

Classifying Types of Gesture and Inferring Intent

Christopher L. Nehaniv

Adaptive Systems Research Group
School of Computer Science
Faculty of Engineering and Information Sciences
University of Hertfordshire
College Lane, Hatfield Herts AL10 9AB
United Kingdom
C.L.Nehaniv@herts.ac.uk

Abstract

In order to infer intent from gesture, a rudimentary classification of types of gestures into a main class is introduced. The classification is intended as a basis for incorporating the understanding of gesture into human-robot interaction (HRI). Some requirements for the operational classification of gesture by a robot interacting with humans are also suggested.

1 Introduction: The Need for Classifying Gesture

The word *gesture* is used for many different phenomena involving human movement, especially of the hands and arms. Only some of these are interactive or communicative. The pragmatics of gesture and meaningful interaction are quite complex (cf. Kendon (1970); Mey (2001); Millikan (2004)), and an international journal *Gesture* now exists entirely devoted to the study of gesture. Applications of service or 'companion' robots that interact with humans, including naive ones, will increasingly require human-robot interaction (HRI) in which the robot can recognize *what* humans are doing and to a limited extent *why* they are doing it, so that the robot may act appropriately, e.g. either by assisting, or staying out of the way. Due to the situated embodied nature of such interactions and the non-human nature of robots, it is not possible to directly carry over methods from human-computer interaction (HCI) or rely entirely on insights from the psychology of human-human interaction. Insights from proxemics and kinesics, which study spatial and temporal aspects of human-human interaction (Hall, 1983; Condon and Ogston, 1967; Kendon, 1970) and some insights of HCI, e.g. recognizing the diversity of users and providing feedback acknowledgment with suitable response timing (e.g. (Shneiderman, 1998)), may also prove to be extremely valuable to HRI. Notwithstanding, the nascent field of HRI must develop its own methods particular to the challenges of embodied interaction between humans and robots. New design, validation, evaluation methods and principles particular to HRI must be developed to meet challenges such as *legibility*, making the robot's actions and behaviour understandable and predictable to a human, and *robotiquette*, respecting human activities and situations (e.g. not interrupting a conversation between humans

or disturbing a human who is concentrating or working intensely (without sufficient cause), as well respecting social spaces, and maintaining appropriate proximity and levels of attention in interaction.

Part of meeting these challenges necessarily involves some understanding of human activity at an appropriate level. This requires the capabilities of recognizing human gesture and movement, and inferring intent.

The term *intent* is used in this paper in a *limited* way that refers to *particular motivation(s) of a human being that result in a gestural motion as relevant for human-robot interaction*.

In inferring the intent from a human's gesture it is helpful to have a classification of which type of gesture is being observed. Without a sufficiently broad classification, understanding of gesture will be too narrow to characterize what is happening and appropriate responses will not be possible in many cases.

While this paper does not attempt a comprehensive survey of the role and recognition of gesture in human-robot interaction, it does suggest inherent limitations of approaches working with a too narrow notion of gesture, excluding entire classes of human gesture that should eventually be accessible to interactive robots able to function well in a human social

74

environment.

The questions of how gestures are acquired and come to be recognized as meaningful by particular individuals in the course of their development (ontogeny of gesture and its recognition), and conventionalized, elaborated, or lost within particular cultures

(evolution of gesture) are large and deep issues, but will not be addressed within the scope of this paper.

Knowing how to recognize and classify gesture may also serve to inform the design of robot behaviour, including gestures made by the robot to achieve legibility and convey aspects of the robot's state and plans to humans. This in turn will contribute to robot interaction with humans that is legible, natural, safe, and comfortable for the humans interacting with the robot.

2 Classification of Gestures

The following is a rough, tentative classification.

Gestures are classed into ve major types with some subtypes.

To begin to approach the complexity of gesture in the context of situated human-robot interaction, the rough classes of gesture described below are developed in order to provide a broad level of description and the rst steps toward a pragmatic, operational definition that could be used by an autonomous system such as a robot to help it (1) to infer the intent of human interaction partners, and, as an eventual goal, (2) to help the robot use gestures itself (if

possible) to increase the legibility of its behaviour.

Ambiguity of Gesture. It should be stressed that a single specific instance of a particular the kind of physical gestural motion could, depending on context and interaction history, reflect very different kinds of human intents. It will not always be possible to infer intent based solely on based the mechanical aspects of human movements (such as changes in joint angles) without taking context into account.

To approach this problem, a classification of gesture for inferring intent and assisting in the understanding of human activity should closely relate gesture with limited categories of intent in situated human activity. The classes of the tentative classification presented here thus correspond to and allow the (limited) attribution of intent on the part of humans. The classification is developed as an aid for helping robots to achieve limited recognition of situated human gestural motion so has to be able to respond appropriately if required, while these robots are working in an environment of ambient human activity (such as a home or office), in which, at times, the robots are also assisting or cooperating with the humans. *Applications of this classification will require the mapping of physical aspects of gestural motion in interactional contexts to the above gestural classes (and their subtypes) suggested here.*

2.1 Five Classes (with Subtypes)

1. **Irrelevant/Manipulative Gestures.** These include *irrelevant gestures, body / manipulator motion, side-effects of motor behaviour, and actions on objects*. Broadly characterized, manipulation by a human is here understood as *doing something to influence the non-animate environment or the human's relationship to it (such as position)*. Gestural motions in this class are manipulative actions (in this sense) and their side effects on body movement. These 'gestures' are neither communicative nor socially interactive, but instances and effects of human motion. They may be salient, but are not movements that are primarily employed to communicate or engage a partner in interaction. Cases include, e.g. motion of the arms and hands when walking; tapping of the fingers; playing with a paper clip; brushing hair away from one's face with one's hand; scratching; grasping a cup in order to drink its contents. (Note it may be very important to distinguish among the subtypes listed above for robot understanding of human behaviour.)

2. **Side Effect of Expressive Behaviour.** In communicating with others, motion of hands, arms and face (changes in their states) occur as part of the overall communicative behaviour, but without any specific interactive, communicative, symbolic, or referential roles (cf. classes 3-5)

Example: persons talk excitedly raising and moving their hands in correlation with changes in voice prosody, rhythm, or emphasis of speech.

3. **Symbolic Gestures.** Gestural motion in symbol gesture is a *conventionalized signal in a communicative interaction*. It is generally a member of a limited, circumscribed set of gestural motions that have specific, prescribed interpretations. A symbolic gesture is used to trigger certain actions by a targeted perceiver, or to refer to something or substitute as for another signal according to a code or convention. Single symbolic gestures are analogous to discrete actions on an interface, such as clicking a button.

75

Examples: waving down a taxi for it to stop; use of a conventional hand signals (a command to halt indicated open palm; a military salute); nodding 'yes'; waving a greeting 'hello' or 'goodbye'.

Note that the *degree of arbitrariness* in such gestures may vary: The form of the gesture may be an arbitrary conventional sign (such as holding up two fingers to mean 'peace', or the use of semaphores for alphabetic letters). On the other hand, a symbolic gesture may resemble to a lesser or greater extent iconically or, in ritualized form, a referent or activity.

Further examples: holding up two fingers to indicate 'two'; opening both (empty) hands by turning palms down to indicate a lack of something.

Nearly all symbolic gestures are used to convey *content* in communicative interactions.

4. **Interactional Gestures.** These are gestures used to *regulate interaction with a partner*, i.e. used to initiate, maintain, invite, synchronize, organize or terminate a particular interactive, cooperative behaviour: raising an empty hand toward the partner to invite the partner to give an object; raising the hand containing an object toward the partner inviting them to take it; nodding the head indicating that one is listening. The emphasis of this category is neither reference nor communication but on gestures as mediators for cooperative action.¹ Interactional gestures thus concern regulating the form of interactions, including the possible regulation of communicative interactions but do not generally convey any of the content in communication. Interactional gestures are similar to class 1 manipulative gestures in the sense that they influence the environment, but in contrast to class 1, they influence the animated environment. *doing something to influence human agents (or other agents) in the environment*, but not by conveying symbolic or referential content.²

¹Note that we are using the word *cooperative* in a sense that treats regulating communication or interaction as an instance of cooperation.

²Some more subtle examples include putting one's hand on another person's arm to comfort them. Such actions, and others involving physical contact, may be quite complex to interpret as understanding them may require understanding and modeling the intent of one person to influence that state of mind of another. At this point, we class simply them with interactional gestures recognizing that future analysis may reveal deep issues of human-human interaction and levels of complexity beyond the rudimentary types of human intent considered here. A special case worthy of note is human contact with the robot, unless this is directly a manipulation of the robot's state via an interface - e.g. via button presses

5. Referential/Pointing Gestures. These are *used to refer to or to indicate objects (or loci) of interest*

. either physically present objects, persons, directions or locations the environment .

by pointing (*deixis* . showing), or indication of locations in space being used as proxies to represent absent referents in discourse.

Table 1 summarizes the *ve* classes.

2.2 Target and Recipient of a Gesture

If a gesture is used interactively or communicatively (classes 2-5), it is important to recognize whether the gesture is directed toward the current interaction partner (if any) . which may be the robot, another person

(or animal) present in the context, or possibly neither (*target*). If pointing, what is the person pointing to?

Who is the pointing designed to be seen by? (*recipient*)

If speaking, to whom is the person speaking? If the gesture is targetted at or involves a contact with an object, this suggests it may belong to class 1 (or possibly 5, even without contact). A gesture of bringing an object conspicuously and not overly quickly toward an interaction partner is manipulative (in the sense explained in the discussion of class 1, since an object is being manipulated), but it may well at the same time also be a solicitation for the partner to take the object (class 4). Similarly if the partner has an object, an open hand conspicuously directed toward the partner or object may be a solicitation for the partner to give the object (class 4).

2.3 Multipurpose Gestures

It is possible for a single instance of a particular gesture to have aspects of more than one class or to lie intermediate between classes. As mentioned above, handing over an object is both class 1 and 4. And, for example, holding up a yellow card in football has aspects of classes 1 and 3, object manipulation and .which would fall into class 3 (symbolic gesture), non-accidental human contact with the robot is likely to be indicative of an intent to initiate or regulate interaction with the robot (class 4). Physical contact between humans might also involve expression of affection (kissing), or aggression (slapping, hitting) . which generally indicate types human-human interaction it would be better for a robot to steer clear of!

³Deixis can involve a hand, *nger*, other directed motion, and/or eye gaze. Checking the eye gaze target of an interaction partner is commonly used to regulate reference and interaction; it

develops and supports joint attention already in preverbal infants. Language, including deictic vocabulary (e.g. demonstratives such as the words .these. and .that.), and other interactional skills, typically develop on this scaffolding (see Kita (2003)).

76

conventional symbolic signal. Many ritualized symbolic gestures (class 3) also can be used to initiate or regulate interaction (class 4), e.g. the 'come here' gesture: with palm away from the recipient, moving the fingers together part way toward the palm; waving forearm and open hand with palm facing recipient to get attention. More complex combinations are possible, e.g. a gesture of grasping designed by the human to be seen by a recipient interaction partner and directed toward a heavy or awkwardly-shaped target object as a solicitation of the partner to cooperatively carry the object with the gesturer (classes 1, 4, 5).

2.4 Ritualization: Movement into Classes 3 and 4

Gestures that originate in class 1 as manipulations of the non-animate environment and the person's relationship to it may become *ritualized* to invite interactions of certain types, e.g., cupping the hand next to the ear can indicate that person doing it cannot hear, so that the interaction partner should speak up. Originally cupping the hand near the ear served to improve a person's ability to hear sounds in the environment from a particular direction (class 1), but it may be intended to be seen by a conversational partner who then speaks up (class 4). The hand cupped at the ear can even be used as a conventionalized symbol meaning 'speak up' (class 3). Other examples of ritualization toward regulation of interaction and also symbolic gesture include mimicking with two hands the motions of writing on a pad as a signal to a waiter to ask for the bill; miming a zipping action across the mouth to indicate that someone should be 'shut up'; or placing a raised index finger over lips which have been pre-formed as if to pronounce /sh/.

2.5 Cultural and Individual Differences

Different cultures may differ in their use of the various types of gesture. Some symbolic gestures such as finger signs (e.g. the .OK. gesture with thumb and index finger forming a circle) can have radically different interpretations in other cultures, or no set interpretation depending on the culture of the recipient (e.g. crossing fingers as a sign of wishing for luck, or the Chinese finger signs for some numbers such as 6, 7, 8). Tilting the head back (Greece) or nodding the head (Bulgarian) are used symbolically for 'no', but would certainly not be interpreted that way in many other cultures. Cultures also differ in their types and scope of movement in (class 2) expressive gestures: Consider, for example, the differences of rhythm, prosody, hand motions, eye contact, and facial expressions accompanying speech between

British, Italian, Japanese, and French speakers. Within cultures, differences between different individuals' uses of gestures can be regional, restricted to particular social groups within the culture, and vary in particularities (such as speed, repertoire, intensity of movement, etc.) between individuals according to preference or ontology. Elderly and young may employ gestures in different ways.

3 Some Related Work on Recognizing Gesture and Intent

The important role of gesture for intent communication in human-robot interaction is increasingly being acknowledged, although some approaches still focus only on static hand poses rather than dynamic use of more general types of gesture in context. A survey of hand gesture understanding in robotics appears in Miners (2002).

Multimodal and voice analysis can also help to infer intent via prosodic patterns, even when ignoring the content of speech. Robotic recognition of a small number of distinct prosodic patterns used by adults that communicate praise, prohibition, attention, and comfort to preverbal infants has been employed as feedback to the robot's 'affective' state and behavioural expression, allowing for the emergence of interesting social interaction with humans (Breazeal and Aryananda, 2002). Hidden Markov Models (HMMs) have been used to classifying limited numbers of gestural patterns (such as letter shapes) and also to generate trajectories by a humanoid robot matching those demonstrated by a human (Billard et al., 2004). Multimodal speech and gesture recognition using HMMs has been implemented for giving commands via pointing, one-, and two-handed gestural commands together with voice for intention extraction into a structured symbolic data stream for use in controlling and programming a vacuuming cleaning robot (Iba et al., 2002). Many more examples in robotics exist.

Most approaches use very limited, constrained, and specific task-related gestural repertoires of primitives, and do not attempt to identify gestural classes. They have tended to focus on a fixed symbolic set of gestures (possibly an extensible one, in which new gestures can be learned), or focus on only a few representatives from one or two of the gestural classes identified here (e.g. symbolic and pointing gestures).

Knowledge of specific conventional codes and

77

signs can help the identification of particular signs within class 3, and also in determining that the gesture in fact belongs to class 3, i.e. is a symbolic communicative signal. Machine learning methods such as Hidden Markov Models may be used successfully to learn and classify gestures for a limited finite set

of *g* gestures (e.g. (Westeyn et al., 2003)). It seems likely that HMM methods would be most successful with class 3 (symbolic gestures), but how successful they would be at differentiating between classes or within other classes remains uninvestigated at present.

4 Inferring the Intent of Gesture

Being able to classify gesture into one of the above classes gives us only a starting point for inferring the intent of the person making the gesture due to frequent ambiguity. Resolving this points to the important roles of context and interactional history. Thus, it is necessary to develop operational methods for recognizing the class of gesture in a particular context.⁴

To this it should help when

- (a) the activity of the gesturer is known,
- (b) previous and current interaction patterns are remembered to predict the likely current and next behaviour of the particular person,
- (c) objects, humans and other animated agents in the environment are identified and tracked.
- (d) the scenario and situational context are known (e.g. knowing whether a gesture occurs at a tea party or during a card game).

Knowing the above could help the robot classify the gesture and infer the intent of the human. Information on the state of human (e.g. working, thirsty, talking, ...) often can limit the possibilities.

4.1 Recognizing Intent from Gesture Given Interactional Context

If the interactional context of recent activity in which a gesture occurs is known, this can suggest possibilities for which classes (and subtypes) of gesture might be involved. Even giving data on the interactional context, including data on context, culture, individual differences, models of human activity and ⁴Knowledge of the immediate context in some cases needs to be augmented by taking into account of the broader *temporal horizon* of interactional history (cf. Nehaniv et al. (2002)).

task aspects that relate to gesture, does not necessarily completely constrain the possible gesture nor its intent (if any). If the context suggests a particular identifying class (and subtype) for the gesture identified, this does not immediately lead to any certain knowledge of human intent behind it.

Data on the interaction history and context may help in determining the class of a gesture. If the class is known, then the set of possible gestures can remain large, or be narrowed significantly. Symbolic gestures (class 3) correspond to discrete symbols in a finite set, of which there may be only a small number according to context or size of the given repertoire of the given symbolic gestural code. Interactional gestures (class 4) are likely to comprise a small, constrained class. Class 1 gestures are either irrelevant, or to be understood by seeking the intent of the

associated motor action or object manipulation (e.g. grasping or throwing an object, arms moving as a side effect of walking). Class 5 (referential and pointing gestures) comprise a very limited class.

4.2 Typical Interactional Context of Gestures

A programme to apply the above classification can be developed as follows.

1. Identify the many, particular gestural motions that fit within each of the five classes. Some gestural motions will appear in more than one class. For example, the same mechanical motion of putting a hand and arm forward with the forearm horizontal and the hand open could indicate preparation to manipulate an object in front of the human (class 1), to show which object is being referred to (class 5), or to greet someone who is approaching, or to ask for an object to be handed over (both class 4).
2. Gestural motions identified as belonging to several classes need to be studied to determine in which contexts they occur: determining in which class(es) particular an instance of the gesture is being used may require consideration of objects and persons in the vicinity, the situational context, and the history of interaction.
3. Systematic characterizations of a physical gestural motion together with interactional contexts in which they occur could then be used to determine the likely class.

78

CLASSIFICATION OF GESTURAL CLASSES AND ASSOCIATED (LIMITED) CATEGORIES OF HUMAN INTENT

CLASS NAME DEFINING CHARACTERISTICS AND ASSOCIATED INTENT

1 'IRRELEVANT' OR MANIPULATIVE INFLUENCE ON NON-ANIMATE ENVIRONMENT GESTURES OR HUMAN'S RELATIONSHIP TO IT;

manipulation of objects, side effects of motor behavior, body motion

2 SIDE EFFECT OF EXPRESSIVE MARKING, BEHAVIOUR (NO SPECIFIC DIRECT INTERACTIVE, SYMBOLIC, REFERENTIAL ROLE)

associated to communication or affective states of human

3 SYMBOLIC GESTURES CONVENTIONALIZED SIGNAL IN COMMUNICATIVE INTERACTION; communicative of semantic content (language-like)

4 INTERACTIONAL GESTURES REGULATION OF INTERACTION WITH A PARTNER; INFLUENCE ON HUMAN (OR OTHER ANIMATED) AGENTS IN ENVIRONMENT BUT GENERALLY WITH LACK OF ANY SYMBOLIC/REFERENTIAL CONTENT

used to initiate, maintain, regulate, synchronize, organize or or terminate various types of interaction

5 REFERENTIAL/POINTING GESTURES DEIXIS; INDICATING OBJECTS, AGENTS OR (POSSIBLY PROXY) LOCI OF DISCOURSE TOPICS, TOPICS OF INTEREST; pointing of all kinds with all kinds of effectors (incl. eyes):

referential, topicalizing, attention-directing

Table 1: **Five Classes of Gesture.** See text for explanation, details and examples. Note that some occurrences of

the same physical gesture can be used different classes depending on context and interactional history; moreover, some gestures are used in a manner that in the same instance belongs to several classes (see text for examples).

79

4.3 Updating the Interaction History

Attribution of intent related to gesture can then feedback into understanding of the situational context, including motivational state of the human performing the gesture, and becomes part of the updated interaction history, which can then help in inferring intent from ensuing gestures and activity.

5 Conclusions

In order to infer the intent of a human interaction partner, it may be useful to employ a classification of gesture according to some major types. In the tentative classification proposed here, whose intent may be, in the above classes, absent / directed to objects or environment, incidentally expressive, symbolic, interactional, or deictic. A summary of the classes is given by Table 1.

In order to deploy the inference of intent on robots interacting with humans it will be necessary to operationalize the distinctions between these (sometimes overlapping) classes. This may require the use of knowledge of human activity, recognition of objects and persons in the environment, and previous interactions with particular humans, as well as knowledge of conventional human gestural referencing and expression, in addition to specialized signaling codes or symbolic systems.

The classification presented here suggests some requirements for the design and implementation of systems inferring intent from gesture based on this classification. These requirements might be realized in a variety of different ways using, e.g. continuous lowkey tracking or more detailed analysis, event-based and/or scenario-based recognition, and prediction of human activity based on models of human activity flows (with or without recognition of particular humans and their previous interactions), depending the particular needs of the given human-robot interaction design and the constraints and specificity of its intended operational context. Design of a robot restricted to helping always the same user in the kitchen environment would be quite different from one that should be a more general purpose servant or companion in a home environment containing several adults, children and pets, but the classification presented here is applicable in informing the design of gesture recognition for inferring intent in either type of system, and for designing other HRI systems.

Finally, effective human-robot interaction will require generation of gestures and feedback signals by the robot. The classification given here can suggest

categories of robotic gestures that could be implemented to improve the *legibility* to humans of the robot's behaviour, so that they will be better able to understand and predict the robot's activity when interacting with it.

Acknowledgments

The work described in this paper was conducted within the EU Integrated Project COGNIRON (.The Cognitive Robot Companion.) and was funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020.

The classification presented here is developed by the author in response to discussions in the COGNIRON project, especially with Rachid Alami, Kerstin Dautenhahn, Jens Kubacki, Martin Haegele, and Christopher Parlitz. Thanks to Kerstin Dautenhahn for useful comments on an earlier draft of this paper. This paper extends and supercedes University of Hertfordshire School of Computer Science Technical Report 419 (December 2004).

References

- A. Billard, Y. Epars, S. Calinon, G. Cheng, and S. Schaal. Discovering optimal imitation strategies. *Robotics & Autonomous Systems*, Special Issue: Robot Learning from Demonstration, 47(2-3): 69.77, 2004.
- Cynthia Breazeal and Lijin Aryananda. Recognition of affective communicative intent in robot-directed speech. *Autonomous Robots*, 12(1):83.104, 2002.
- W. S. Condon and W. D. Ogston. A segmentation of behavior. *Journal of Psychiatric Research*, 5:221. 235, 1967.
- Gesture*. ISSN: 1568-1475. John Benjamins Publishing Co., The Netherlands
http://www.benjamins.com/cgi-bin/t_seriesview.cgi?series=GEST.
- Edward T. Hall. *The Dance of Life: The Other Dimension of Time*. Anchor Books, 1983.
- Soshi Iba, Christian J. J. Paredis, and Pradeep K. Khosla. Interactive multi-modal robot programming. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation, Washington D.C., May 11-15, 2002*, 2002. 80
- Adam Kendon. Movement coordination in social interaction: Some examples described. *Acta Psychologica*, 32:100.125, 1970.
- Sotaro Kita, editor. *Pointing: Where Language, Culture and Cognition Meet*. Lawrence Erlbaum Associates, Inc, 2003.
- Jacob Mey. *Pragmatics: An Introduction*. Blackwell Publishers, 2001.
- Ruth Garrett Millikan. *The Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT Press/Bradford

Books, 2004.

William Ben Miners. *Hand Gesture for Interactive Service Robots*. MSc Thesis, The University of Guelph, Faculty of Graduate Studies, August 2002.

Chrystopher L. Nehaniv, Daniel Polani, Kerstin Dautenhahn, Ren'e te Boekhorst, and Lola Ca.namero.

Meaningful information, sensor evolution, and the temporal horizon of embodied organisms. In *Artificial Life VIII*, pages 345.349. MIT Press, 2002.

Ben Shneiderman. *Designing the User Interface: Strategies of Effective Human-Computer Interaction*. Addison-Wesley, 3rd edition, 1998.

Tracy Westeyn, Helene Brashear, Amin Atrash, and Thad Starner. Georgia tech gesture toolkit: Supporting experiments in gesture recognition. In *ICMI 2003: Fifth International Conference on Multimodal Interfaces*. ACM Press, 2003.

81