuration, giving you already some list of alternatives that are the most probable ones."

7.8. Limitations

Sample Size and Generalisability

One of the primary limitations was the small sample size of five UX designers. While this enabled a deep exploration of each unique experience, it limits the generalisability of the findings. The results risks not being representative of the wider UX designer population, particularly those with different levels of experience or from different cultural backgrounds.

Time Constraints on the Design Task

The 20-minute time limit for the design task may have constrained the depth and quality of the participants' outputs. In professional settings, designers typically have extended periods for ideation, research, and iteration. The limited time may not have allowed participants to fully engage with the patterns or develop comprehensive solutions.

Lack of End-User Evaluation

As the UX flows did not evolve into any form of UI representation or any more user-friendly output, the study did not include usability testing or feedback from end-users. As a result, the findings are based on the designers' outputs and the expert's technical assessment. For future work, expanding the study into the next design phases would enable incorporating end-user evaluations and can provide further insights into the usability and acceptance of the resulting final designs by users.

Single Expert Assessment

Relying on a single IoT expert may introduce individual biases and limit the scope of technical perspectives. Multiple expert assessments could improve the reliability and depth of the technical evaluation, offering a more comprehensive understanding of feasibility and implementation challenges.

Researcher Bias

As the sole researcher conducting the study, there is a risk of unconscious bias influencing data collection, analysis, and interpretation. Although I took measures to mitigate bias, such as using structured mechanisms, it is not possible to guarantee complete objectivity.

Remote Workshop Challenges

Conducting workshops remotely introduced potential limitations, including technical issues such as connectivity problems and audio disruptions. For instance, during the sessions, there were occasions when we experienced audio problems with Microsoft Teams. Remote settings also reduced my ability to observe non-verbal cues and body language which are often crucial in understanding participants' reactions and engagement levels. Additionally, participants might have faced potential distractions in their environments, which could have affected their focus and the quality of interaction. These factors may have influenced the designers' ability to fully engage with the patterns and could have impacted the depth of their design outputs.

Limited Scope of Patterns Tested

Due to time and resource constraints, only two UX patterns – the Preference Learner and the Geographic Trigger – were evaluated in this study. These proof-of-concept findings provide detailed insights into the performance of these specific patterns and will certainly inform changes to all other patterns presented in this research project. However, some of these findings may not be entirely applicable to the other patterns and future research is recommended to expand the assessment to the remaining patterns.

Assumption of Pattern Comprehension

The methodology assumed that participants fully understood the patterns and the associated documentation provided during the workshops. Despite my efforts to explain the materials and encourage questions, misunderstandings or misinterpretations could have occurred. Such misunderstandings could have impacted the effectiveness of the designs produced, as participants might not have accurately applied the patterns as intended. In the future, implementing comprehension checks, such as summarisation tasks, or providing additional

support during the workshops might address this limitation and ensure that participants have a clear understanding of the patterns before commencing the design task.

Absence of Longitudinal Data

The study captured a single interaction with the patterns during the one-hour workshop sessions. This snapshot does not assess how designers' proficiency or perceptions might evolve over time with continued use of the patterns. In practice, designers might become more adept at applying the patterns or may discover new challenges after extended use. A longitudinal study could provide insights into learning curves, sustained utility, and the long-term impact of using such patterns in practice, offering a more comprehensive understanding of their effectiveness over time.

7.9. Discussion

The validation study provides an encouraging proof-of-concept account of how the documented patterns were used by experienced UX designers in a constrained workshop setting. Across the five cases, the resulting UX flows were generally assessed positively against the rubric, particularly in relation to flexibility and customisation, communication clarity, and the inclusion of user-control considerations. These findings suggest that the patterns can function as a useful early-stage design resource, but they remain bounded by the small sample, the short workshop format, and the interpretive nature of the rating process.

For that reason, the descriptive and correlational results should be read cautiously. They point to a small number of possible tendencies within this dataset, especially the apparent association between user-control/privacy considerations and communication with engineers. This tendency may be worth examining in future studies, but the present evidence remains exploratory.

The expert assessment strengthens the practical interpretation of the study by indicating that all five concepts were technically feasible in principle using current Consumer IoT technologies. At the same time, the expert feedback also identified important areas for refinement, including setup complexity, feedback mechanisms, and multi-user scenarios. Taken together, the findings suggest that the pattern language is promising as an early-stage design resource, while also confirming the need for broader and more rigorous future validation.

Reflecting on the outputs produced by the designers, the patterns appear to have helped participants engage with some of the key considerations involved in integrating machine learning into UX design. In particular, the resulting flows

suggest that the documentation prompted attention to issues such as adaptive behaviour, user autonomy, permission management, privacy, and the configuration of triggers and conditions. Within the limits of this proof-of-concept setting, this indicates that the patterns may encapsulate useful design considerations for early-stage concept development.

These findings should be interpreted in light of the study's limitations, especially the small sample size, the brief workshop duration, and the remote setting. Even so, the convergence of rubric-based evaluation, qualitative case analysis, and expert feedback provides an encouraging exploratory basis for the claim that the ML-enhanced UX pattern language may support the design of feasible Consumer IoT interaction concepts. This conclusion, however, should be understood as provisional and bounded by the proof-of-concept nature of the study.

While there are clear areas for further refinement, the patterns represent a promising step towards supporting designers in navigating the complexities of contemporary UX design involving machine learning and Consumer IoT. By documenting and testing these patterns, this research offers an initial resource for designers seeking to create more adaptive, personalised, and user-centred experiences, while also opening productive avenues for future work

8. Conclusions

8.1. Introduction

In this final chapter, I synthesise the outcomes of my research project. As set out in Chapter 1, the aims and motivations of this study led to the formulation of two Guiding Propositions, which were then explored through two research questions that defined the scope and direction of the empirical work. Guiding Proposition 1 (GP1) concerned whether a pattern language encapsulating machine-learning-driven interaction guidance for Consumer IoT could be identified, thereby helping designers conceptualise adaptive and user-centred features. Guiding Proposition 2 (GP2) concerned whether, if identified, such a pattern language could support the creation of feasible Consumer IoT interaction concepts when employed by designers with no deep expertise in machine learning. These propositions were explored through RQ1, which addressed the identification and formal documentation of such a pattern language, and RQ2, which addressed its practical usefulness as a design resource for non-ML expert designers. Across the preceding chapters, I have shown how the machine learning experiments, the pattern mining

study, the formal pattern documentation, and the proof-of-concept validation study together addressed these questions and, in doing so, provided a coherent body of evidence in relation to both guiding propositions. In the following sections, I summarise the research project, discuss its theoretical and practical contributions, reflect on its limitations, and outline directions for future research.

8.2. Summary of the Research Project

My research set out to address a critical gap in the domain of Consumer IoT design: the need for accessible design frameworks that integrate machine learning (ML) in a manner that does not require deep technical expertise from UX professionals. I began by reviewing the literature on UX design, pattern theory (drawing on Alexandrian principles), and the integration of ML into interactive systems. I then undertook a pattern mining study involving experienced UX designers, during which I uncovered and refined initial pattern seeds (early concepts of patterns, essential ideas that form the basis of a pattern), using structured workshops leveraging the “What Would a Human Do” (WWHD) thought experiment. This process led to the identification and subsequent formal documentation of a coherent ML-enhanced UX pattern language, comprising key patterns namely, **Preference Learner**, **Geographic Trigger**, **Stop Learning Preferences**, **Intelligent Access Sharing**, the proto pattern **Smart Shortcuts** and the methodology pattern **Personal Spark**. Finally, I validated the proposed pattern language through a proof-of-concept study in which designers, with limited to no ML expertise, used the patterns to produce UX flows for Consumer IoT scenarios. The resulting flows were evaluated both quantitatively and qualitatively using a detailed rating rubric and an independent IoT expert’s assessment. Taken together, these findings provided a coherent and rigorous exploratory body of evidence relevant to both Guiding Propositions.

8.3. Contributions to Theory and Practice

My research makes contributions to both theory and practice of UX design in the context of Consumer IoT. These contributions emerge from the documentation of an ML-enhanced UX pattern language, and the empirical validation of its application by designers with limited ML expertise. This novel pattern language assists with the integration of AI principles into UX design, and although my work does not include the development of new machine learning algorithms, it provides valuable insights into ML-enhanced interaction design.

8.3.1. Theoretical Contributions:

Application and Operationalisation of Alexandrian Pattern Theory:

While the use of Alexandrian Pattern Theory in digital domains is well documented, my research demonstrates its applicability in a novel context: the integration of

machine learning into UX design for Consumer IoT. I have not sought to extend the fundamental theoretical framework itself, but rather to operationalise it in a way that yields a dynamic pattern language. Importantly, my work highlights that the value of a pattern language lies in the interconnections among its patterns rather than in a fixed number of elements.

Documentation of a Novel, Dynamic and Extendable Pattern Language:

Central to my theoretical contribution is the documentation of key patterns – namely, **Preference Learner**, **Geographic Trigger**, **Stop Learning Preferences**, **Intelligent Access Sharing**, the proto-pattern **Smart Shortcuts**, and notably the **Personal Spark** pattern. The Personal Spark pattern is particularly important because it functions both as a structured ideation strategy and as a mechanism for continued pattern mining. I examined this pattern through a series of one-to-one workshops involving 23 participants, and the findings provide encouraging support for its usefulness in identifying interaction concepts amenable to ML augmentation. This, in turn, offers future researchers a replicable and practically grounded method through which the pattern language may be extended. By operationalising Alexandrian Pattern Theory in this way, I present a pattern language that is capable of evolving alongside emerging technological developments, thereby contributing to a contemporary application of Pattern Theory rather than treating pattern-language documentation as static.

Advancement in ML-Enhanced Interaction Design:

By incorporating relevant AI and human-AI interaction principles into UX design, my research offers insights into how adaptive systems may be conceived for Consumer IoT. Considered in light of the studies reported here, the novel pattern language appears capable of assisting designers in balancing technical feasibility with user-centred design concerns. Within the scope of this thesis, this represents a meaningful contribution to **ML-enhanced interaction design in the Consumer IoT domain**.

8.3.2. Practical Contributions:

A Practical Toolkit for UX Designers

The ML-enhanced UX pattern language offers a potentially useful design resource for UX designers, particularly those with limited ML expertise. In the validation study, participants were able to produce coherent UX flows after only a brief introductory session (approximately 10–20 minutes), followed by a short design task (circa 20-minute long). Within the limited scope of this proof-of-concept study, this suggests that the patterns may be usable as an early-stage aid for concept generation.

Enhanced Interdisciplinary Communication:

By offering a structured and well-documented format, the pattern language may be leveraged to improve communication between engineers, designers and other stakeholders. The documentation generated, emphasising user control, privacy, and technical feasibility, can be used to facilitate interdisciplinary collaboration – an essential factor for successful digital product development, which includes Consumer IoT experiences.

Empirical Insights for ML Integration in Consumer IoT:

The findings gathered from the quantitative evaluation, qualitative feedback, and independent expert assessment provide an initial set of insights into how ML may be integrated into UX design for Consumer IoT. These insights should be understood as exploratory rather than definitive, however they can still help to inform future refinement of adaptive and personalised digital products, potentially even beyond the Consumer IoT realm.

8.4. Limitations and Future Research

Although my research provides evidence in support of the overarching guiding propositions, there are limitations that warrant mention, spanning all stages of the project – from the initial ML experiments and pattern mining approach to the final validation study.

8.4.1. Limitations of the Machine Learning Experiments

Dataset Constraints and Model Transparency:

For the machine learning experiments, I utilised publicly available datasets (such as the Rico dataset) which may not be fully relevant to the Consumer IoT context. Moreover, the technical limitations of tools like the IBM Cloud Annotations platform – particularly the restricted customisation and limited ability to directly tune internal model parameters, constraining the performance of the automated labelling processes.

Simplified Experimental Setup:

Although the user identification experiment achieved a high level of accuracy (95% with a k-nearest neighbours classifier), it was conducted under controlled conditions using data that had already been collected in the UCI dataset.

This simplified experimental set up might not fully reflect the complexities encountered in real-world scenarios, where sensor noise and environmental

variability are more significant and more robust data pre-processing considerations are needed.

8.4.2. Limitations of the Pattern Mining Method

Subjectivity in a Designer-centric Approach

The pattern mining phase was based on a specifically developed, qualitative approach using the **Personal Spark** pattern (which leverages the WWHD thought experiment). While this novel method successfully elicited rich interaction concepts from 23 participants, it is inherently subjective. The particular perspectives, background and expertise of the participating designers may have influenced which pattern seeds were identified and prioritised, potentially introducing bias.

Incomplete Pattern Language

I do not claim that the documented pattern language is exhaustive. Although I have captured an initial set of key patterns, many additional patterns remain to be uncovered and thoroughly documented. The dynamic interconnection among patterns is what confers value to the language; however, this incompleteness may limit its immediate generalisability across all Consumer IoT contexts.

8.4.3. Limitations of the Validation Study

Limited Training Duration

The validation study provided only 10–20 minutes of introduction on the patterns followed by a 20-minute design session. While it is encouraging the fact that the participants were able to create feasible and user-centred UX flows within this short timeframe, it may also have limited the depth and sophistication of the design outcomes. Extended exposure periods might reveal additional benefits or challenges in the application of the pattern language.

Sample Size and Diversity

Due to time and resource constraints, the validation study was conducted with a relatively small sample of designers, which, although sufficient for a proof-of-concept perspective, may limit the immediate generalisability of the findings. A larger, more diverse sample would be advisable to confirm the broader applicability of the pattern language.

Limited Scope of Scenarios

The design tasks used in the validation study were limited to specific Consumer IoT scenarios. While these were chosen to represent common challenges, they may not fully cover the diverse situations in which the pattern language might be applied.

Future research should explore a wider range of scenarios, including those that involve complex, multi-user environments.

Remote Workshop Constraints

The design sessions (as well as the IoT expert assessment and some of the pattern mining sessions), were conducted remotely, which may have impacted on the nuances of communication and the capture of non-verbal cues. This limitation could have affected the quality of feedback captured and, consequently, the outcomes of the sessions.

8.4.4. Limitations of the Consumer IoT Expert Assessment

Reliance on a Single Expert

The technical feasibility of the proposed UX flows was assessed by a single Consumer IoT expert. While his insights were invaluable and provided critical independent validation of the design solutions, relying on one expert introduces a potential bias and limits the breadth of technical perspectives. Future studies would benefit from incorporating evaluations from multiple experts to ensure a more comprehensive and balanced assessment.

Potential for Subjective Interpretation

As with any expert review, the IoT expert's assessment is subject to individual interpretation and personal experience. Although his evaluation largely corroborated the feasibility of the UX flows, variations in expertise and perspective among different experts could potentially achieve different assessments of the same design artefacts.

Limited Scope of Feedback

The expert assessment, while thorough, primarily focused on technical feasibility and naturally covered aspects such as initial setup and multi-user considerations. There is scope for future evaluations to examine additional dimensions – such as long-term maintainability, scalability, and integration challenges – that are critical for real-world implementations in complex Consumer IoT ecosystems.

8.4.5. Recommendations for Future Research

To address these limitations and build upon the findings of this research, I propose the following directions for future inquiry:

Incorporating AI in the methodology

Recent advances of Generative AI present opportunities that were not available at

the time of this research study began. For instance, the 2024 study '**Flowy: Supporting UX Design Decisions Through AI-Driven Pattern Annotation in Multi-Screen User Flows**' (Lu et al., 2024) introduces an application that aids designers by identifying underlying patterns used in user flows they present, in addition to providing further examples of flows using the identified patterns. Similar approaches could be leveraged accelerate pattern identification by focusing on key unique characteristics of the type of patterns relevant for this research project, such as the integration of machine learning and the context of Consumer IoT.

Expanding and Refining the Pattern Language

Additional pattern seeds could be documented and integrated into the language, particularly by leveraging the Personal Spark pattern as an ongoing pattern-mining mechanism. Collaborative approaches that involve other stakeholders such as ML engineers may help mitigate subjectivity and enhance the comprehensiveness of the pattern language.

Conducting Longitudinal and Larger-Scale Validation Studies

Subsequent validation studies could engage a larger, more diverse cohort of designers and extend over longer periods. This would enable a deeper assessment of how designers' proficiency with the pattern language evolves and how the language performs in varied, real-world contexts.

Exploring Multi-User and Complex Scenarios

Future studies could explicitly investigate the application of the pattern language in multi-user settings and more complex Consumer IoT scenarios, where conflicting user needs are prevalent. Further enhancing the language by increasing the focus on these complexities will help to broaden its applicability.

Incorporating Multiple Expert Assessments

To reduce the risk of bias and provide a more balanced technical evaluation, future research should incorporate assessments from multiple Consumer IoT experts. This would provide more nuanced understanding of the practical feasibility and potential identification of a wider range of challenges associated with the proposed designs.

8.4.6. Final Reflections on Limitations

While this research presents a coherent body of evidence in relation to the guiding propositions and makes meaningful theoretical and practical contributions, that evidence is limited by the exploratory design of the project. The limitations discussed in the preceding sections indicate clear directions for refinement and extension. Addressing them in future work would strengthen the evidential basis of

the pattern language and clarify its applicability across a wider range of Consumer IoT contexts.

8.5. Final Concluding Remarks

In conclusion, this research project has addressed the two research questions through which the guiding propositions were explored. Taken together, the pattern mining study, the formal documentation process, and the proof-of-concept validation indicate that an ML-enhanced UX pattern language for Consumer IoT can be articulated and meaningfully applied by designers without deep machine learning expertise. These conclusions should be understood within the qualitative and exploratory scope of the project; nevertheless, they provide a strong foundation for further refinement, validation, and extension. The first research question (RQ1) asked:

Can a pattern language that encapsulates the use of machine learning to support user interactions with Consumer IoT be identified?

Yes. The findings presented in this thesis indicate that the answer is affirmative in principle, within the scope of this inquiry.

I demonstrated that by utilising an innovative pattern mining technique aligned with the classic Alexandrian Pattern Theory as a framework, it was possible to identify and formally document a novel, ML-enhanced UX pattern language applied to the Consumer IoT context. This language includes key UX patterns: **Personal Spark**, **Preference Learner**, **Geographic Trigger**, **Smart Shortcuts**, **Stop Learning Preferences** and the proto-pattern **Intelligent Access Sharing**, which collectively address **Guiding Proposition 1 (GP1)**:

A pattern language that encapsulates machine-learning-driven interaction guidance for Consumer IoT can be identified, thereby helping designers conceptualise adaptive and user-centred features.

I have shown that these patterns are not isolated entities but are interconnected, and that the unique versatility of the **Personal Spark** pattern enables this pattern language to be a dynamic and evolving resource that could indefinitely align to the complexity of modern Consumer IoT UX design challenges.

The second research question (RQ2) inquired:

If identified, can such a pattern language support the effective design of user interaction concepts for Consumer IoT scenarios by designers with no deep machine learning expertise?

Yes. Within the scope of this proof-of-concept study, the findings suggest that the pattern language can support this form of design practice.

The validation study provides encouraging evidence. The five designers produced UX flows that were generally assessed positively against the rubric, and the independent expert judged those concepts to be technically feasible in principle. These findings suggest that the pattern language may help designers with limited ML expertise to develop feasible Consumer IoT interaction concepts. However, this conclusion should be understood within the limits of a small, exploratory study and therefore as a basis for further validation.

This interpretation remains consistent with **Guiding Proposition 2 (GP2)**:

When employed by designers with no deep expertise in machine learning, this pattern language may support the creation of feasible Consumer IoT interaction concepts.

There is one final aspect I want to address, however. One may have noticed that unlike all the patterns formally documented through this research project, I did not provide a name for this novel pattern language. That is intentional. My final (symbolic) contribution to the field is to encourage the community of practice to decide and converge towards the most suitable name. Names can introduce subjective interpretations of relevance that may not always be shared across an entire community. For example, the name I gave to the first pattern I authored, **Personal Spark**, was directly linked to one aspect of Alexanders' work – the **quality without a name**⁴³ – from which, I derived relevance for a contemporary application: namely, how the outcome of combining good UX practices is experienced by end users – 'it just feels right'. This perspective may not be shared amongst the UX community or even the Pattern Authoring community. Recognising that subjectivity, the subsequent patterns I authored followed a more 'relevant' and generalisable naming convention, reached during the shepherding⁴⁴ process.

It has been a privilege to lay the foundations for what has the potential of becoming an ever-growing resource to the UX design community, and the ultimate proof of its utility will be when this pattern language begins to be adopted and referred to by

⁴³ Alexander's Quality Without a Name refers to an indescribable quality in design that permeates a space or object with a sense of natural coherence and vitality – one that he describes as much more than a simple sum of its parts and as a self-sustaining 'fire'. In his views, that almost abstract attribute makes a place feel intuitively 'right' and resonates with human needs, even though it cannot be measured or named.

⁴⁴ The special pattern authoring review process where the reviewer is an experienced pattern author <https://www.europlop.net/submission/>

whatever name it is given. This name will encapsulate the very essence of what makes patterns practical tools – **their organic emergence from a community of practice**.

Ultimately, I am confident that my research has successfully addressed the research questions and guiding propositions. The findings demonstrate that integrating machine learning with user-centred design practices can lead to the development of adaptive, feasible, and innovative interaction concepts for Consumer IoT. This work not only advances theoretical understanding but also provides practical tools for industry practitioners, therefore bridging the gap between academic research and real-world application. I believe that the framework developed herein will serve as a strong foundation for future innovation and exploration in this evolving field.

9. References

- 15288-2023 - ISO/IEC/IEEE International Standard - Systems and software engineering--System life cycle processes, 2023.
- Akado, Y., Kogure, S., Sasabe, A., Hong, J.-H., Saruwatari, K., Iba, T., 2015. Five patterns for designing pattern mining workshops, in: Proceedings of the 20th European Conference on Pattern Languages of Programs. Presented at the EuroPLOP 2015: 20th European Conference on Pattern Languages of Programs, ACM, Kaufbeuren Germany, pp. 1–13.
<https://doi.org/10.1145/2855321.2855331>
- Albeshar, A.S., 2023. Reviewing the Usability of Web Authentication Procedures: Comparing the Current Procedures of 20 Websites. *Sustainability* 15, 11043.
<https://doi.org/10.3390/su151411043>
- Alexander, C., 1979. *The Timeless Way of Building*. Oxford University Press.
- Alexander, C., Alexander, P. in the D. of A.C., Silverstein, M., Angel, S., Ishikawa, S., Abrams, D., 1975. *The Oregon Experiment*. Oxford University Press.
- Alexander, C., Ishikawa, S., Silverstein, M., 1977. *A pattern language: towns, buildings, construction*. Oxford University Press, New York.

- Amershi, S., Cakmak, M., Knox, W.B., Kulesza, T., 2014. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 105–120. <https://doi.org/10.1609/aimag.v35i4.2513>
- Amershi, S., Inkpen, K., Teevan, J., Kikin-Gil, R., Horvitz, E., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., 2019. Guidelines for Human-AI Interaction, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*. Presented at the the 2019 CHI Conference, ACM Press, Glasgow, Scotland Uk, pp. 1–13. <https://doi.org/10.1145/3290605.3300233>
- Arifin, W.N., Yaacob, N.M., 2025. Sample size determination for machine learning in medical research. <https://doi.org/10.48550/arXiv.2503.05809>
- Barbosa-Hughes, R., 2020a. Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT, in: *Proceedings of the European Conference on Pattern Languages of Programs 2020, EuroPLOP '20*. Association for Computing Machinery, New York, NY, USA, pp. 1–9. <https://doi.org/10.1145/3424771.3424777>
- Barbosa-Hughes, R., 2020b. Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT 9.
- Barbosa-Hughes, R., 2019. A Pattern Approach for Identification of Opportunities for Personalisation and Automation of User Interactions for the IoT, in: *Proceedings of the 24th European Conference on Pattern Languages of Programs, EuroPLOP '19*. ACM, New York, NY, USA, p. 8:1-8:9. <https://doi.org/10.1145/3361149.3361157>
- Borchers, J., 2001. *A Pattern Approach to Interaction Design*. John Wiley & Sons, Inc., USA.
- Borchers, J.O., 2000a. A Pattern Approach to Interaction Design, in: *Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques, DIS '00*. ACM, New York, NY, USA, pp. 369–378. <https://doi.org/10.1145/347642.347795>
- Borchers, J.O., 2000b. Interaction Design Patterns: Twelve Theses, in: *In Pattern Languages for Interaction Design: Building Momentum, CHI 2000, The Hague, the Netherlands, Apri*. pp. 2–3.
- Bryant, A., Charmaz, K., 2010. *The SAGE Handbook of Grounded Theory: Paperback Edition*. SAGE.
- Buschmann, F., Meunier, R., Rohnert, H., Sommerlad, P., Stal, M., 1996. *Pattern-Oriented Software Architecture - Volume 1: A System of Patterns*. Wiley Publishing.
- Caivano, D., Fogli, D., Lanzilotti, R., Piccinno, A., Cassano, F., 2018. Supporting end users to control their smart home: design implications from a literature review and an empirical investigation. *Journal of Systems and Software* 144, 295–313. <https://doi.org/10.1016/j.jss.2018.06.035>

- Card, S., Moran, T., Newell, A., 1980. The keystroke-level model for user performance time with interactive systems. CACM. <https://doi.org/10.1145/358886.358895>
- Cavoukian, A., 2010. Privacy by design: the definitive workshop. A foreword by Ann Cavoukian, Ph.D. IDIS 3, 247–251. <https://doi.org/10.1007/s12394-010-0062-y>
- Chollet, F., 2017. Deep Learning with Python. Manning Publications Company.
- Chuenjitwongsa, S., 2017. How to Conduct a Delphi Study. Educational Research.
- Chung, E.S., Hong, J.I., Lin, J., Prabaker, M.K., Landay, J.A., Liu, A.L., 2010. Development and Evaluation of Emerging Design Patterns for Ubiquitous Computing 10. <https://doi.org/https://doi.org/10.1145/1013115.1013148>
- Corbin, J., Strauss, A., 2014. Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory. SAGE Publications.
- Cover, T., Hart, P., 1967. Nearest neighbor pattern classification. IEEE Transactions on Information Theory 13, 21–27. <https://doi.org/10.1109/TIT.1967.1053964>
- Dankers, F.J.W.M., Traverso, A., Wee, L., van Kuijk, S.M.J., 2019. Prediction Modeling Methodology, in: Kubben, P., Dumontier, M., Dekker, A. (Eds.), Fundamentals of Clinical Data Science. Springer International Publishing, Cham, pp. 101–120. https://doi.org/10.1007/978-3-319-99713-1_8
- Dawes, M.J., Ostwald, M.J., 2017. Christopher Alexander's A Pattern Language: analysing, mapping and classifying the critical response. City, Territory and Architecture 4, 17. <https://doi.org/10.1186/s40410-017-0073-1>
- Dearden, A., Finlay, J., 2006. Pattern languages in HCI: A critical review. Human-computer interaction 21, 49–102.
- Deka, B., Huang, Z., Franzen, C., Hibschan, J., Afergan, D., Li, Y., Nichols, J., Kumar, R., 2017. Rico: A Mobile App Dataset for Building Data-Driven Design Applications. ACM Press, pp. 845–854. <https://doi.org/10.1145/3126594.3126651>
- Denzin, N.K., 2017. The Research Act: A Theoretical Introduction to Sociological Methods. Routledge, New York. <https://doi.org/10.4324/9781315134543>
- Deschamps-Sonsino, A., 2018. Smarter Homes: How Technology Will Change Your Home Life. Apress.
- Dibitonto, M., Tazzi, F., Leszczynska, K., Medaglia, C.M., 2018. The IoT Design Deck: A Tool for the Co-design of Connected Products, in: Ahram, T., Falcão, C. (Eds.), Advances in Usability and User Experience, Advances in Intelligent Systems and Computing. Springer International Publishing, Cham, pp. 217–227. https://doi.org/10.1007/978-3-319-60492-3_21
- Dix, A., 2017. Human-computer interaction, foundations and new paradigms. Journal of Visual Languages & Computing 42, 122–134. <https://doi.org/10.1016/j.jvlc.2016.04.001>

- Dix, A., Finlay, J., Abowd, G.D., Beale, R., 2003. Human-Computer Interaction, 3 edition. ed. Pearson, Harlow, England ; New York.
- Doosti, B., Dong, T., Deka, B., Nichols, J., 2018. A Computational Method for Evaluating UI Patterns. arXiv:1807.04191 [cs].
- Dove, G., Halskov, K., Forlizzi, J., Zimmerman, J., 2017. UX Design Innovation: Challenges for Working with Machine Learning As a Design Material, in: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17. ACM, New York, NY, USA, pp. 278–288.
<https://doi.org/10.1145/3025453.3025739>
- Duyne, D.K.V., Landay, J.A., Hong, J.I., 2007. The Design of Sites: Patterns for Creating Winning Web Sites. Prentice Hall Professional.
- Erickson, T., 2000. Lingua Francas for design: sacred places and pattern languages, in: Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques, DIS '00. Association for Computing Machinery, New York, NY, USA, pp. 357–368.
<https://doi.org/10.1145/347642.347794>
- Fogarty, J., Tan, D., Kapoor, A., Winder, S., 2008. CueFlik: interactive concept learning in image search, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08. Association for Computing Machinery, New York, NY, USA, pp. 29–38.
<https://doi.org/10.1145/1357054.1357061>
- Fogli, D., Lanzilotti, R., Piccinno, A., 2016. End-User Development Tools for the Smart Home: A Systematic Literature Review, in: Streitz, N., Markopoulos, P. (Eds.), Distributed, Ambient and Pervasive Interactions, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 69–79.
https://doi.org/10.1007/978-3-319-39862-4_7
- Gamma, E., Helm, R., Johnson, R., Vlissides, J., 1995. Design Patterns: Elements of Reusable Object-Oriented Software. Pearson Education.
- Gartner, 2015. Gartner's 2015 Hype Cycle for Emerging Technologies Identifies the Computing Innovations That Organizations Should Monitor [WWW Document]. URL <https://www.gartner.com/en/newsroom/press-releases/2015-08-18-gartners-2015-hype-cycle-for-emerging-technologies-identifies-the-computing-innovations-that-organizations-should-monitor> (accessed 1.6.25).
- Google, 2014. This is material design- Google Developers Blog [WWW Document]. URL <https://developers.googleblog.com/en/this-is-material-design/> (accessed 2.16.25).
- Gross, M.D., 1998. Smart House And Home Automation Technologies. Encyclopedia of Housing, design machine group, University of Washington 6.
- Guo, B., Zhang, D., Wang, Z., Yu, Z., Zhou, X., 2013. Opportunistic IoT: Exploring the harmonious interaction between human and the internet of things. Journal of

- Network and Computer Applications 36, 1531–1539.
<https://doi.org/10.1016/j.jnca.2012.12.028>
- Habib, H., Cranor, L.F., 2022. Evaluating the Usability of Privacy Choice Mechanisms. Presented at the Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022), pp. 273–289.
- Harley, A., 2018. UX Expert Reviews [WWW Document]. Nielsen Norman Group. URL <https://www.nngroup.com/articles/ux-expert-reviews/> (accessed 4.27.25).
- Harrison, N., 2003. Advanced Pattern Writing., in: EuroPloP. pp. 809–828.
- Harrison, N.B., 1999. The language of shepherding. Pattern languages of program design 5, 507–530.
- Hassenzahl, M., 2004. The Interplay of Beauty, Goodness, and Usability in Interactive Products. Human–Computer Interaction 19, 319–349.
https://doi.org/10.1207/s15327051hci1904_2
- Hasson, F., Keeney, S., McKenna, H., 2000. Research guidelines for the Delphi survey technique. J Adv Nurs 32, 1008–1015.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition. Springer Science & Business Media.
- Hastie, T., Tibshirani, R., Friedman, J.H., 2001. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media.
- Hentrich, C., Zdun, U., Hlupic, V., Dotsika, F., 2015. An approach for pattern mining through grounded theory techniques and its applications to process-driven SOA patterns, in: Proceedings of the 18th European Conference on Pattern Languages of Program, EuroPloP '13. Association for Computing Machinery, New York, NY, USA, pp. 1–16. <https://doi.org/10.1145/2739011.2739020>
- Hussain, A., Mkpojiogu2, E., Husin, M., 2021. Dimensioning UX Models for Design and Evaluation. Turkish Journal of Computer and Mathematics Education (TURCOMAT) 12, 1878–1883.
- Iba, T., 2014. A Journey on the Way to Pattern Writing: Designing the Pattern Writing Sheet. Presented at the 21st Conference on Pattern Languages of Programs (PloP2014)., p. 17.
- Iba, T., Yoder, J., 2014. Mining Interview Patterns: Patterns for Effectively Obtaining Seeds of Patterns, in: 10th Latin American Conference on Pattern Languages of Programs (SugarloafPloP2014). Presented at the SugarloafPloP, p. 13.
- Iqbal, S., Pison-Young, L., 2009. The Delphi method [WWW Document]. BPS. URL <https://www.bps.org.uk/psychologist/delphi-method> (accessed 5.26.25).
- James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. An Introduction to Statistical Learning: with Applications in R. Springer Science & Business Media.
- Johnson, R.B., Onwuegbuzie, A.J., 2004. Mixed Methods Research: A Research Paradigm Whose Time Has Come. Educational Researcher 33, 14–26.
<https://doi.org/10.3102/0013189X033007014>

- Jordan, M.I., Mitchell, T.M., 2015. Machine learning: Trends, perspectives, and prospects. *Science* 349, 255–260. <https://doi.org/10.1126/science.aaa8415>
- Kirwan, B., Ainsworth, L.K., 1992. *A Guide To Task Analysis: The Task Analysis Working Group*. CRC Press.
- Kolko, J., 2010. *Thoughts on Interaction Design*. Morgan Kaufmann.
- Kuckartz, U., Rädiker, S., 2019. *Analyzing Qualitative Data with MAXQDA: Text, Audio, and Video*. Springer International Publishing, Cham.
<https://doi.org/10.1007/978-3-030-15671-8>
- Kuniavsky, M., 2019. *The UX of Predictive AI in the IoT (Rosenfeld To Be Designed)*.
- Kuniavsky, M., 2010. *Smart Things: Ubiquitous Computing User Experience Design*. Elsevier.
- Lee, I., Lee, K., 2015. The Internet of Things (IoT): Applications, investments, and challenges for enterprises. *Business Horizons* 58, 431–440.
<https://doi.org/10.1016/j.bushor.2015.03.008>
- Leitner, H., 2015. *Pattern Theory: Introduction and Perspectives on the Tracks of Christopher Alexander: Volume 1, 1 edition*. ed. CreateSpace Independent Publishing Platform, Graz.
- Lew, G., Schumacher Jr, R.M., 2020. *AI and UX: Why Artificial Intelligence Needs User Experience*. Apress.
- Li, Z., Yang, W., Yuan, J., Wu, J., Chen, C., Ming, Y., Yang, F., Zhang, H., Liu, S., 2024. A Scalable Matrix Visualization for Understanding Tree Ensemble Classifiers.
<https://doi.org/10.48550/arXiv.2409.03164>
- Liao, Q.V., Vaughan, J.W., 2024. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *Harvard Data Science Review*.
<https://doi.org/10.1162/99608f92.8036d03b>
- Lu, Y., Tong, Z., Zhao, Q., Oh, Y., Wang, B., Li, T.J.-J., 2024. Flowy: Supporting UX Design Decisions Through AI-Driven Pattern Annotation in Multi-Screen User Flows. <https://doi.org/10.48550/arXiv.2406.16177>
- Malizia, A., Paternò, F., 2023. Why Is the Current XAI Not Meeting the Expectations? *Commun. ACM* 66, 20–23. <https://doi.org/10.1145/3588313>
- Mays, N., Pope, C., 2000. Assessing quality in qualitative research. *BMJ* 320, 50–52.
<https://doi.org/10.1136/bmj.320.7226.50>
- Mesaros, G., Doble, J., 1997. A pattern language for pattern writing, in: *Proceedings of International Conference on Pattern Languages of Program Design (1997)*. p. 164.
- Meyer, E., Wachter-Boettcher, S., 2016. *Design for Real Life. A Book Apart*, New York.
- Mohanraj, I., Ashokumar, K., Naren, J., 2016. Field Monitoring and Automation Using IOT in Agriculture Domain. *Procedia Computer Science, Proceedings of the 6th International Conference on Advances in Computing and Communications* 93, 931–939. <https://doi.org/10.1016/j.procs.2016.07.275>

- Myers, B.A., 1998. A brief history of human-computer interaction technology. *interactions* 5, 44–54. <https://doi.org/10.1145/274430.274436>
- Nielsen, J., 2000. Why You Only Need to Test with 5 Users [WWW Document]. Nielsen Norman Group. URL <https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/> (accessed 1.23.25).
- Nielsen, J., 1993. *Usability Engineering*, 1 edition. ed. Morgan Kaufmann, Amsterdam.
- Nielsen, J., Mack, R.L. (Eds.), 1994. *Usability Inspection Methods*. John Wiley & Sons, New York.
- Nielsen, J., Molich, R., 1990. Heuristic evaluation of user interfaces, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Empowering People - CHI '90*. Presented at the the SIGCHI conference, ACM Press, Seattle, Washington, United States, pp. 249–256. <https://doi.org/10.1145/97243.97281>
- Norman, D., 2013. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.
- Norman, D.A., 1988. *The psychology of everyday things, The psychology of everyday things*. Basic Books, New York, NY, US.
- Osgood, C.E., Suci, G.J., Tannenbaum, P.H., 1957. *The Measurement of Meaning*. University of Illinois Press.
- Perera, C., Zaslavsky, A., Christen, P., Georgakopoulos, D., 2014. Context Aware Computing for The Internet of Things: A Survey. *IEEE Commun. Surv. Tutorials* 16, 414–454. <https://doi.org/10.1109/SURV.2013.042313.00197>
- Poslad, S., 2011. *Ubiquitous Computing: Smart Devices, Environments and Interactions*. John Wiley & Sons.
- Prechelt, L., Unger, B., Tichy, W.F., Brossler, P., Votta, L.G., 2001. A controlled experiment in maintenance: comparing design patterns to simpler solutions. *IEEE Transactions on Software Engineering* 27, 1134–1144. <https://doi.org/10.1109/32.988711>
- PricewaterhouseCoopers, 2017. *Smart home, seamless life: Unlocking a culture of convenience*.
- Reenskaug, T.M.H., 1979. *The original MVC reports*.
- Resnick, M., Albert, W., 2014. The Impact of Advertising Location and User Task on the Emergence of Banner Ad Blindness: An Eye-Tracking Study. *International Journal of Human-Computer Interaction* 30, 206–219. <https://doi.org/10.1080/10447318.2013.847762>
- Riehle, D., Züllighoven, H., 1995. A pattern language for tool construction and integration based on the tools and materials metaphor, in: *Pattern Languages of Program Design*. ACM Press/Addison-Wesley Publishing Co., USA, pp. 9–42.

- Robson, C., 2011. Real World Research, 3rd edition. ed. John Wiley & Sons, Chichester.
- Rosson, M.B., Carroll, J.M., 2002. Usability Engineering: Scenario-Based Development of Human-Computer Interaction. Elsevier Science.
- Rowland, C., Goodman, E., Charlier, M., Light, A., Lui, A., 2015. Designing Connected Products: UX for the Consumer Internet of Things. O'Reilly Media, Inc.
- Sasabe, A., Kaneko, T., Takahashi, K., Iba, T., 2016. Pattern mining patterns: a search for the seeds of patterns, in: Proceedings of the 23rd Conference on Pattern Languages of Programs, PLoP '16. The Hillside Group, USA, pp. 1–16.
- Schümmer, T., Matschke, C., 2013. Proto-patterns for growing a community of practice, in: Proceedings of the 18th European Conference on Pattern Languages of Program, EuroPLOP '13. Association for Computing Machinery, New York, NY, USA, pp. 1–8. <https://doi.org/10.1145/2739011.2739033>
- Sharp, H., Preece, J., Rogers, Y., 2019. Interaction Design: Beyond Human-Computer Interaction. John Wiley & Sons.
- Shneiderman, B., 2022. Human-Centered AI. Oxford University Press, Oxford, New York.
- Shneiderman, B., Plaisant, C., 1987. Designing the user interface: strategies for effective human-computer interaction. IGBIO Newsl.
- Smith, R.L., 1988. Smart House: The Coming Revolution in Housing. GP Pub.
- Strauss, A., Corbin, J.M., 1998. Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory. SAGE Publications.
- Swan, M., 2012. Sensor Mania! The Internet of Things, Wearable Computing, Objective Metrics, and the Quantified Self 2.0. Journal of Sensor and Actuator Networks 1, 217–253. <https://doi.org/10.3390/jsan1030217>
- Taylor, M.J., Reilly, D., Wren, C., 2018. Internet of things support for marketing activities. Journal of Strategic Marketing.
- Tidwell, J., 2005. Designing Interfaces: Patterns for Effective Interaction Design. O'Reilly Media, Inc.
- Tidwell, J., 1998. Interaction design patterns. PLoP'98 Conference on Pattern Language of Programming, Illinois.
- Turchi, T., Malizia, A., Paternò, F., Borsci, S., Chamberlain, A., 2024. Adaptive XAI: Towards Intelligent Interfaces for Tailored AI Explanations, in: Companion Proceedings of the 29th International Conference on Intelligent User Interfaces, IUI '24 Companion. Association for Computing Machinery, New York, NY, USA, pp. 119–121. <https://doi.org/10.1145/3640544.3645253>
- van Duyne, D.K., Landay, J.A., Hong, J.I., 2003. The Design of Sites: Patterns, Principles, and Processes for Crafting a Customer-centered Web Experience. Addison-Wesley Professional.
- van Lent, M., Laird, J., 1999. Developing an Artificial Intelligence Engine. p. 12.

- van Welie, M., van der Veer, G.C., Eliëns, A., 2001. Patterns as Tools for User Interface Design, in: Vanderdonckt, J., Farenc, C. (Eds.), Tools for Working with Guidelines. Springer London, London, pp. 313–324.
https://doi.org/10.1007/978-1-4471-0279-3_30
- Vega-Barbas, M., Pau, I., Augusto, J.C., Seoane, F., 2018. Interaction Patterns for Smart Spaces: A Confident Interaction Design Solution for Pervasive Sensitive IoT Services. IEEE Access 6, 1126–1136.
<https://doi.org/10.1109/ACCESS.2017.2777999>
- Weber, T., Zoitl, A., Hußmann, H., 2019. Usability of Development Tools: A CASE-Study, in: 2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C). Presented at the 2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C), pp. 228–235.
<https://doi.org/10.1109/MODELS-C.2019.00037>
- Welie, M., Trætteberg, H., 2000. Interaction Patterns in User Interfaces. Proceedings of Seventh Pattern Languages of Programs Conference (PLoP'00).
- Wellhausen, T., Fießer, A., 2012. How to write a pattern?: A rough guide for first-time pattern authors, in: Proceedings of the 16th European Conference on Pattern Languages of Programs. ACM, p. 5.
- Wigdor, D., Wixon, D., 2011. Brave NUI World: Designing Natural User Interfaces for Touch and Gesture, 1st ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Yang, Q., 2018. Machine Learning as a UX Design Material: How Can We Imagine Beyond Automation, Recommenders, and Reminders?, in: 2018 AAAI Spring Symposium Series. Presented at the 2018 AAAI Spring Symposium Series.
- Yang, Q., 2017. The Role of Design in Creating Machine-Learning-Enhanced User Experience, in: 2017 AAAI Spring Symposium Series. Presented at the 2017 AAAI Spring Symposium Series.
- Yin, R.K., 2018. Case study research and applications: design and methods, Sixth edition. ed. SAGE, Los Angeles London New Delhi Singapore Washington DC Melbourne.
- YouGov, 2018. The dawn of the connected home. YouGov.
- Zhang, C., Budgen, D., 2012. What Do We Know about the Effectiveness of Software Design Patterns? IEEE Transactions on Software Engineering 38, 1213–1231.
<https://doi.org/10.1109/TSE.2011.79>
- Zhang, X., Yao, L., Zhang, S., Kanhere, S.S., Sheng, Q.Z., Liu, Y., 2018. Internet of Things Meets Brain-Computer Interface: A Unified Deep Learning Framework for Enabling Human-Thing Cognitive Interactivity. arXiv:1805.00789 [cs].

Appendix A

The Role of the Consumer IoT Experience in Extending the Notion of Home

Rubem Barbosa-Hughes

University of Hertfordshire, School of Creative Arts

Hatfield, Hertfordshire, United Kingdom

r.barbosa-hughes@herts.ac.uk

Abstract

The notion of smart homes is often associated with intelligent spaces that can improve people's lives. These improvements may materialise in a wide range of ways, from basic applications such as automation of mundane tasks and energy saving to more critical ones such as health monitoring and assistance to vulnerable adults.

With more people welcoming the proposed benefits of the smart home concept, naturally the adoption of smart devices is set to increase exponentially. However, in order to support the interaction between people and these smart objects, user experience (UX) designers are faced with the task of designing digital interactions with physical devices, many of which are common everyday objects that people are accustomed to interacting with.

This paper reviews the early concepts of smart homes and explores how the interdisciplinary field of UX design is dealing with the challenge of conceiving effective, efficient and pleasing interaction solutions for the modern idea of smart homes. Firstly, we explore approaches using human-to-human interactions as a benchmark for human-to-digital interface interactions along with other approaches exploring opportunities for using AI and Machine Learning. Secondly, we explore how the resulting designs may impact the notion of "being at home in the city in the twenty-first century".

Introduction

The Internet of Things (IoT) is the name given to the network of devices or 'Things' featuring sensors and/or actuators, software and other technological mechanisms that allow them to connect with other Things and systems as well as to the internet.

IoT is the central idea behind Industry 4.0 which boosts efficiency in areas such as agriculture, logistics, manufacturing, healthcare and more.

As an example, IoT solutions have been assisting farmers to monitor plant growth and calculate water needs and crop profit (Mohanraj et al., 2016), tasks which would have been significantly more challenging without this technology.

The Consumer IoT focuses on the context of products targeted at the consumer market. Often the mainstream media refer to Consumer IoT devices as 'Smart Home devices' or 'Connected devices'. We use the term 'smart device' and 'connected device' interchangeably throughout this paper.

Unsurprisingly, due to the popularity of its parent field, the Industry 4.0, the consumer IoT emerged with high expectations.

In the next chapters we will discuss the idiosyncratic challenges of designing experiences for the Consumer IoT, as well as how the focus on human interaction is critical not only to incentivise its adoption, but also in impacting the notion of being at home.

Initially, it was all about automation

The idea of automating mundane tasks in people's homes is nothing new. In fact, the creation of the first home automation device is credited to Jim Sutherland who, in 1966, created the ECHO IV (an acronym for Electronic Computing Home Operator). The device, which comprised four large cabinets⁴⁵ (6' x 2' x 6'), controlled temperature and appliances, and could save shopping lists, recipes, and other family memos.

The term 'Smart house' was trademarked on the 21st October 1984 by the National Association of Home Builders (NAHB) in the USA. The NAHB developed the Smart House project in partnership with some industry representatives. Among other main selling points, it was meant to be significantly safer than 'non smart' counterparts. The electricity supply was only live on power outlets which had an appliance plugged in, and only while it was turned on. The Smart house also featured sensors to detect smoke, gas and water leaks, and any other scenarios deemed abnormal. The term 'Smart Home' was then adopted by those companies involved in the competing projects (Deschamps-Sonsino, 2018; Gross, 1998).

The common motivator behind these early concepts of Smart Homes was automation, freeing users from the burden of mundane tasks. However, similarly to the Smart House project in the 1980's, the modern notion of Smart Homes still

⁴⁵ The ECHO IV Home Computer: 50 Years Later, 2016. CHM. URL <https://computerhistory.org/blog/the-echo-iv-home-computer-50-years-later/>

faces the challenge that truly 'smart' homes have to be built from the ground up or else entail substantial conversions. Moreover, the infrastructure that supports it all would need to be reviewed if we were to extend this notion to the scale of modern cities. This means all appliances, water and heating systems would be in some way connected to each other and to the internet and thereby integrated with intelligent cloud systems to ensure the appropriate levels of personalisation and resource optimisation.

Introducing personalisation is potentially as important as addressing the technological and logistical challenges. Implementing automation with a 'one size fits all' approach would be ineffective even in neighbourhoods with a homogenous demographic profile. If we take electricity usage as an example, a general trend can certainly be identified when analysing data from entire countries or cities, however, when we fine tune to the level of the individual households, significant variations in distinct routines and preferences would show up in the same geographical area. This likely suggests that data from industrial level, large scale applications has limited relevance to inform the design of consumer level implementations.

The challenge of automation for Smart Homes is no longer solely about enabling devices to perform tasks autonomously; rather it is about them performing these tasks autonomously when relevant to the user. As a trivial example, infrared motion sensors can easily be used to notify the user of any movement of their front door. To the user, however, knowing that 'there was movement' in the front door is an incomplete solution. A camera could be useful, but the cognitive task of examining the image to conclude whether there was a 'friend or foe' at the door still lies with the user.

This need for a more intelligent solution could be driving growth in the adoption of single smart devices. Instead of the more comprehensive Smart Home concept that has been envisaged since the 1960's, an aspect that is actually impacting the notion of home in the 21st century is the steady rate of introduction of these single Smart Things into peoples' houses, one by one, as well as the shift in behaviour that their adoption brings.

The smarter home

As Mike Kuniavsky stated, 'connecting stuff to the internet is easy... and pointless' (Kuniavsky, 2019). The implication is that remote operation is not such a novel interaction and does not essentially produce value. Smart speakers are the most popular device owned in the UK. Overall, nearly a quarter of the British population (23%) own a smart home device and 11% of these devices are smart speakers (YouGov, 2018). Interestingly, one of the main reasons for the popularity of Smart

Speakers is how well they work with other devices, indicating that the appetite for a more complete ecosystem still exists. This is still a far cry from the ubiquitous computing exhibited in certain notions of Smart Spaces though, where users constantly interact with multiple intelligent services (Vega-Barbas et al., 2018).

Controlling a compatible device using a smart speaker as an interface, on the other hand, has become a trivial task in many UK homes. However, the usability concerns around this newer interaction scenario spans from the ability of the Natural Language Processing model (to deal with varied utterances) through to how well the interface handles issues with latency in the communication. In addition, the experience is no longer limited to the house as a location, but can extend to anywhere users take a mobile phone with them.

In such a context, the technical challenges merge seamlessly with the user interaction challenges. For those households adopting smart devices, having a good experience at home is now becoming increasingly dependent on the previous work of user experience (UX) designers, as much as the work of electricians or plumbers.

Designing smarter home interactions

When electricians and plumbers produce poor designs for basic home facilities such as placing light switches or water taps in awkward places, the usability of the home is compromised. When the user interaction with smart home devices is not carefully considered, the usability of the home suffers equally. The severity of the impact depends on how much the user relies on those devices, and more critically, when the interaction fails.

In recent years, a growing body of work has been dedicated to addressing the challenges associated with designing good experiences with the Consumer IoT (Alshammari et al., 2019; Barbosa-Hughes, 2019; Dibitonto et al., 2018a; Kuniavsky, 2010; Rowland et al., 2015). One of the approaches explored in previous works was to compare human-human interaction against human-computer interaction to accomplish a given task. The method is supported by the 'What Would a Human Do?' (WWHD) thought experiment and allows designers to identify suitable elements of human-human interactions which could be translated into digital interactions by using technologies such as machine learning and location services.

This is how it works: if we consider a smart lock as an example, a typical use case would be to notify the user when a suspicious attempt to unlock the door is happening or has happened. Using the WWHD approach to reframe this interaction problem, the designer could sketch, wireframe or use any other way to document the initial idea for the user interaction. Then the designer needs to ask the help of a

colleague (or another person who is familiar with the problem the design concept is trying to solve) to perform a simple thought experiment, i.e.

‘to imagine that, instead of the (mobile/voice/VR/AR) user interface as a point of interaction, there is a human being performing the role of the interface.’

In this example, a mobile app that accompanies the device could allow the user to set up notifications if the lock is tampered with in some way, but these notifications should not be sent if it is an authorised person. The WWHD approach would typically lead to comments such as ‘I would call the (imaginary) person by their name, then ask them to please let me know if they can see any stranger trying to unlock the door’. This simple analogy of an interaction with another person may suggest that one of the most critical steps to solve the interaction problem is in reality to produce a mechanism to identify who is and who isn’t a stranger.

This approach is suitable for early, conceptual stages of the Design Thinking process, when solutions are being explored. It helps to address the right user interaction problem from a human-centric perspective, rather than being technology led.

Other approaches have been explored to address other stages of Consumer IoT UX design, such as the use of patterns and Card Decks. Patterns in Human-Computer interaction are formal ways of describing a proven solution to a problem, in this case an interaction problem. Patterns addressing user interactions with entire Smart Spaces (Vega-Barbas et al., 2018) or with individual smart devices (Barbosa-Hughes, 2020b) can accelerate the design process by using tried and tested solutions. On the other hand, the IoT Design Deck is defined as ‘a method for the co-design of connected products’ (Dibitonto et al., 2018). It was developed to assist with three stages of the IoT design challenge: Design, Facilitation and Prototype. It anticipates that designing for IoT should be a multi-disciplinary effort of co-creation.

These and many more novel methods are behind the design of the experiences people have in their homes today. While some succeed, many fail in delivering an efficient, effective and pleasant user experience.

In late 2013 the UK Government commissioned a report on the usability of connected heating systems. The result was not very positive: all six of the then major connected heating devices failed in delivering a good user experience. A few years later, in 2018, a smart oven owner shared a photo of his oven on social media

displaying a 'Terms of Service' screen he had to accept before using the appliance⁴⁶ - clearly not a good example of user-centred design.

The Consumer IoT brands are slowly learning that investing in the UX of its devices is a critical step if they want to participate in the steady growth in consumer adoption.

The role of smartphones in extending the notion of home
With smartphones frequently used as one of the primary interfaces to control connected devices, users can interact and intervene with their household routine even when they are not physically at home.

In the past, smartphones have received a negative reputation for enabling work to invade our private lives. However, in a redeeming role, smartphones can now enable people to enjoy a small portion of the emotional experience of being at home whenever they access their smart cat flap log or by simply overriding and adjusting temperature controls if the kids are coming home earlier, for example.

Through remote operation, users of smart devices can continue to exert their influence over both positive and negative factors at home, even while away from it. This novel notion of home experience could be compared in some ways to communicating via two different dialects of a language. People may broadly understand and communicate basic ideas, but more complex, refined interactions are more challenging to achieve. Nonetheless, good examples of connected device interactions through smartphones allow for that level of dialect-like interaction, not quite as refined as the physical interactions, but still valid and meaningful.

Smartphones may even be considered a catalyst for another evolution of the IoT: The IoTAR or Internet of Things Augmented Reality, which can further blur the lines between the physical and the digital, potentially adding another dimension to the notion of being at home in the twenty-first century city.

REFERENCES

- Alshammari, A., Alhadeaf, H., Alotaibi, N., Alrasheedi, S., Prince, M., 2019. The Usability of HCI in Smart Home 06, 7.
- Barbosa-Hughes, R., 2020. Interaction Patterns using Machine Learning and Location Services in User Interfaces for the Consumer IoT 9.

⁴⁶ <https://twitter.com/dotMorten/status/1018243235364274176/photo/1>

- Barbosa-Hughes, R., 2019. A Pattern Approach for Identification of Opportunities for Personalisation and Automation of User Interactions for the IoT, in: Proceedings of the 24th European Conference on Pattern Languages of Programs, EuroPlop '19. ACM, New York, NY, USA, p. 8:1–8:9. <https://doi.org/10.1145/3361149.3361157>
- Deschamps-Sonsino, A., 2018. Smarter Homes: How Technology Will Change Your Home Life. Apress.
- Dibitonto, M., Tazzi, F., Leszczynska, K., Medaglia, C.M., 2018a. The IoT Design Deck: A Tool for the Co-design of Connected Products, in: Ahram, T., Falcão, C. (Eds.), Advances in Usability and User Experience, Advances in Intelligent Systems and Computing. Springer International Publishing, Cham, pp. 217–227. https://doi.org/10.1007/978-3-319-60492-3_21
- Gross, M.D., 1998. Smart House And Home Automation Technologies. Encyclopedia of Housing, design machine group, University of Washington 6.
- Kuniavsky, M., 2019. The UX of Predictive AI in the IoT (Rosenfeld To Be Designed).
- Kuniavsky, M., 2010. Smart Things: Ubiquitous Computing User Experience Design. Elsevier.
- Mohanraj, I., Ashokumar, K., Naren, J., 2016. Field Monitoring and Automation Using IOT in Agriculture Domain. Procedia Computer Science, Proceedings of the 6th International Conference on Advances in Computing and Communications 93, 931–939. <https://doi.org/10.1016/j.procs.2016.07.275>
- Rowland, C., Goodman, E., Charlier, M., Light, A., Lui, A., 2015. Designing Connected Products: UX for the Consumer Internet of Things. O'Reilly Media, Inc.
- Vega-Barbas, M., Pau, I., Augusto, J.C., Seoane, F., 2018. Interaction Patterns for Smart Spaces: A Confident Interaction Design Solution for Pervasive Sensitive IoT Services. IEEE Access 6, 1126–1136. <https://doi.org/10.1109/ACCESS.2017.2777999>
- YouGov, 2018. The dawn of the connected home. YouGov.

Appendix B

(Conference paper excerpt reproduced for reference)

Professional Doctorates in Design and the use of Computational Tools

Silvio Carta, Barbara Brownie, Ian W. Owen, Rubem Barbosa-Hughes, Sahar Khajeh
University of Hertfordshire

This contribution focuses on the relevance of Professional Doctorates in Design within the larger context of the Professional Doctorates. This work analyses the importance of digital culture and computational tools as a way to create new knowledge and allow doctoral students to make original contributions in their research. We claim that doctoral students need to innovate in the ways in which they use digital tools as a part of their professional practice in their research. This may include creating new tools, modifying existing ones or even applying tools developed in a specific field into other areas. This idea is firstly compared with recent shifts in large architectural practices where specialised teams are created to develop new computational tools. Secondly, we compare it with other fields of design. Finally, we present a case study of one of our doctoral students who has developed new tools for his research using data analytics and machine learning techniques as a way to produce new knowledge in his design work.

This study concludes with a number of pointers for those involved in professional doctorates in general, and in art and design in particular, where digital components are becoming increasingly relevant and, inevitably common.

The digital culture and the importance of digital tools

Digital technologies and, more specifically, their application to computational design are increasingly permeating into many aspects of practice and research in design. The more new techniques, devices and computational capacity are available, the more automated, optimised and exploratory design tools become. Designers are increasingly able to perform their task quickly and more accurately, resulting in new and extraordinary projects. Researchers have new tools at their disposal to generate, analyse and manage data and information. There is a risk, however, that using established tools may inevitably lead to the repetition of processes that have been carried out before, and the reproduction of familiar aesthetics, with the result that it is difficult to make an original contribution to design practice.

Shifts in design practices in the pursuit of new knowledge

We looked at an increasing number of design practices that heavily rely on algorithmic design and computational techniques as a normal part of their design process. This trend has been extensively observed (see for example Hopfe et al. 2006, Besserud and Hussey 2011, Heumann 2014). The use of computer-based techniques in practice can be divided into two main categories. The first is the use of computers to simplify laborious tasks usually performed by designers to acquire a higher level of precision and control over the process and to save time and resources. The use of computer-aided design (CAD) developed throughout the

1990s and 2000s epitomises this approach, wherein designers use software to replicate hand-drafting more efficiently and accurately. The second category includes the use of algorithmic and computational logic in the design.

Examples throughout the history of design suggest that in order to make an original contribution to the field, one must often first develop an original technology, or an original method of using established technology. Such is the dependence on computational tools that innovation in design is increasingly tied into innovation in technology, with the adoption or adaption of computational tools and processes becoming part of the process of innovation. Commentary in the field of graphic design has long identified the problem of “rampant” technological determinism (Heller and Vienne, 2012, p. 212). Where innovation has occurred in design, it has frequently involved seeking liberation from the constraints imposed by a technology, or exploring the possibilities presented by a new technology that has not previously been employed in design. The rejection of established technologies has been key to innovations in design, and has characterised contributions to the field. In the field of graphic design, for example April Greiman’s adoption of early Apple Mackintosh computers transformed graphic design from a two-dimensional to three-dimensional practice (Meggs, 2011, p. 468). In the field of architecture, not only are design practices using a greater variety of computational tools to evaluate, simulate, test and find variations of their design using experimental pieces of software (such as the numerous plug-ins and libraries for the Grasshopper and Dynamo programs), but further to this they are increasingly developing their own tools in-house. Examples of this are the group CODE within the Zaha Hadid Architects, Kohn Pedersen Fox Associates (KPF) Urban Interface (UI) and the Applied Research and Development group at Foster and Partners.

Where innovation in design arises from innovation in use of computational tools, the process of conducting research in design relies upon a greater understanding of how computers work, and the use of computational thinking as a fundamental part of the design process. This approach includes an awareness of algorithmic logics, datasets and statistic models. An example of this is the use of machine learning (ML) techniques or agent-based modelling (ABM) as a part of the design practice. The technique of form-finding (Piker 2013, Senatore & Piker 2015), for instance, is a clear use of such an approach, whereby designers develop a process (intended as a series of tasks that a computer will perform), setting certain conditions to be met, and use computers to run a series of simulations/tests that will eventually return an optimised shape.

Approaches appropriate for doctoral-level study

The transition to doctoral-level study necessitates critical reappraisal of the processes and techniques that may have been previously employed at Masters level or in professional practice. While Masters level education goes some way to developing a critical approach to tools and methods, at this level we tend to observe rejection of established tools (such as a return to the handmade) rather than innovation in their use, a tendency that must be overcome in order to achieve the original contribution to knowledge that characterizes doctoral output. Steven Heller (2012, p. 212) describes how students are taught to identify these limitations as a core part of Higher Education in graphic design, particularly at postgraduate level, and encouraged to critically reflect upon the influence of digital technologies on design practice. Many of these students therefore embark on doctoral study with a critical approach to the technology that industry takes for granted in design practice, prepared to challenge the use of these tools. Critical re-evaluation of the

available digital tools becomes an integral part of their doctoral study, and can be viewed as a crucial step towards ensuring their original contribution to the field.

Case study: Pattern Writing Sheet

We present here the case of one DDes (Professional Doctorate in Design) candidate whose work pertains primarily to the field of design, but methods from social sciences and computer science have been employed in order to explore the research topic originally. Rubem Barbosa-Hughes's work hinges on human-centred design as an *"approach to interactive systems development that aims to make systems usable and useful by focusing on the users, their needs and requirements, and by applying human factors/ergonomics, and usability knowledge and techniques"* (ISO 9241-210 2010). The Human-centred design approach aims to produce better user experience (UX) in recognition that, rather than designing artefacts, designers are crafting the experience which users will have with those design products. In addition to the qualitative methods used in HCI, Barbosa-Hughes decided to employ the "Pattern Writing Sheet" (Iba 2014). This method is used to guide the initial structure and aid the pattern writing process, during which, the writer is encouraged to 'think more deeply about the actual problem that the pattern solves and what consequences arise by applying the solution' (Wellhausen and Fießer, 2012). Once documented, the aim is to submit the patterns to a rigorous process of review conducted by experts in the relevant field called 'Shepherding' (Wellhausen and Fießer, 2012), in the Patterns Language of Program (PLoP) conference (<https://europlop.net/>). This case illustrates how the use of methods originated and widely used in other disciplines, can bring new insights on design problems.

Pattern Name (7) Give a Name to this pattern (6-2)

(Ideas of Pattern Name)

A good name expresses the essence and is memorable. Usually very ordinary language with two nouns or nouns and adjective. Utter the name in order to check whether it is easy to say as common language.

Image (6-1) Try to sketch the pattern showing the conflicting forces and the solution that resolves the conflict. Refer to the image when thinking of a Name. Think of new words to express this pattern. Imagine the essence of this pattern. Is there an important tip or technique in the theme or domain?

Subject (1) What kind of theme or domain do you want to write a pattern about?

Context (4-1) When or where does the problem occur? Specify the context. In this context

Forces (4-2) Forces in a pattern are laws or tendencies that we cannot change and which make the problem difficult because they can be incompatible. Because of these Forces Why does the problem occur? What kind of forces are at work?

Problem (3) What kind of situation is it when the problem occurs? Therefore What will happen if you don't implement the Solution?

Solution (2-1) Identify one important thing you really want to share with colleagues and newcomers. First think a lot, then choose just one.

Actions (2-2) Be concrete, for example Be abstract. It can be said as For example What is the Consequence of the Solution? + : Positive Consequence, generating living quality - : Negative Consequence, side effects

Consequence (5) As a Result

Pattern Writing Sheet with Instructions Ver 0.01

This work "Pattern Writing Sheet" by Takashi Iba is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. See the details about the license at the site <http://creativecommons.org/licenses/by-nc-sa/4.0/>. Contact us by E-mail in advance, if you want to use this sheet for commercial or business. E-Mail: contact[at]unishiro.jp

Pattern Writing Sheet, Iba (2014: pp 11)

Impact on Practice



- A sharper focus on what tools are relevant for our Prof Docs students
- A bigger awareness of ways in which work can be prepared, shared and delivered
- A clearer alignment between design tools used in practice and within the Professional Doctorate in Design
- Better understanding in design practice of how to conduct a Professional Doctorate in Design

References:

- Besserud, Keith, and Thomas Hussey. "Urban design, urban simulation, and the need for computational tools." *IBM Journal of Research and Development* 55.1.2 (2011): 2-1.
- Heller, S. & Vienne, V. 2012. *The Education of a Graphic Designer*, Allworth Press: New York.
- Heumann, Andrew. "COMPUTATIONAL APPROACHES TO PROCESS, PRACTICE, AND PRODUCT." *Acadia 2014*. Design Agency. Acadia
- Hopfe, C. J., et al. "Computational tools for building services design-professional's practice and wishes." *Proceedings of the 17th Int. Air-conditioning and Ventilation Conference*. STP—Society of Environmental Engineering, 2006.
- Iba, T., 2014. *Journey on the Way to Pattern Writing, Designing the Pattern Writing Sheet*. 21st Conference on Pattern Languages of Programs, 2014.
- Meggs, P.B., Purvis, A.W., 2011. *Meggs' History of Graphic Design*. 5th edn, John Wiley & Sons: Hoboken.
- Piker, Daniel. "Kangaroo: form finding with computational physics." *Architectural Design* 83.2 (2013): 136-137.
- Senatore, Gennaro, and Daniel Piker. "Interactive real-time physics: an intuitive approach to form-finding and structural analysis for design and education." *Computer-Aided Design* 61 (2015): 32-41.
- Wellhausen, T., Fießer, A., 2012. How to write a pattern?: A rough guide for firsttime pattern authors, in: Proceedings of the 16th European Conference on Pattern Languages of Programs. ACM, p. 5.

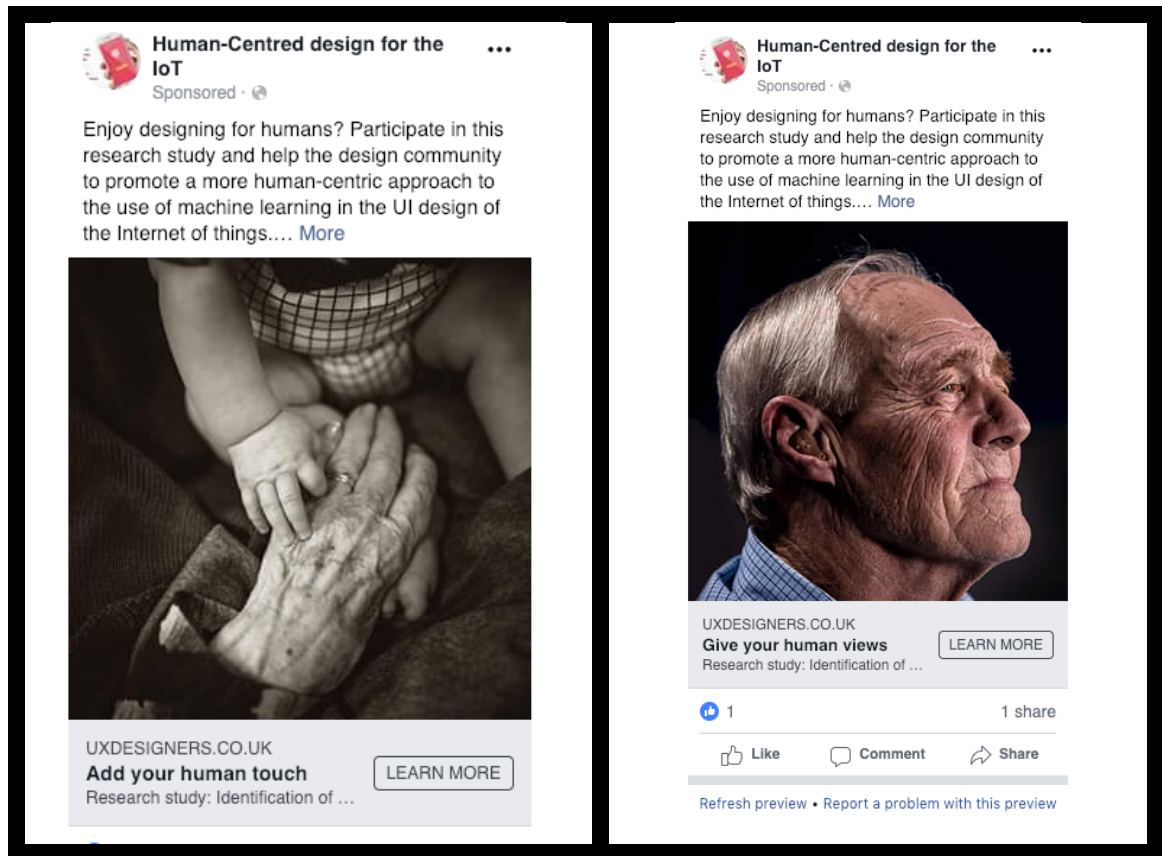
Appendix C

Variations of Ad texts tested showing focus on IoT, Machine Learning and both IoT and Machine Learning.

UX Designers: Help to promote human-centricity to the the UI design of the Internet of Things. Visit uxdesigners.co.uk to find out how to get involved.		Research project: University of Hertfordshire UH
UX Designers: Help to promote human-centricity to the use of machine learning in the UI design of the Internet of Things. Visit uxdesigners.co.uk to find out how to get involved.		Research project: University of Hertfordshire UH
UX Designers: Help to promote human-centricity to the use of machine learning. Visit uxdesigners.co.uk to find out how to get involved.		Research project: University of Hertfordshire UH

Appendix D

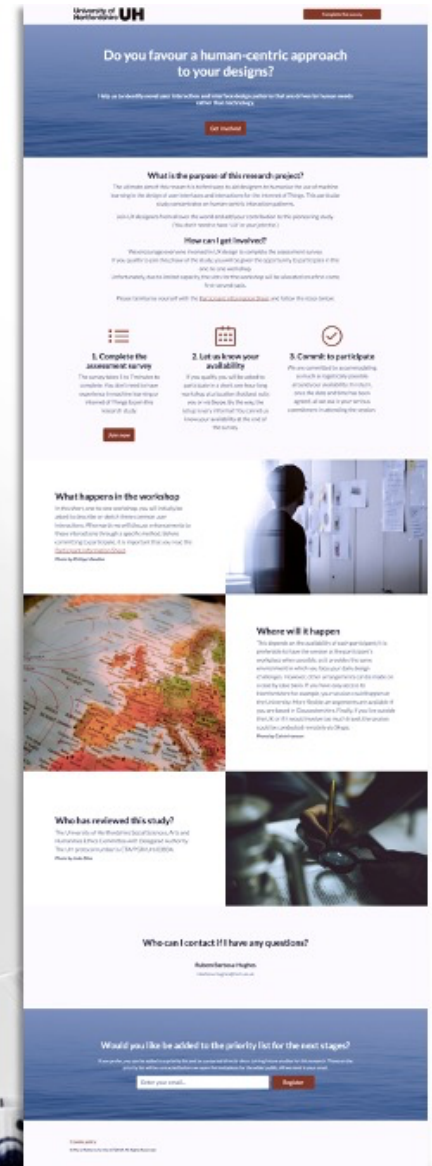

The two creative approaches used in the A/B test to determine the most effective ad at attracting UX professionals to participate in the pattern mining study.



Appendix E


The research study's information and recruitment website.


The study page: uxdesigners.co.uk





Appendix F

The UI pattern sheet shared with participants of the pattern mining workshops featuring Material Design components.



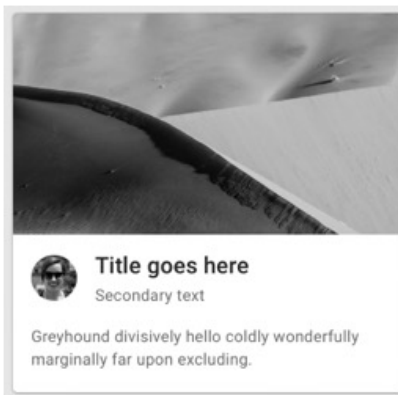




Title Label

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad


Action Button



Title goes here

Secondary text

Greyhound divisively hello coldly wonderfully marginally far upon excluding.



Elizabeth Park


Work contact



(650) 555-1234
Mobile



heyfromelizabeth@gmail.com
Work







Title goes here




Secondary text

Greyhound divisively hello coldly wonderfully marginally far upon excluding.

ACTION 1 ACTION 2



Progress indicator



Appendix G

Validation Study information page also hosted at uxdesigners.co.uk

University of
Hertfordshire

UH

Join the study

Machine Learning-Driven UX for IoT: Designers, share your views!

At this PhD research delves deep into Machine Learning-driven UX patterns, your design expertise becomes crucial.

We've identified ML-enhanced UX patterns for Consumer IoT experiences, but they need your professional insights and invaluable feedback. Challenge our findings, share your experiences, and help us shape the intelligent design of tomorrow.

This research is about the design practice, and not technology - so you are not required to have experience with Machine Learning or Internet of Things.

[Get involved](#)

What is the purpose of this research project?

The ultimate aim of this research is to find ways to aid designers to humanise the use of machine learning in the design of user experiences for the Consumer Internet of Things. This particular study concentrates on gathering feedback on formally documented patterns.

Join UK designers from all over the world and add your contribution to this pioneering study.

How can I get involved?

We encourage everyone involved in UX design for more than 5 years to join this study. Unfortunately, due to limited capacity the slots for the workshop will be allocated on a first come, first served basis.

Please follow the steps below:



1. Complete the assessment survey

The survey takes less than 5 minutes to complete. You don't need to have experience in machine learning or Internet of Things to join this research study.

[Join now](#)



2. Let us know your availability

If you qualify, you will be asked to participate in a short online workshop. By the way, the setup is very informal. We will get in touch to agree the most suitable time for your session.



3. Commit to participate

We are committed to accommodating as much as logistically possible around your availability. In return, once the date and time has been agreed, all we ask is your serious commitment in attending the session.

What happens in the workshop

In this short, one-to-one workshop, you will initially be asked to use one or more UX patterns. Afterwards we will discuss different aspects around the use of these patterns.

Photo by Philipp Mueller





Where will it happen

Your session will be conducted fully remotely, with initial contact to organise it made through whatever contact means you share.

Photo by Calvin Hansen

Who has reviewed this study?

The University of Hertfordshire Social Sciences, Arts and Humanities Ethics Committee with Delegated Authority. The UH protocol number is CTARGA/UN058/71.

Photo by Julie Siles



Who can I contact if I have any questions?

Rubem Barbosa-Hughes
r.barbosa-hughes@hertsm.ac.uk



[Cookie policy](#)
© Machine Learning-Driven UX for IoT 2021. All Rights Reserved.

Appendix H

IoT Expert Questionnaire

Personal Information

1. Current Job Title: (Optional for confidentiality if preferred. This information will potentially be placed in the final PhD Dissertation)
2. Organisation: (Optional for confidentiality if preferred. This information will potentially be placed in the final PhD Dissertation)
3. Years of Experience in UX Design/Research or related fields:
4. What is your highest level of academic qualification?
5. What field is your degree in?
6. Areas of Specialisation within UX Design: (e.g., Interaction Design, User Research, Information Architecture, etc.)
7. Have you worked on UX projects specifically related to IoT? If yes, could you briefly describe one or two notable projects?
8. Have you authored any publications or presented at conferences related to UX or IoT? If yes, please provide titles or links where applicable.
9. Have you previously participated as an expert in academic research studies? If yes, please briefly describe the nature of your involvement.

10. Have you received any awards or recognitions for your work in UX design?
11. Are you a member of any professional organisations or associations related to UX design or technology?
12. What do you believe are the most pressing challenges currently facing UX design in IoT?
13. What trends or innovations in UX design do you find most exciting or promising?

Appendix I

Pattern Mining Recruitment & Screening Survey Structure

The complete list of survey questions is as follows:

1. If you are interested in taking part, please read the statements below and then select all statements below to record your consent to participate. [checkbox]
2. First name [single line free text]
3. Surname [single line free text]
4. Contact details (please give sufficient information to enable the investigator to get in touch with you, such as email address) [single line free text]
5. What is your age group [multiple choice (single answer)]
6. How long have you been working in UX or Interaction design? [multiple choice (single answer)]
7. Have you ever been involved in the UX design of an Internet of Things (IoT) product or app? [multiple choice (single answer)]
8. Have you ever been involved in the UX design of an app which used machine learning to help with any aspect of the user experience (e.g., facial recognition, personalised recommendations based on previous interactions, search based on an image, voice assistants' integration, etc)? [multiple choice (single answer)]

[Logic flow: If "Yes" the participant was presented with Q. 9, if "No" presented with Q. 10]


9. Were you involved in any stage of planning or designing the role of machine learning in that project? [multiple choice (single answer)]
10. Do you currently use any Design System? (i.e., a library of reusable components and patterns that follow a particular style guide) [multiple choice (single answer)]

[Logic flow: If "Yes" the participant was presented with Q. 11, if "No" the survey was ended.]

11. Is it an open source (e.g., Material Design) or custom made for a company or brand you work for? [multiple choice (single answer)]

Logic Flow of the Survey

The logic flow of the survey was designed to tailor questions based on participants' responses, ensuring relevance and efficiency. The diagram below exported from JISC surveys illustrates the logic flow.



Logic flow of the screening survey for the validation study