Anticipating Future Experience using Grounded Sensorimotor Informational Relationships

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

Adaptive Systems Research Group, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, U.K.
c.l.nehaniv@herts.ac.uk

Abstract

Operational definitions and applications of the sensorimotor *experience* of an artificial embodied organism are presented along with a mathematical metric for distance between experiences based on Shannon information. We describe a simple robotic experiment that illustrates how an artificial embodied agent can use its own history of experience combined with the experience metric to predict future experience. Present sensorimotor experience is used to find the most similar past experience using the geometry of its growing and changing experience metric space. This is then used to ground the ontogeny of autonomous prospective capability in interacting with the environment, e.g. to anticipate forthcoming changes in environment based on temporally extended past experiences.

Introduction

Increasingly, the importance of embodiment and situatedness within complex and rich environments are becoming recognized as a crucially important factors in engendering intelligence in an artifact (cf. for example Clancey (1997); Pfeifer and Bongard (2007), and the philosophical position regarding 'structural coupling' of Maturana and Varela (1987)). Living organisms in particular experience and re-experience particular recurring patterns of trajectories of interactions with the environment through their sensing and acting; and these habitual trajectories can form the basis of prospection, further development, and adaptation (Varela et al., 1991). [1]

Moreover, it is in how an artificial agent develops its capabilities over its life-time of interactions (*ontogeny*) that is important in building a *grounded* intelligence, able to adapt to unknown and changing environments (including long- and short-term variations in its embodiment and in its sensory or motor repertoire). Especially given the complexity of interactions in natural environments, and the richness of sensors available to modern robots, whose properties change over time in different environments or with changing embodiment, it is largely infeasible and impractical to attempt to foresee and model the situations a robot (or other artificial agent) may encounter and how to adapt to them in advance (e.g. Brooks (1999)). Instead, autonomous methods for bootstrapping development without prior knowledge of the structural coupling relationship based on enactive construction and development of intelligence behaviour warrant investigation, both from the perspectives of engineering applications as well as from the viewpoint of a generalized biology. Building on basic 'phylogenetic' capabilities, such an approach is hypothesized to allow for a basic of autonomous, enactive development in embodied models of developmental cognitive systems with expanded temporal horizon of their perception and action (Nehaniv et al. (2002),Vernon et al. (2007), Mirza et al. (2007)).

Our goal is to research methods that can be used by an artificial embodied agent to develop its capabilities through its ongoing interactions with its environment, while scaffolding its adaptation on the basis of previous experience and previously achieved adaptation. In earlier work we introduced formal mathematical metrics on sensorimotor experience and its geometry, as well as heir use as part of a developmental architecture for robots that bases future action on previous experience (Nehaniv, 2005; Mirza et al., 2005a, 2007). In this paper we present results from a robotic experiment that illustrates how a history of embodied experience, combined with a metric measure for comparing experiences, can be used to predict temporally extended future experience. This is an important result for our developmental architecture as it demonstrates the efficacy of the metric measure, and in turn its suitability for directing future action and behaviour based on the individual's past experience.

Other Related Work. Olsson et al. (2006) use information distance to develop basic sensorimotor maps in interaction with the environment, beginning from raw uninterpreted sensors. Independently of our work, Oates et al. (2000) have also described experiences as a time-series of multi-variate sensorimotor data (which is essentially identical to our operational definition of experience), but computing distance

[1] This work was conducted within the EU Integrated Project RobotCub ("Robotic Open-architecture Technology for Cognition, Understanding, and Behaviours"), funded by the EC through the E5 Unit (Cognition) of FP6-IST under Contract FP6-004370.

between time-series and clustering experiences to produce prototypes. Experiences are associated with the actions that initiated them, so robot can generalize about potential outcomes of its actions. Distances between experiences are calculated by using Dynamic Time Warping followed by measuring the area between the curves, and clusters formed by taking averages of time-warped experience curves. In contrast, our framework uses an information-theoretic metric on such experiences.

Kaplan and Hafner (2005) use information distances between sensors in an Aibo robot to compare simple behaviours of the robot. In that method, rather than reducing the dimension by summation within groups as we have done, they consider distances between different behaviours as distances between the full matrix of distances between all sensors. Long continuous examples of each behaviour (1000 timesteps) are used, and the whole sequence used rather than a moving window. The resulting distances between behaviours are shown as a projection onto a two-dimensional map, and they find that similar behaviours group together. This research supports the view that robot behaviour can be clustered using information relationships between sensor time-series. However, the incremental formulation of our approach allows us to propose a system that can be used for ontogeny, and the use of the experience metric allows for better comparison of past behaviour and experience.

Continuous Case-Based Reasoning (CCBR) (Ram and Santamaria, 1997) has many similarities to the approach described here. However, in our approach the information metric allows for a more robust comparison of sensorimotor details concentrating on the statistics of the particular time-series, and so better able to recognize regularities in time-series than a simple Euclidean metric. Also, the metric nature of the space is also able to recommend a number of increasingly distant matches (neighbours) and is able to weight their similarity along with a qualitative value from the environmental feedback to provide, potentially, more appropriate actions.

Sensorimotor Experience and Metric

A robot or other embodied agent's entire view of the world is experienced through its sensors, including those that measure internal factors such as temperature, actuator positions, and other more general internal variables. Any sensor can be modelled as a random variable $\mathcal{X}$ changing with time, taking values $X(t) \in \mathcal{A}_{\mathcal{X}} = \{x_1, \ldots, x_m\}$ from a probability distribution $\mathcal{P}_{\mathcal{X}}$. Time is taken to be discrete (i.e. t will denote a natural number). A robot's experience, then, can be considered as the stream of all readings $(X^1(t), \ldots, X^n(t))$ from all these variables $\mathcal{X}^i$ over a given time period (i.e. $t \in [t', t' + h]$ for some *temporal horizon* $h > 0$). This is a purely operational sensorimotor view of experience and, by itself, says nothing about the quality or meaning of that experience.

Formally, an agent's *experience* from time t over a temporal horizon h can be defined as

$$E(t,h) = (\mathcal{X}^1_{t,h}, \ldots, \mathcal{X}^N_{t,h}) \quad (1)$$

where $\mathcal{X}^1_{t,h}, \ldots, \mathcal{X}^N_{t,h}$ is the set of random variables available to the agent constructed or estimated according to time-series of sensorimotor readings from N sensorimotor variables $(X^1, \ldots, X^N)$ ending at time t with a horizon h timesteps (from time $t-(h-1)$ to t).

Experience Metric

Given a definition of Sensorimotor Experience and the information metric, a formal measure of distance between experiences can be defined. This is useful as it allows a direct, scaled comparison between different sets of sensorimotor readings of a robot or agent. A metric for comparison of sensorimotor experiences is important as it is then possible to talk of proximity and distance between different experiences in a quantitative and geometrically meaningful way.

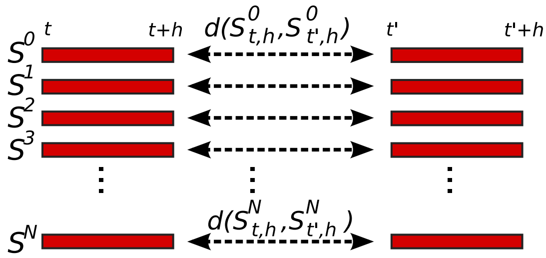

Figure 1: *Experience Metric.* A visual illustration of the experience metric. Each experience is shown as a collection of sensor readings of length h starting at time t and t'. The information distance between each respective sensor over time is summed to give the Experience Metric.

We define the *Experience Metric*, a metric on experiences of temporal horizon h, as

$$D(E, E') = \sum_{k=1}^{N} d(\mathcal{X}^k_{t,h}, \mathcal{X}^k_{t',h}) \quad (2)$$

where $E = E(t,h)$ and $E' = E(t',h)$ are experiences of an agent at time t and t' over horizon h, and d is the Crutchfield-Rényi information metric (Crutchfield, 1990), or more simply, the *information distance* between jointly distributed random variables. That is, $d(\mathcal{X}, \mathcal{Y}) = H(\mathcal{X}, \mathcal{Y}) - I(\mathcal{X}, \mathcal{Y})$, where H denotes entropy and I denotes mutual information (see (Cover and Thomas, 1991) for an introduction to these concepts of information theory)[2]. D is measured in *bits*; see also Figure 1. That D is a metric follows from the fact that the metric axioms (equivalence, symmetry, and the triangle

[2] $d(\mathcal{X}, \mathcal{Y}) = 2H(\mathcal{X}, \mathcal{Y}) - H(\mathcal{X}) - H(\mathcal{Y})$ and is estimated directly from the frequency distributions of binned sensor values.

inequality) hold for each of the components in the summation, since d is a metric (Nehaniv, 2005). For a visual proof that d (and hence D) is a metric, see (Nehaniv et al., 2007).

Earlier Experiments

In Mirza et al. (2005b) we describe an experiment showing ball-path prediction using the experience distance measure. In that experiment an Aibo robot (see Figure 2 and below) remained stationary while a ball was moved in view of its head mounted camera. The predicted ball path was plotted in real-time overlaid on the images from the camera. This experiment illustrated that sensor experience can be used to match experience successfully. This experiment builds on that result, but uses the full embodied experience to match previous experience. The camera images do not, by themselves, give information about the position of the ball so self-experience is important.

Experiment

Interactive Path Prediction

A simple robotic experiment was devised that would illustrate how an artificial embodied agent can use its own history of experience combined with the experience metric described above to predict future experience. The robot follows the motion of a ball moved in front of it by using a simple reactive behaviour to adjust its head motors to attempt to centre the ball in its field of vision. The robot continually builds a metric space of experiences from its ongoing sensorimotor experience, including its own proprioceptive sense of movement arising through interaction with the environment. A closest historical experience, in terms of experience distance, to the current one is then found. Experiences temporally following the historically closest experience then provide a model for anticipation of future experience. How good this model is depends on both the predictability and consistency of the environmental interaction as well as how "good" the historical matching is. Thus, the analysis of the experiment focuses on measuring how well matched the historical experience is to the current one. Note that predicting the trajectory of the tracked object corresponds to prospection regarding part of a future temporally extended interval of sensorimotor experience.

It is important to note that, the robot is not matching current ball position with previous ball position, rather all sensory and motor variables are used as information sources to detect similarity between experiences.

Implementation and Experimental Setup

The robot used was a Sony Aibo ERS-7. The control and sensory collection software was implemented in Java with URBI (Baillie, 2005) providing the robot control layer and ball detection. Sensor readings are sent over wireless to a personal computer approximately every 80-120ms. Reception of each frame of data defines a *timestep*. Video images were received from the robot head camera approximately every 400ms, however visual sensors were computed at the rate of the sensor data using the most recent image from the camera. Experiences were formed from data streams from 33 internal sensors (including proprioceptive motor positions and infrared distance measurements, and 9 sensors formed from average pixel values in a 3×3 grid over the image.

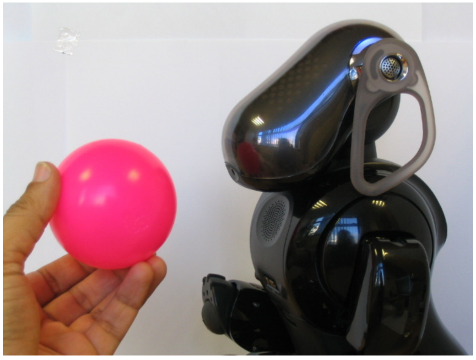

Figure 2: *Sony Aibo ERS-7, and Pink Ball*

The robot was stationary in a "sitting" position, with the head pointed forward (Figure 2). A pink ball was moved in the air in view of the robot's head camera at a distance of approximately 30cm. No particular effort was made to "sanitize" the environment to aid ball-detection against the background. Thus, it is likely that other items in the environment provided potentially useful information about any interaction. The robot executes a continuous reactive behaviour to follow the motion of a ball with its head. The algorithm is simple, making appropriate incremental adjustments to the neck, headTilt and headPan motors, such that the position of the ball is brought closer to the centre.

The metric space creation and prediction was implemented in Java and ran on-line in real-time. The horizon length of the experiences was $h = 20$ timesteps or approximately 1700ms. The data was quantized into $Q = 10$ bins in the probability distribution estimation algorithm.

The ball was moved such that the time for the ball to describe a circle (or to move horizontally or vertically for a complete cycle) was 6-7 seconds. Thus the horizon length was shorter than, but of the same order of magnitude as, a single cycle of the repeated behaviour and the experiences would comprise approximately a half of a cycle.

The full interaction sequence lasted 965 timesteps (~ 84 seconds) constituting 945 experiences of horizon length $h = 20$. The movements of the ball consisted of a number of horizontal and vertical movements, and a number of clockwise circles; see Table 1.

Visualizing Ball Path: A projection of the current ball position relative to the robot is plotted in two dimensions by estimating the direction in which the head is pointed from

Table 1: Path Prediction Experiment - Sequences of Movements (TS denotes time step number)

Start TS	End TS	Movement Type	Iterations
91	185	Horizontal, Left to Right	2 full
201	272	Vertical movements, Top to Bottom	2 full
283	361	Horizontal, Right to Left	1 full
376	453	Vertical, Top to Bottom	2 full
463	534	Horizontal, Right to Left	1 full
548	593	Vertical, Top to Bottom	1 full
607	852	Circular, Clockwise	4 full
866	929	Vertical, Bottom to Top	2 full

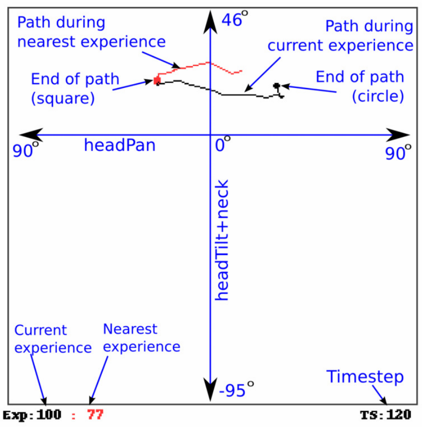

Figure 3: *Ball Path Traces.* The diagram shows the parts of the ball path diagrams used to visually analyse the traces of the ball in a neck-centred coordinate system derived from motor positions. See Figures 6 and 7.

the positions of three motors contributing to head motion. The coordinates for the ball position in the plot are given by:

$$(x, y) = (W \times headPan, H \times (headTilt + neck)/2)$$

where W and H are the image width and height, and $headPan, headTilt$ and $neck$ are the motor values at any instant normalized into the range $(0, 1)$. See the explanatory diagram of Figure 3. Note that the plots are created for analysis of the experiments, and this abstraction of the sensoriomotor flow is *not* available to the robot. Instead it allows us as external observers to gain insight into what the robot 'expects' will happen in an interval of the near future based on its own previous experiences, and how accurate these expectations are (again to an external observer).

Error Measurements: Two different measurements of path error were used. The first measured the sum of the Euclidean distance between each corresponding point of the paths. The second calculated a vector direction for each path and returned the angular difference in radians between the vectors as the error.

Table 2: Improvement of Experience Matching Over Time

Type	Iteration	Number $< \pi/4$	Total Number	Percentage $< \pi/4$
HORIZ	1	0	41	0.0%
HORIZ	2	27	73	37.0%
HORIZ	3	25	75	33.3%
HORIZ	4	27	72	37.5%
VERT	1	0	34	0.0%
VERT	2	8	51	15.7%
VERT	3	15	30	50.0%
VERT	4	42	61	68.9%
VERT	5	32	52	61.5%
VERT	6	27	49	55.1%
CIRCLE	1	9	65	13.8%
CIRCLE	2	13	54	24.1%
CIRCLE	3	27	66	40.9%
CIRCLE	4	31	63	49.2%

Results and Analysis

Figures 4 and 5 show, using different methods of error estimation, the error between the current path and the path corresponding to the nearest previous experience in terms of information distance. Figures 6 and 7 show traces of the paths from experiences in regions where horizontal and vertical movements were taking place. As can be seen from the traces, which are selected from regular intervals, it is often the case that the paths are similar and so the experiences are well matched. However, the objective measure of error indicates that the actual path is not exactly the same. This is to be expected as there do not exist any *precisely* identical experiences in a real situation.

The opposite direction path (but of the same type) is regularly matched. As the sensors are not biased left or right, and the experience distance measure is the sum of information distances between variables, then a symmetric error such as this is likely. Indeed, such experiences are *informationally* very close to their 'opposites'. Out-of-phase periodic variables can have a small or zero[3] information distance.

In terms of angle, the error is less than $\pi/4$ (*i.e.* closer to parallel than orthogonal) 55.13% of the time and is greater

[3]Variables that have a zero information distance are *recoding equivalent* and are not necessarily identical (see Crutchfield, 1990).

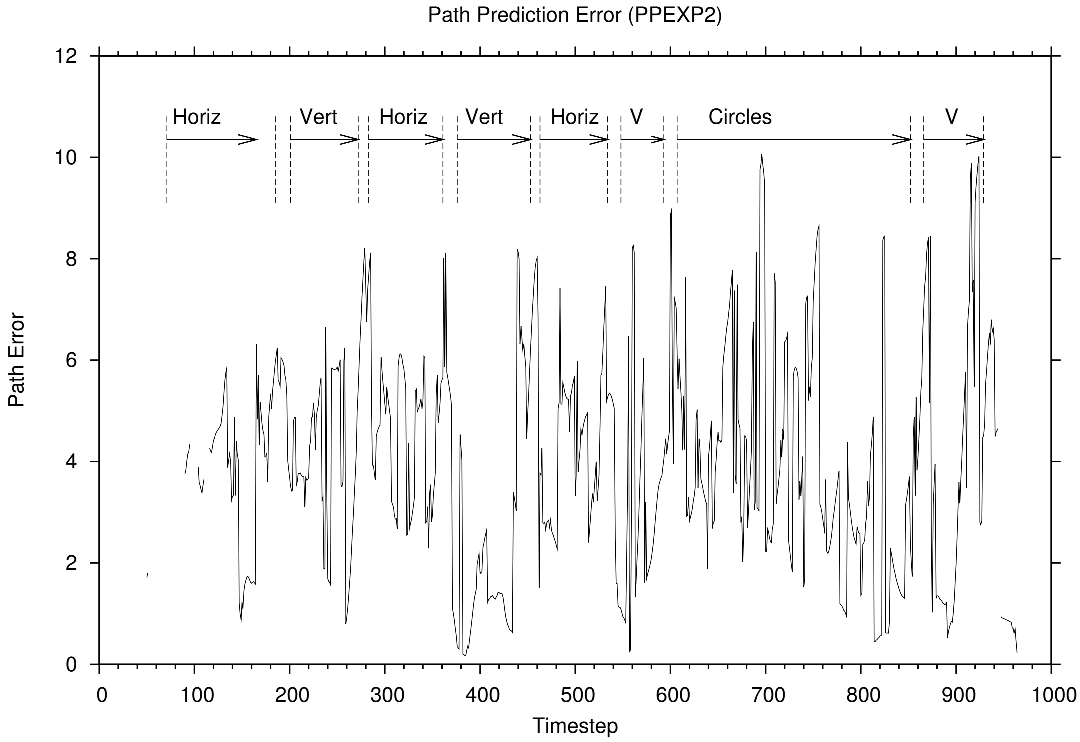

Figure 4: Euclidean distance (error) between the paths of the ball during the current and nearest previous experience. The error is often exaggerated as experiences of paths of the same type but opposite direction are often matched. The top part of the graph shows the behaviour (See Table 1). The *Path Error* (pixels) in this case is the sum of the Euclidean distance between corresponding points. Temporal horizon $h = 20$, number of bins $Q = 5$.

than $3\pi/2$ (*i.e.* closer to opposite than orthogonal) 29.21% of the time. This indicates that the path and therefore the experience is generally well matched, however due to the nature of the measure, experiences from the opposite phase in a cycle are often selected. This error is compensated for in Figure 5 by reflection about $\pi/2$. It is interesting to note the opposite phase corresponds to time-reversed motion, and that the present metric relies on probability distributions constructed from sensorimotor flow and that these distributions do not encode the directionality of time.

Examining the progression of the error over time in these data, one would expect to see an improvement as the same kinds of behavioural interaction are re-experienced. How the matching of experiences improves over time is examined, referring to Table 2 and Figure 5. During the horizontal motions after one full cycle, 37% of experiences can be matched to similar ones in the history. Vertical motions show that the success rate peaks at 68.9% with the 4th presentation. The success rate drops slightly thereafter as there are more experiences to select from. The Circle movements also show marked improvement as experience grows. The initial 13.8% success rate of the very first circular motion reflects the fact that parts of the circular motion are being matched with previous horizontal and vertical experiences, with some limited success, even before any such motions had been observed.

Conclusions

The work describing the construction and use of information metrics for the comparison of robot behaviour demonstrates achievement of a degree of temporally extended prospection by an embodied agent, based on its raw sensorimotor experience. The experience metric was first described in (Mirza et al., 2005a) and with mathematical proofs of the mathematical metric properties along with some alternative metrics on experience in (Nehaniv, 2005). As mentioned, an operational formulation of experience (but not of the metric) was previously described in (Oates et al., 2000). A non-metric measure of distance between experiences was described there that used the area between time-warped experience curves. The fact that independent research groups both developed essentially the same notion operationalizing an agent-centred definition of experience suggests that this definition is a natural one.

Experiments were described that use fairly large numbers of robotic sensors to describe robotic experience such that a simple sort of prediction can be achieved by the matching of present experience with experiences in the history and

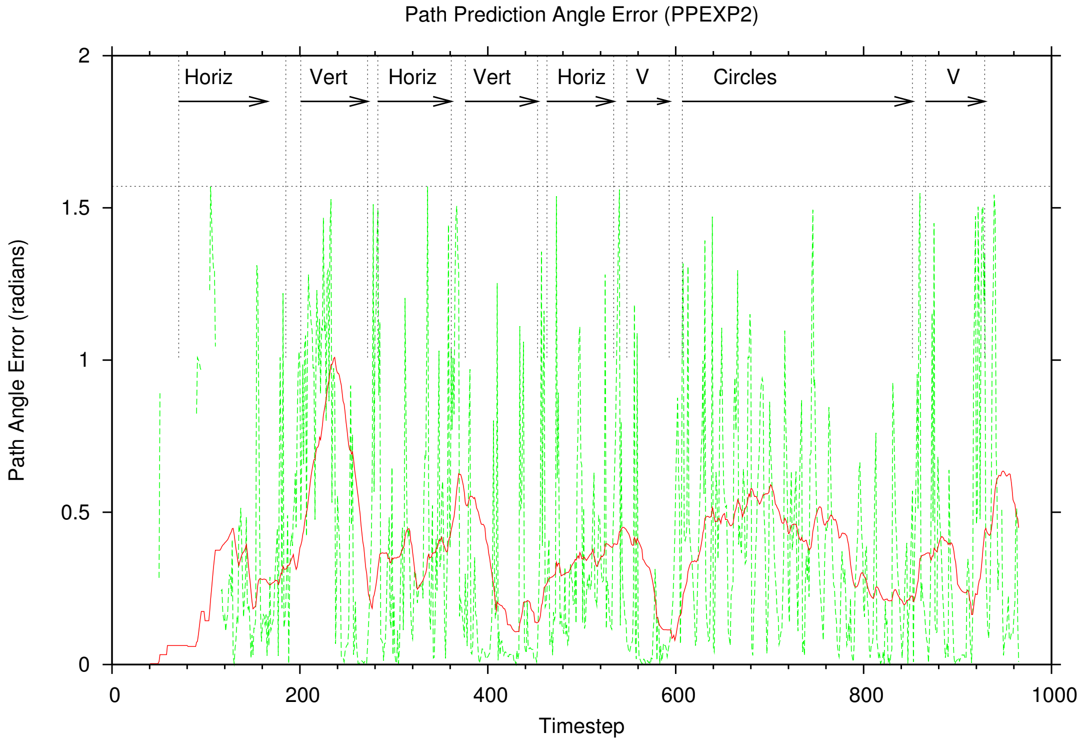

Figure 5: Angle error and the average angle error (over the last 40 timesteps) between the paths of the ball during the current and nearest previous experience. The graph shows the error reducing, on average, *within* a given behaviour sequence. The top part of the graph shows the behaviour (See Table 1). The *angle error* is the difference in radians between the vector direction of each path. For errors$>\pi/2$, $\pi - error$ is shown (reflection about $\pi/2$). Temporal horizon $h = 20$, number of bins $Q = 5$.

extrapolating forward from the matched past experience. It was found that *proximity in terms of experience metric corresponds well with an external observer's notion of similarity of experience.* Future research may consider using the anticipated experience for active perception and in human-robot interaction.

The sensorimotor variables were treated by the autonomous robot in an uninterpreted "agnostic" manner, that is, no sensor is regarded as being different from any another or special in any way, in terms of finding close experiences. This performance was achieved despite many of the sensors not providing any seemingly useful information about the current experience. Proprioceptive motor experience was important in this experiment in determining the experience and matching it to the appropriate past experience.

The capability of the experience metric to find suitable matching experiences was found to increase as more examples of a particular type of behaviour were presented. This appears to level-off, and potentially become worse as more examples are presented. However, the experiments described had too short a run time for a definitive conclusion to be drawn on the latter observation. Another important aspect of the experience metric is that it appears to confuse a behaviour with its 'opposite' (phase-shifted or time-reversed counterparts), as these are informationally nearly identical. This can be seen clearly in both the simple and interactive ball-path prediction experiments as opposite direction of path.

Needless to say, the ontogeny of prospective ability of children and other mammals is an extended process lasting years and we cannot yet hope to mirror its complexity and success in artificial systems, although the work presented here suggests that we have made a small start in this direction.

References

Baillie, J.-C. (2005). URBI: Towards a universal robotic low-level programming language. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems 2005*. http://www.urbiforge.com.

Brooks, R. (1999). *Cambrian Intelligence: The Early History of the New AI*. MIT Press, Cambridge, Mass.

Clancey, W. J. (1997). *Situated Cognition: On human knowledge and computer representations*. Learning in doing: Cognitive and compuational perspectives. Cambridge University Press.

Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. Wiley, New York.

Figure 6: *Head Movement Traces and Matched Historical Traces for Prediction.* Images are from evenly spaced timesteps from three separate *horizontal* movement regions starting at timestep TS=120, 290 and 470. Each diagram shows the path of the ball, as determined by robot head movements, for both the current experience at that timestep (dark line) and for the matched (nearest previous) experience (red/grey line). Path direction indicated by circle/square at the end of the path. (See Figure 3). The angle error between the path directions is used to analyse how well the path and thus experience are matched. Temporal horizon $h = 20$, number of bins $Q = 5$.

Crutchfield, J. (1990). Information and its metric. In Lam, L. and Morris, H., editors, *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves*, pages 119–130. Springer-Verlag, New York.

Kaplan, F. and Hafner, V. (2005). Mapping the space of skills: An approach for comparing embodied sensorimotor organizations. In *Proc. 4th IEEE International Conference on Development and Learning (ICDL-05)*, pages 129–134. IEEE.

Maturana, H. R. and Varela, F. J. (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. New Science Library (Shambhala), Boston.

Mirza, N. A., Nehaniv, C. L., Dautenhahn, K., and te Boekhorst, R. (2005a). Using temporal information distance to locate sensorimotor experience in a metric space. In *Proc. 2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 150–157, Edinburgh, Scotland. IEEE Press.

Mirza, N. A., Nehaniv, C. L., Dautenhahn, K., and te Boekhorst, R. (2007). Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny. *Adaptive Behaviour*, 15(2):167–187.

Mirza, N. A., Nehaniv, C. L., te Boekhorst, R., and Dautenhahn, K. (2005b). Robot self-characterisation of experience using trajectories in sensory-motor phase space. In *Proc. 5th International Workshop on Epigenetic Robotics*. Lund University Cognitive Studies.

Nehaniv, C. L. (2005). Sensorimotor experience and its metrics. In *Proc. 2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 142–149, Edinburgh, Scotland. IEEE Press.

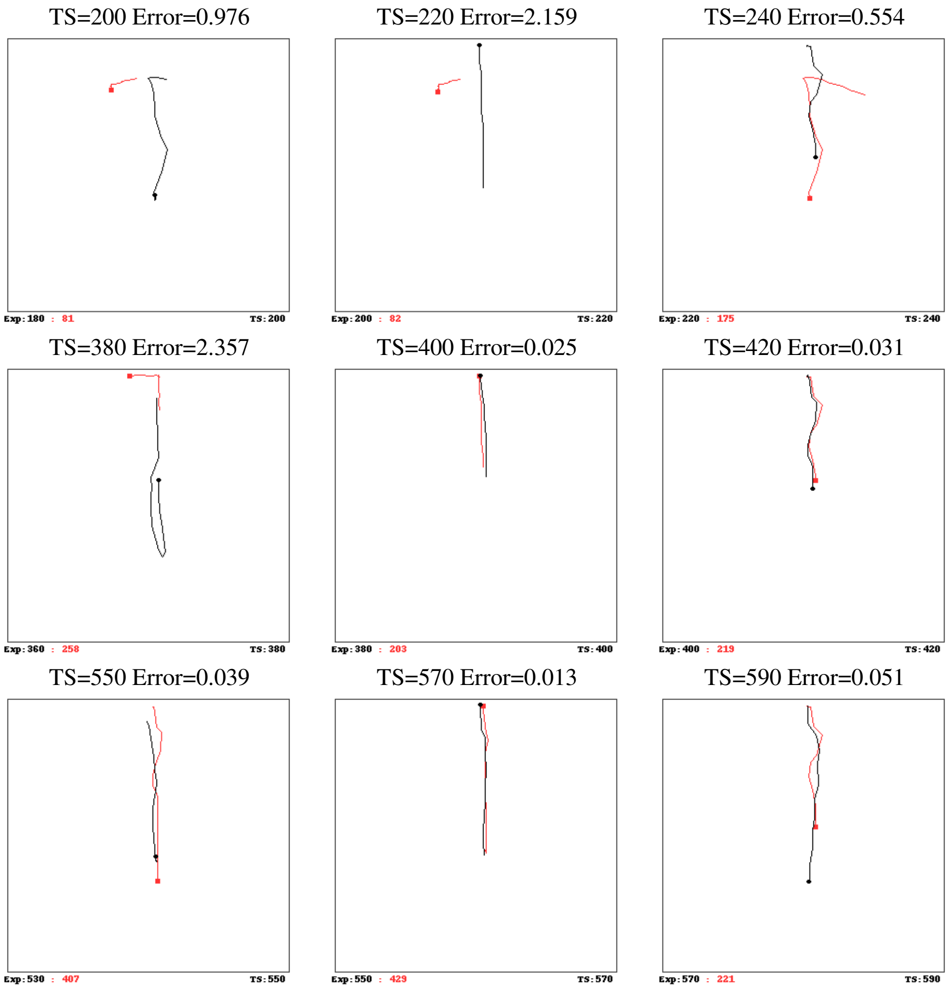

Figure 7: *Head Movement Traces and Matched Historical Traces for Prediction.* Images are from evenly spaced timesteps from three separate *vertical* movement regions starting at timestep TS=200, 380 and 550. Each diagram shows the path of the ball, as determined by robot head movements, for both the current experience at that timestep (dark line) and for the matched (nearest previous) experience (grey line). Path direction indicated by circle/square at the end of the path. (See Figure3). The angle error between the path directions is used to analyse how well the path and thus experience are matched. Temporal horizon $h = 20$, number of bins $Q = 5$.

Nehaniv, C. L., Mirza, N. A., and Olsson, L. (2007). Development via information self-structuring of sensorimotor experience and interaction. In *50 Years of Artificial Intelligence*, volume 4850 of *LNCS*, pages 87–98. Springer.

Nehaniv, C. L., Polani, D., Dautenhahn, K., te Boekhorst, R., and Cañamero, L. (2002). Meaningful information, sensor evolution, and the temporal horizon of embodied organisms. In *Artificial Life VIII*, pages 345–349. MIT Press.

Oates, T., Schmill, M. D., and Cohen, P. R. (2000). A method for clustering the experiences of a mobile robot that accords with human judgements. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence*, pages 846–851, Menlo Park, California. AAAI Press.

Olsson, L., Nehaniv, C. L., and Polani, D. (2006). From unknown sensors and actuators to actions grounded in sensorimotor perceptions. *Connection Science*, 18(2):121–144.

Pfeifer, R. and Bongard, J. (2007). *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press.

Ram, A. and Santamaria, J. C. (1997). Continuous case-based reasoning. *Artificial Intelligence*, 90(1-2):25–77.

Varela, F. J., Thompson, E., and Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press, Cambridge, MA, USA.

Vernon, D., Metta, G., and Sandini, G. (2007). The icub cognitive architecture: Interactive development in a humanoid robot. In *IEEE International Conference on Development and Learning, Imperial College, London*.