

Using Temporal Information Distance to Locate Sensorimotor Experience in a Metric Space

Naeem Assif Mirza, Christopher L. Nehaniv, Kerstin Dautenhahn, René te Boekhorst

Adaptive Systems Research Group

University of Hertfordshire

College Lane, Hatfield, AL10 9AB. UK

{N.A.Mirza,C.L.Nehaniv,K.Dautenhahn,R.teBoekhorst}@herts.ac.uk

Abstract- Information distance is used to measure how similar sensorimotor experience is to past experience within a certain temporal horizon. Applied to groups of sensors this gives a mathematical metric on sensorimotor experience over time. We show that for complex data from a robot, large scale similarity of experience can be discovered from the robot perspective, providing a means of building an experiential interaction history.

1 Introduction and Background

We view cognitive systems as embodied dynamical systems structurally coupled (in the sense of Maturana and Varela [1]) to their environments. Cognition then is the development and activity of such a system in response to a history of interactions with its environment (including the social environment). Cognitive structures arise from the recurrent sensorimotor patterns that enable and scaffold increasingly complex perceptually guided action. This view follows the dynamical systems approach of Thelen and Smith [2] that views the dynamic interactions between individuals and environments, and among cognitive processes as the fundamental units of behaviour.

An *autobiographic agent* as defined in [3] dynamically constructs and reconstructs histories of its experience as an enactive embodied dynamical system. These interaction histories are grounded in the physical world, and modify behaviour of the agent while also being modified themselves by further experience [4].

Histories of autobiographic agents that make use of these histories in guiding their actions can be thought of as extending the *temporal horizon* of an agent beyond that of a simple reactive agent, and beyond that of an affective agent driven by emotions, hormones and the like [5]. These agents become *post-reactive* systems acting with respect to a broad temporal horizon by making use of temporally extended episodes in interaction dynamics.

From the perspective of development during the lifetime of an individual (ontogeny), a dynamically constructed interaction history can facilitate development at the borders of known experience. It is known that this is the case for human development which is continually scaffolded by building new capabilities on top of existing ones, with learning proceeding at the periphery of known experience and already mastered interaction skills (“zone of proximal development”) [6].

Thus, it is our long term aim that a robotic agent will be able to use an extended interaction history to modify future

behaviour by recognising that a particular previously experienced episode (for example playing with a familiar toy or interaction partner) is reoccurring, attempting to re-engage in that previous dynamic interaction, and exploring new possibilities beyond that interaction by some modification of behaviour. In doing this the agent will reinforce the memory of the experience while also modifying it with the new experience.

This paper presents a methodology for locating sensorimotor experience in a metric space and making relative comparisons of distance between experiences.

In the remaining sections we explain our motivation and describe the formal metrics used (Section 2), describe the robotic platform used and the investigations conducted (Section 3) followed by a presentation and discussion of the results (Section 4). Finally Section 5 presents our conclusions and summary while Section 6 looks at future directions.

2 Motivation and Metrics

While it is true that an engineer will know in exact detail what a robot has been instructed to do, because of the dynamical structurally coupled nature of the robot environment interaction, it is not trivial to know what the robot is “experiencing” in a particular interaction. We require a method, therefore, that characterises *experience*, in terms of sensorimotor readings, from the robot perspective.

Additionally, it is useful to have a general, universal method that (1) does not rely on particular sensor and actuator configurations or types of robot architecture, and, moreover, (2) allows for changes over time in the sensors, e.g. in terms of layout, size, or resolution. This could be harnessed in ontogenetic development and learning by experience from past interaction history with the environment, where the kinds of interaction have changed or some of the sensors or motors have passed through various ontological developmental stages.

The method we present uses a metric measure of the dissimilarity of temporal sequences of sensors building a metric space in which experiences can be placed and thus compared. Previous work has used the information metric to find structure in uninterpreted sensor-motor data [7],[8] and our work uses similar techniques. We extend the concept of *average information distance sensor-motor phase-plots* introduced in [8] which quantitatively represents the changing

internal geometry of “clouds of experience” and how they change over time, to look instead at the relation in time of where those “clouds of experience” are in the metric space and how their trajectory varies in time.

Related work includes [9] where sensory-motor coordination is investigated and agent-environment interaction “fingerprints” are derived, as well as [10] where mutual information is used to characterise interactions of a simulated agent. An important distinction between this and previous work is that time-developed interaction histories are played out as dynamical trajectories in a geometric space of sensorimotor experiences, whose points correspond to tuples of information sources, using complex physically grounded data from a real robot.

2.1 Information Distance

For any pair of information sources for which a joint distribution can be calculated, it is possible to measure the *information distance* between them. Information distance was shown to be a metric by Crutchfield in [11] (in contrast to related measures such as mutual information) and it is this property in particular we wish to exploit. The method, as used in [7] aligns the discretised time-series of two sensors from a robot to generate their joint distribution by assigning sensor values to bins and counting the frequency of occurrence in each bin.

The *information distance* metric is defined for two information sources X and Y as

$$d(X, Y) = H(X|Y) + H(Y|X) \quad (1)$$

and is measured in *bits*, where *conditional entropy* $H(X|Y)$ is given by

$$H(X|Y) = - \sum_{x \in X} \sum_{y \in Y} p(x, y) \log_2 p(x|y) \quad (2)$$

and similarly for $H(Y|X)$. In this paper X and Y are derived from continuous robot sensors and the joint and conditional probability distributions $p(x, y)$ and $p(x|y)$ are estimated from the time-series of discrete values produced by sampling the continuous values of the sensors for a given period of time. The time period is a window of τ events and the normalised sensor values are quantised into Q bins of equal size. Simultaneously taken readings are used to estimate the joint probabilities.

The effect of quantisation on the entropy estimation from a series of data can produce false peaks of high entropy at the transitions between bins. This can be illustrated by computing the entropy of a hypothetical sensor whose time-series values are taken from a sine-curve. Figure 1 shows a single period of a sine curve constructed from 500 points (time-steps) for which the entropy has been estimated for a moving window of 25 time-steps using 4 different numbers of bins. The entropy over the time-window is plotted on the vertical axis at the point on the horizontal axis corresponding to the end of the time-window. When the numbers of bins is small there are artificial peaks in the entropy curve, for a larger number of bins the curve is closer to the




Figure 1: *Effect of quantisation on entropy estimation.* Sine curve of 500 points (one per timestep) for which the entropy for a moving window of 25 timesteps is estimated for 4 different numbers of bins: 4, 8 16 and 64. Increasing the number of bins reduces the quantisation effect.

expected curve, i.e. low entropy over flat regions, entropy increasing as the slope increases.¹

The quantisation problem is very marked with such ideal smoothly changing data, however; real data has a great deal of variation and noise in it and this serves to mitigate this effect. We demonstrate later that the general shape of entropy curves for real data is less sensitive to the number of bins.

2.2 Temporal Information Distance

At a particular instance in time a group of n sensors will have a particular value in an n dimensional space. Further, for a given time window τ , each sensor will have a time-series of recent values (at times determined by the sampling

¹Note, this is only the case where entropy is estimated by discrete sampling.

rate). Taking all sensors together this can be thought of as describing the *experience of the agent over the time window*.

Using the individual time-series of the sensors for the given window size, the information distance between any two sensors can be estimated as was detailed in [8]. There, distances between component random variables in the experience define a metric space, and morphological measures can reveal the structure of the collection of experiential random variables in that space (e.g. tightness of the collection).

Further, we propose that it is possible to compare any one sensor with itself at another instance in time. Thus, two time-series of equivalent window size can be taken from two different regions in the overall time-series of a particular sensor and the information distance calculated. This would indicate by how much the sensor has changed between the two time regions examined and defines another metric space. Temporal comparisons using information distance have also been used by [12] to detect visual or tactile flow

Thus, for a sensor s , the *temporal information distance* of the sensor between two time regions $(t - \tau, t)$ and $(t' - \tau, t')$ where τ is the window size, using Q bins is given by the information distance $d(\mathbf{s}^t, \mathbf{s}^{t'})$, where \mathbf{s}^t is a random variable whose distribution is estimated from the time-series of values $(s^{t-\tau}, s^{t-\tau+1}, \dots, s^t)$, $\mathbf{s}^{t'}$ is the random variable whose distribution is estimated from $(s^{t'}, s^{t'-1}, \dots, s^{t'-\tau})$ and $d(X, Y)$ is the information distance between two random variables described above.

For a collection of n sensors $S = (s_1, s_2, \dots, s_n)$ two further metric measures can be defined; the *total temporal information distance*

$$D(\mathbf{s}^t, \mathbf{s}^{t'}) = \sum_{i=1}^n d(s_i^t, s_i^{t'}) \quad (3)$$

and the *average temporal information distance*

$$\bar{D}(\mathbf{s}^t, \mathbf{s}^{t'}) = \frac{1}{n} \sum_{i=1}^n d(s_i^t, s_i^{t'}) \quad (4)$$

Both measures satisfy the mathematical axioms of metrics (equivalence, similarity and triangle inequality) because the information distance is a metric [13]. $D(\mathbf{s}^t, \mathbf{s}^{t'})$ is measured in *bits* indicating the total information distance for all sensors between time regions and $\bar{D}(\mathbf{s}^t, \mathbf{s}^{t'})$ is measured in *bits per sensor* indicating on average how much the sensors have changed.

Instead of considering all sensors together it is also possible to consider groups of sensors and look at how they vary with time and also how they vary with respect to each other. As we are interested in agent-environment interactions, a useful split of a robot's sensors would be into read-only sensors (roughly corresponding to exteroceptors or just the environment) and read-write sensors (roughly corresponding to proprioceptors/effectors or internal states of the agent specifically). With these two groups it is then possible to consider the dynamical variation over time of the agent-environment interaction.

Table 1: AIBO Telemetry Collected

Sensors	#	Motors	#
IR-Distance	1	Leg Joint Positions	12
Accelerometers	3	Head Joint Positions	4
Temperature/Battery	2	Tail Joint Positions	2
Buttons	8	Motor Force / Duties	18
Visual	27		
Total Sensors	41	Total Motors	36

3 Platform and Experiments

For our experiments we use the commercially available SONY AIBO² robot. Behaviours were written using the Open Source software Tekkotsu [14] and executed on the AIBO. Sensor/motor data was transmitted at regular intervals (on average 10 frames/sec.) to a workstation over wireless LAN where the data was processed in real-time. For experimental purposes, data was also reprocessed off-line with different parameter values. Experiments were carried out in a low walled 2m×2m arena.

Table 1 summarises the sensors available to the Aibo from which data was collected. The data was grouped into 36 proprioceptive variables representing the agent and 41 sensor variables representing the environment as perceived by the robot. Visual images from the head mounted camera were converted into 27 individual sensors by taking an average of each of the red, green and blue values in each region of a 3×3 grid over the image.

The following preliminary investigations were conducted:

- Test of metric using a sine curve.
- Investigation of effect of varying the window size using real robot data.
- Investigation of effect of varying number of bins (quantisation effect) using real robot data.

Finally, using data from a robot performing a series of behaviours, we use the temporal information distance to construct local and global views of the experiential metric space and to identify experiences which are close in the metric space.

4 Results and Discussion

Test of metric using a sine curve. To test the utility of using the temporal information distance to be able to place experience in a metric space, the method was tested using a simple sine-curve with probabilities estimated using 32 bins. All possible time windows of size 25 were compared with each other and Figure 2 shows a plot of the temporal information distance (on a grey-scale) for every pair of time-regions.

The results indicate that the very tops and bottoms of a sine curve appear different to other regions but very similar to themselves with a gradual change in the degree of similarity. Also, sloping regions appear similar to each

²AIBO is a registered trademark of SONY Corporation




Figure 2: Temporal information distances between all temporal regions of size 25. Data is a sine curve with 1 period covered over 500 timesteps (see Figure 1). Temporal information distance estimated for 32 bins. Information Distance represented on grey-scale. Note that the top and bottom of the sine curve are distance zero from each other.

other although not exactly the same. It is also worth noting that the information distance metric finds similarity with contrary sloping regions in the same manner as if the slopes were in the same direction.

Investigation of effect of varying window size using real robot data. The AIBO robot executed a short simple behaviour (in this case turning in-place) and the sensor data was analysed to investigate the effect of the window size on the method.

The behaviour consisted of 113 timesteps (11.3 seconds) of data for the 77 sensors while the robot turned around approximately 4.5 times. A time region of a given size near the middle was chosen, and the temporal information distance calculated for each sensor from all other possible time regions of the same size within the 113 timesteps available. Note that a complete window is needed to make a comparison. An average for all the sensors was plotted against time. This was repeated for 4 different window sizes and the results shown in Figure 3.

Two effects can be seen. Firstly, increasing window size increases the overall temporal information distance (averaging around 1.3 bits for a window of 10 time-steps to around 2.1 bits for a window of 40). This effect is probably due to the increased number of samples from which to estimate the probability distributions as the window increases. This in turn leads to more accurate (higher) estimations of information distance. Secondly, the overall shape of the curve is qualitatively similar for all the window sizes. This shows that, for this simple behaviour, the window size does not have a great qualitative effect on the information distance. However, for much larger window sizes, there is further smoothing of the curve.

The same tests were run for other examples of turning and also of executing other behaviours, and we conclude that the qualitative similarity over a range of window sizes is typical.

Another point of interest is that only when a region is compared to itself does the temporal information distance become zero; at all other times it is fairly high. This is




Figure 3: Effect of window size on temporal information distance for Aibo data. Graphs show temporal information distances from a central time region (always ending at timestep=80), averaged over all sensors. Data is taken from a robot turning on the spot. The horizontal axis marks the timestep at the end of the time-region being compared. Top to bottom; window sizes 10, 20, 30 and 40.

true even for a time-shift of 1 timestep. Nevertheless, a cyclic pattern of increasing and decreasing similarity can be seen and this corresponds to the cyclic pattern of leg movements and those of the repeating visual sensoric input from turning on the spot.

Investigation of effect of varying number of bins (quantisation effect) using real robot data. The same data used for Experiment 1 was processed to get information distances between a central time region and all other time regions using a range of numbers of bins. Again an average was taken for all of the sensors and the results are shown in Figure 6.

Here, as for window size, it can be seen that the number of bins does not have a qualitative effect on the information distance for real robot data, except for a difference in

Information distances between all regions of window size = 20




Figure 5: Information distances from all regions of size 20 to all other regions. Data from a robot exploring the environment.

experiences⁴ for window size 20 and 10 bins are: 310 (turn), 389 (turn), 402 (turn), 112 (turn), 105 (turn), 693 (turn), 700 (turn), 290 (turn), 363 (walk), 694 (turn). Again, the only experience that is “misclassified” here is 363 which occurs 4 timesteps into a walking phase after a phase of turning. The radius of the sphere containing these experiences is $r = 1.1090$ bits/sensor.

Many seemingly similar experiences are farther away, but this might be because they differ in a manner not immediately apparent to an external observer.

From these results, it appears that the closest experiences in terms of the temporal information distance are of the same type. The results indicate that it might be necessary to compensate for differences in observed behaviour and that given by the experience metric due to the information distance considering experience over a temporal period.

5 Conclusion

The measure *temporal information distance* is described as an information distance between two equivalently sized temporal periods in the time-series of the sensorimotor variables considered as random variables and can be used to place experience in a metric space. This can be used for subsets of sensorimotor variables to describe the sensorimotor experience and to discover experience near to other ones in terms of their distance in the space. We also show that experiences that are near in their temporal information distance are also similar in terms of how we might subjectively

classify them, thus providing a method by which a robot agent might classify its sensorimotor experience from its own perspective.

6 Future Directions

Our future research will continue working towards an agent-constructed experiential interaction history that can be used to modify, predict and guide future behaviour. An agent would map out its experience in a metric space and learn how to move from one experience to another. The metric nature of the space would allow the agent to know if it was moving closer to or further away from any particular experience, potentially allowing it to navigate to familiar or desired interactions. Further, knowing where in the experiential space this experience is located and knowing how to move through that space could allow for the prediction of what might happen next. We will also look for dynamical “attractors” in the space where repeated cycles of behaviour occur and use these as areas from which to explore new experience.

Acknowledgements

The work described in this paper was conducted within the EU Integrated Project RobotCub (“Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours”) and was funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

We would also like to thank Alexander Klyubin and Lars Olsson their input in clarifying the quantisation issues.

⁴Ignoring start-up time when no behaviour was being executed

Bibliography

- [1] H. R. Maturana and F. J. Varela, *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: New Science Library (Shambhala), 1987.
- [2] L. Smith and E. Thelen, *A Dynamic Systems Approach to Development*. MIT Press, 1993.
- [3] K. Dautenhahn, “Embodied cognition in animals and artifacts,” in *Proc. AAI FS Embodied Cognition and Action*, pp. 27–32, AAAI Press, 1996. Technical report FS-96-02.
- [4] C. Nehaniv and K. Dautenhahn, “Embodiment and memories—algebras of time and history for autobiographic agents,” in *Cybernetics and Systems '98: Proceedings of the Fourteenth European Meeting on Cybernetics and Systems Research* (R. Trappl, ed.), vol. 2, pp. 651–656, Vienna: Austrian Society for Cybernetic Studies, 1998.
- [5] C. L. Nehaniv, D. Polani, K. Dautenhahn, R. te Boekhorst, and L. Cañamero, “Meaningful information, sensor evolution, and the temporal horizon of embodied organisms,” in *Artificial Life VIII*, pp. 345–349, MIT Press, 2002.
- [6] L. Vygotsky, *Mind and society: The development of higher mental processes*. Cambridge, MA: Harvard University Press., 1978.
- [7] L. Olsson, C. L. Nehaniv, and D. Polani, “Sensory channel grouping and structure from uninterpreted sensor data,” in *IEEE NASA/DoD Conference on Evolvable Hardware 2004*, pp. 153–160, IEEE Computer Society, 2004.
- [8] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst, “Using sensory-motor phase-plots to characterise robot-environment interactions,” in *6th IEEE International Symposium on Computational Intelligence in Robotics and Automation* In Press, 2005.
- [9] I. J. A. te Boekhorst, M. Lungarella, and R. Pfeifer, “Dimensionality reduction through sensori-motor coordination,” in *Proceedings of the joint International Conference on Artificial Neural Networks and Neural Information Processing, Istanbul, Turkey* (O. Kaynak, E. Alpaydin, E. Öja, and L. Xu, eds.), Lecture Notes in Computer Science 2114, (Berlin), pp. 496–503, Springer, 2003.
- [10] G. Tarapore, M. Lungarella, and G. Gómez, “Fingerprinting agent-environment interaction via information theory,” in *Proceedings of the 8th International Conference on Intelligent Autonomous Systems* (F. Groen, N. Amato, A. Bonarini, E. Yoshida, and B. Kröse, eds.), (Amsterdam, The Netherlands), pp. 512–520, 2004.
- [11] J. Crutchfield, “Information and its metric,” in *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves* (L. Lam and H. Morris, eds.), pp. 119–130, New York: Springer-Verlag, 1990.
- [12] L. Olsson, C. L. Nehaniv, and D. Polani, “Discovering motion flow by temporal-informational correlations in sensors,” in *5th International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems (EpiRob05)*, in press.
- [13] C. L. Nehaniv, “Sensorimotor experience and its metrics: Informational geometry and the temporal horizon,” in *IEEE Congress on Evolutionary Computation*, 2005.
- [14] D. Touretzky and E. Tira-Thompson, “Tekkotsu,” *Carnegie Mellon University, Online Open Source*, <http://www.tekkotsu.org>, 2004.




Figure 6: Effect of number of bins on information distance for Aibo data. Graphs show temporal information distances from a central time region averaged over all sensors. Data is taken from a robot turning on the spot. The horizontal axis marks the timestep at the end of the time-region being compared. Top to bottom; bins 5, 10, 20, 30, 40 and 60.




Figure 7: Information distances from region 200-220 (walking) to all other temporal regions. Data from a robot exploring the environment. Time 200-220 corresponds to an early part of a walk. For comparison, the observed behaviour is shown on the lower part of the graph with the vertical lines marking changes in behaviour. Top: All sensors, Middle: Exterioceptors, Bottom: Proprioceptors.




Figure 8: Information distances from region 290-310 (turning) to all other regions. Data from a robot exploring the environment. Time 290-310 corresponds to a part of a turn. For comparison, the observed behaviour is shown on the lower part of the graph with the vertical lines marking changes in behaviour. Top: All sensors, Middle: Exterioceptors, Bottom: Proprioceptors.