Using Sensory-Motor Phase-plots to Characterise Robot-Environment Interactions*

Naeem Assif Mirza, Chrystopher L. Nehaniv, Kerstin Dautenhahn, René te Boekhorst
Adaptive Systems Research Group
School of Computer Science, University of Hertfordshire
Hatfield, Hertfordshire, AL10 9AB. United Kingdom
{*N.A.Mirza,C.L.Nehaniv,K.Dautenhahn,R.teBoekhorst*}*@herts.ac.uk*

***Abstract*— Information theoretic methods are used to characterise and identify robot-environment interactions, with a view to using these to build an embodied interaction history from the robot's perspective. A bottom-up approach is taken using uninterpreted raw sensor and motor data. Interactions are analysed by calculating the Average Information Distance (AID) between all sensors and motors over a moving time window and used to create 2-dimensional "phase-plots" that can be thought of as describing the current interaction. Sensor-Motor AID Phase-plots are shown to be able to distinguish simple behaviours among a sequence of behaviours.**

***Index Terms*— Interaction History, Information Theory, Embodied Cognition, Ontogenetic Robotics**

I. Introduction

A goal of research into embodied cognition is to be able to produce systems that act on a time horizon beyond that of the reactive or affective. These might be called *post-reactive systems* [1] or *autobiographic agents* [2]. For an autonomous robot to be able to draw on past experience to affect its future behaviour, it is useful to be able to compare and contrast current experiential trajectories with those experienced in the past. We suggest therefore, that a necessary step toward making use of interaction history is the ability to identify, characterise and distinguish experiences.

This paper describes an experimental investigation of a method by which a robot can characterise and identify interactions with its environment in terms of its particular embodiment. Uninterpreted raw sensory and motor data is simplified using information theory [3] and used to distinguish and characterise periods of behaviour.

We use a simple average of the *information distance metric* [4] between all variables using their conditional entropies estimated over a moving time-window. This results in just two values changing with time, one each for sensors and motors, that can then be plotted to produce a *Sensor-Motor Average Information Distance (AID) Phase-plot*. We investigate whether such radically informationally compressed phase-plots can characterise the type of interactions that the robot is experiencing in the environment.

*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.

The information metric has been used to find structure in uninterpreted sensor-motor data in [5] and our work uses similar techniques. Related work also includes [6] where sensory-motor coordination is investigated and agent-environment interaction "fingerprints" are derived. Our approach also goes beyond this and other related work [7] in that, using an information distance metric, we are able to characterise histories of agent-environment interactions as dynamical trajectories in a geometric space, whose points correspond to information sources, using complex physically grounded data from the perspective of a real robot.

We start by reviewing the relevant background to embodied cognition, interaction history and information theory and then describe our particular method. The experiments are then described and the results presented and discussed. Finally we look at the next steps in producing a method for a robot to build and use an interaction history peculiar to its own embodiment.

A. Embodied Cognition and Research Orientation

Our model of cognition is the development and activity of an embodied dynamical system, structurally coupled to its environment (including the social environment), which develops in its sophistication and capacities in response to particular interaction histories. In line with a dynamical systems perspective [8] [9], social and environmental interaction can be envisaged as "structural coupling" (following the description of Maturana and Varela [10] and as applied to robots [11]). This is also related to enactive cognition [12], the emergence of cognitive structures from the recurrent sensorimotor patterns that enable and scaffold increasingly complex perceptually guided action.

Embodiment, then encompasses both agent and environment and the dynamical coupling of both. The agent for our purposes here comprises the physical body, and all its internal states, and positional configurations. An agent may have access to internal information via proprioception of positions of motors and limbs, and motivational and homeostatic variables among others. Environment in general is everything external to the agent. The agent experiences the environment as a projection onto its sensor surfaces which include touch, vision, sound as well as sensors such as infra-red distance sensors. Here we consider the dynamics of agent-environment interactions in terms of the interplay between internal state and motor variables (under control of the agent) and sensor variables (read-only by the agent)

We also take a bottom-up, agent-centred view, looking at the raw data values presented to the robot and making as few as possible value-based or design-influenced decisions as possible regarding the interpretation of that data. We would like inherent structure in the data to emerge without presupposing the nature of that structure. In this paper we take a simplistic view of sensory-motor data available to the agent but would see future work allowing the agent to infer and develop more structure in the data allowing for more detailed interaction histories to emerge.

Our approach is also to look at creating usable interaction histories from an ontological developmental perspective. That is, the histories and structures emerging from the data will build on previous development, finding new structure at the boundaries of familiar experience. This particular work represents a very early developmental stage where nothing is known about the structure of the data and vision is coarse-grained.

B. Interaction History

This work is concerned with exploring techniques for building embodied agents which dynamically construct and reconstruct their experiential history (*autobiographic agents* as defined in [2]). These histories are grounded in the physical world, and modify behaviour of the agent while also being modified themselves by further experience [13].

Histories of autobiographic agents 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 [14]. These agents become *post-reactive* systems acting with respect to a broad temporal horizon by making use of temporally extended episodes in interaction dynamics.

In our view, the type of information that is encoded in an experiential history extending back along the time horizon of the agent will probably be compressed or informationally reduced in some way, while still encoding the salient aspects of the original interactions in terms of meaningful information [14] and utility to the agent. A mathematical expression of the differing resolutions of information at different distances along the time horizon are *semigroup expansions of time* [13]. Semigroups can be constructed of sequences of experiences, and expansions reveal higher level structure in sequences of events and states.

C. Information Distance Metric

We describe Crutchfield's *information metric*[1] [4] as a method to compare the difference in information content (in the sense of Shannon information theory), between information sources over time. This measure has the properties of a metric on the space of information sources (such as robot sensors and actuators) and has been shown, in certain specific circumstances, to be better able to find correlations and differences between two random variables than other methods including Hamming distance and frequency distribution distance [5].

To define the information metric, we consider the values of any pair of sensor or motor variables as random variables X and Y, then the *information distance (ID)* is given by

$$d(X,Y) = H(X|Y) + H(Y|X) \tag{1}$$

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

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

and similarly for $H(Y|X)$. 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 sensors or motors for a given period of time. The process used here to estimate a probability distribution $p(x)$ for a discrete sampled variable X, involved dividing the input space into Q bins of equal size, assigning each measured value x of X during the time-series considered to the appropriate bin q and then counting the events f_q in each bin q . The probability $p(x=q)$ is then $\frac{1}{N}f_q$ where N is the number of events.

The probability density function, and therefore the ID, depends on the number of bins Q and on the number of time-steps τ used to estimate the frequency. For our purposes, we are interested in an estimate that reflects the current probabilities so we consider only values of x in a moving window of τ previous time-steps. That is, our information sources are sensors or actuators considered over a length τ time window.

II. SENSOR-MOTOR PHASE-PLOTS

To be able to use the information distance metric to characterise interactions in terms of time-series of sensor-motor values, we reduce the content of the information taking an average of the ID between all sensors and also between all motors.

Consider the discretised values any sensor or motor variable as a time-series of values of a random variable X taking values $X^0, X^1 ...$ in the range[2] $[0.0, 1.0]$. Then for a collection of n random variables $\mathbf{X}^t = (X_1^t, X_2^t, ..., X_n^t)$ at time t with probability density functions estimated for a window of τ time-steps using Q bins, the *average information distance* (AID) at time t across all variables $\mathbf{X}^t$ is given by

$$\begin{aligned} \langle d_{\mathbf{X}}\rangle^t &= \frac{1}{n^2}\sum_{a=1}^{n}\sum_{b=1}^{n} d(X_a^t, X_b^t) \\ &= \frac{1}{2n(n-1)}\sum_{1\le a<b\le n} d(X_a^t, X_b^t) \end{aligned} \tag{3}$$

[1]Also referred to as the *information distance* [5]

[2]Linearly normalised

where the computational simplification follows from the metric symmetry of the information distance. Low values of the AID indicate a small information distance on average between all variables and imply a high degree of correlation between them. One situation where we would expect the AID to be very low would be when the variables were unchanging. The highest value of AID would occur between completely uncorrelated random variables.

Now, grouping all environmental sensoric inputs **S** together and all motor and internal variables **M**, we can calculate the AID for each set and plot the average information distances in two dimensions to get a representation of the relation between sensors and motors. If we do this for successive time-steps for a fixed-size moving window we get a representation of how the sensor-motor relationship is changing with time. We call this plot a *Sensor-Motor AID Phase-plot*.

III. Experimental Method

We chose to run experiments on a real robot to avoid artifacts of simulation and to provide rich sensory-motor data. We use the commercially available SONY AIBO[3] robot - see Fig. 1. Behaviours were written using the Open Source software Tekkotsu [15] 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, with an aim to a) investigate the effect of the window-size τ and bin-size Q parameters, b) study phase-plots of some simple behaviours and c) use phase-plots and AID centres of gravity to identify simple behaviours among a series of behaviours .

Fig. 1. SONY AIBO™ ERS210 used in the experiments.

Table I summarises the variables available to the Aibo from which data was collected. The data was grouped into 36 motor (read-write) variables and 14 sensor (read-only) variables. Additionally, visual images from the head mounted camera were converted into a further 27 individual

[3]AIBO is a registered trademark of SONY Corporation

TABLE I
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

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.[4]

IV. Results and Discussion

Prior to running the experiments outlined above, we looked at whether the AID measure provided anything beyond that of a simple average of all normalised sensor values or the Hamming distance between them. A two minute time-series of data taken while a robot moved around the arena interacting with a pink ball was analysed using AID and a simple average. The resulting plots are shown in Fig. 2 and show that the AID measure captures detail and variation in the data that a simple average cannot.

A. Window-size τ and Bin-size Q

The choice of values for the window-size τ (across which the probability density functions and therefore the information distance was estimated), and for the bin-size Q (which sets the resolution of the probability density functions), could be expected to affect the resulting phase-plots. To investigate the effect, we took data from an experimental run where the robot was "exploring" the arena for 90 seconds (904 timesteps). The "exploring" behaviour consisted of walking forwards until an object (wall) was detected and then turning a random amount before repeating the behaviour.

The data was processed to get the AID averaged over a moving time-window for many different values of τ and Q. A selection of the results are shown in Fig. 3 showing AID phase-plots for varying values of τ for a fixed Q and *visa versa*. Firstly, the results show an overall similarity in structure and indicate that the method is fairly robust with respect to the values of τ and Q. The results also indicate that a increasing τ reduces the detail and variability in the plot as more of the data is averaged (smoothing). In the limit, this would result in a single value for the whole time-series. Increasing Q increases the overall information distance. This would be expected as a finer grained estimation of the probability density would find more differences

[4]Although using a 3×3 grid results in very low visual resolution, these values were chosen to balance the number of motor and sensor variables as it was thought that a large number of visual sensors might bias the data in favour of vision. Note however that experiments were also carried out with 108 vision sensors (36 of each R,G,B) and with 36 sensors (either RED or "Effective-RED" $r - \frac{g+b}{2}$, i.e. the amount of red compensating for the effect of green and blue on perception of red [7] [12, ch. 8]), and the results were broadly similar to those presented for 27 sensors.

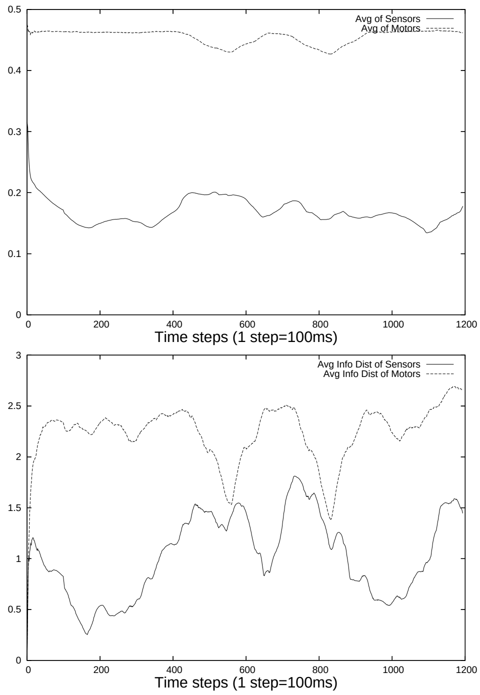

Fig. 2. *Comparison of AID measure and Simple Average.* (Top) Simple moving average of all variables (sensors/motors) over 100 time-step window. *(units of vertical axis are those of the random variables scaled into the range* $[0.0, 1.0]$*)* (Bottom) AID for the same data over a 100 time-step window, *(units of vertical axis are measured in bits)* For AID, previous data in time-window is used to estimate probability distribution. Results show that the AID measure captures detail and variation in the data that a simple average cannot.

in the data (increased maximum entropy). Note also that, depending on the window size τ, and therefore the number of data points used in the estimate, a very large number of bins would be sparsely populated resulting in inaccurate estimations of the probabilities. Thus there is a limit to the amount of detail to be found in the plots.

We used these experiments to choose values of τ and Q with which to conduct the further experiments. A value of $Q = 12$ was chosen with a view to maximising the differentiability in the plots while keeping the computation time to an amount reasonable for on-line computation at 10 frames of data per second.[5] The window-size was kept small $\tau = 20$ to show a large amount of detail only smoothing out short term variations.

B. Characterising Simple Behaviours

Three behaviours were studied; *walking*, *turning* and *observing*, and one control *stationary* (see Table II). Each was repeated a number of times[6], sometimes without change and sometimes with variations in, for example, the direction of walk or location of turn. A phase-plot for one of the walk runs is shown in Fig. 4 along with a plot of the average information distance of sensors and motors against time. This illustrates the utility of the phase-plot method as it makes it instantly clear the changing

[5]Increasing either τ, Q or both results in increased computation time.
[6]except for *stationary* that was only conducted once.

TABLE II

SIMPLE BEHAVIOURS EXECUTED

Behaviour	Description	Runs
walking	Walking from one end of the arena to another	11
turning	Turning on the spot in either direction	7
observing	Robot stationary with activity in environment, e.g. ball or hand waved in front of visual field	5
stationary	Robot remains stationary (with force on motors) in a static environment	1

relations between the two. A feature to note is that due to estimating probabilities over a window, there is a start-up effect while the window becomes populated and also a delay in the phase-plot responding to changes as illustrated by annotations on the figure.

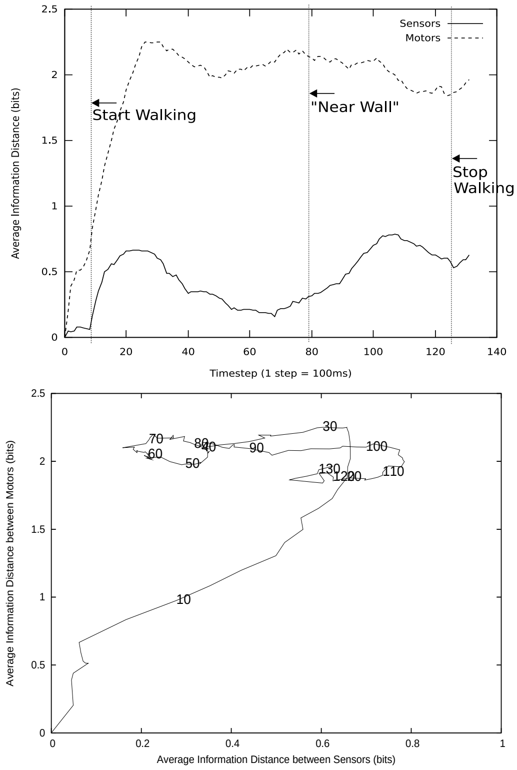

Fig. 4. *Average Information Distance (AID) against time and in sensorimotor phase-plot.* The data is from an AIBO walking from one end of the arena to the other. (Top) Sensors and Motors AID vs time. Annotations mark *Start of walk*, *Near Wall* - which means that wall is in IR sensor range (i.e. 900mm) and *Stop Walking*. This illustrates delay in AID responding to change ($delay \simeq \tau$) (Bottom) Sensor-Motor Phase-plot - plot starts at origin, time steps are marked along path. Data has 130 time-steps, window size τ=20 and bin size Q=12.

One of our goals is to allow the agent or robot to easily characterise and identify behaviours, therefore it needs to be able to simplify a Phase-plot. One way is to take the *Centre of Gravity* (CoG) of the plot by assuming each position on the plot to be a point of unit mass. Additionally we can look at the overall movement of the Phase-plot during a behaviour, a way to do this would be to calculate the overall direction of movement or *Vector* of the phase-plot.

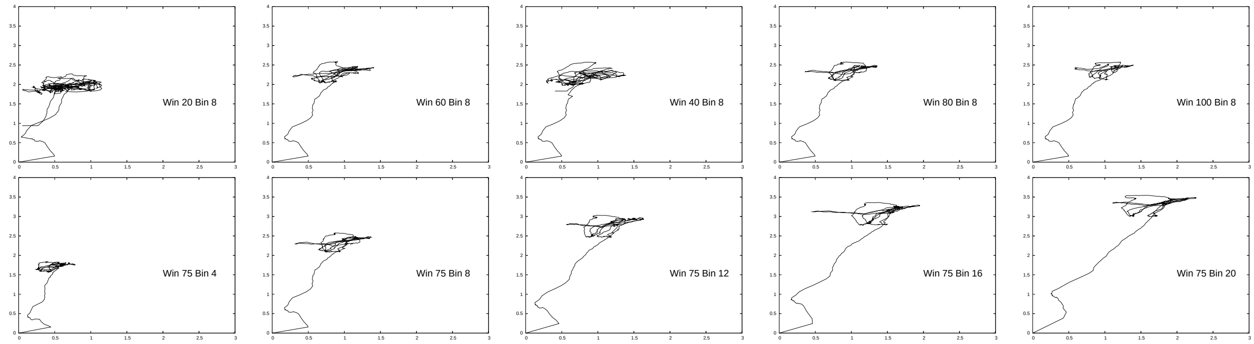

Fig. 3. *Effect on Sensor-Motor AID phase-plots of changing window-size τ and bin-size Q*. Figure shows a selection of results (top) effect of changing window-size $\tau \in \{20, 40, 60, 80, 100\}$ with fixed Q=8 and (bottom) effect of changing bin-size $Q \in \{4, 8, 12, 16, 20\}$ with fixed τ=75. In all cases horizontal and vertical axes are the moving window AID for sensors and motors respectively.

Vector and CoG measures also allow us to easily summarise all the experiments on a single graph. Fig 5 shows the CoG and Vectors of all 24 experimental runs. It is clear from these two measures describing the trajectories that *turning* and *walking* are very different from *stationary* and *observing*. This would be expected due to the activity of the motors. Moreover, the difference between being stationary in a quiescent and a changing environment is shown as a difference in the sensory AID, again as would be expected. This is can also be seen to some extent with *turning* and *walking*, the former being characterised by a far more rapidly changing sensory input from the vision sensors.

Furthermore, it is interesting to note that *turning* and *walking* are further distinguished by how their respective phase-plots change during the behaviour as shown by their Vectors; *walking* has sensory and motor AID reducing while *turning* has (for most of the examples) motor and sensor AID increasing.[7]

C. Identifying Behaviours

The aim of the final experiment was to see if the simple behaviours of section IV-B could be identified within a sequential series of such behaviours. The overall behaviour executed was *exploring* as described in section IV-A consisting of walking and turning behaviours.

The path traversed by the robot, estimated from video, is illustrated in Fig. 6 and annotated with numbered waypoints chosen at points where behaviour changes from walk to turn or *visa versa*. The phase-plot for the complete behaviour was computed and positions of the plot at timesteps corresponding to the waypoints were marked on the plot. The resulting plot is shown in Fig. 7 with the phase-plot itself removed for clarity.

As the transition points mark the end of a period of one kind of behaviour, then the phase-plot positions at that point should reflect the behaviour that has just ended given that the behaviour had been active for enough timesteps prior to that point. The results of Fig. 7 show a clear separation between *walking* and *turning* with the latter showing greater sensor AID than the former. This agrees well with the results of the previous section (IV-B).

[7]Note that even if the CoG and Vector direction were taken into account not all *turning* and *walking* runs could be distinguished.

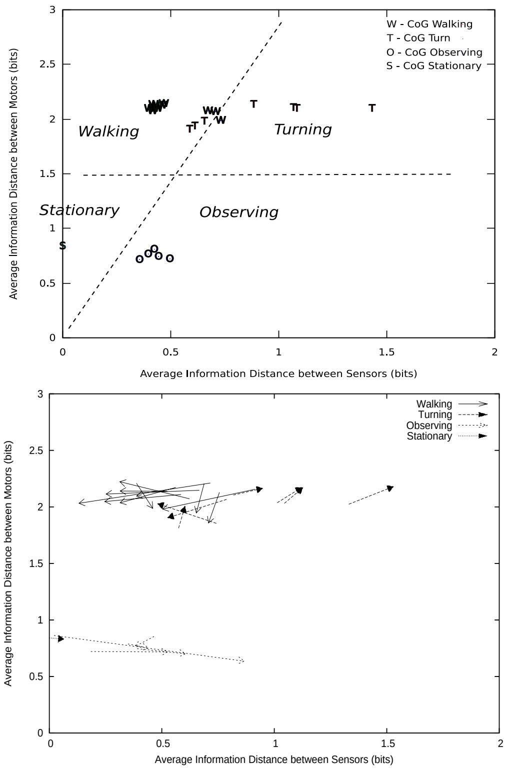

Fig. 5. *Phase-space trajectories. Summary of 24 Simple Behaviour Experiments.* (Top) Centre of Gravity of each phase-plot. Experiments are in 4 categories *walking*, *turning*, *stationary* and *observing* (see Table II). The 4 behaviours appear in 4 quadrants of the geometric space as indicated. (Bottom) Overall direction of movement of each phase-plot as Vectors. These show further distinction between behaviour types. Window size τ=20, bin size Q=12.

V. Conclusions

This paper introduces *Sensory-Motor Average Information Distance (AID) Phase-Plots* as a method for representing and characterising the complex dynamic interactions of an embodied agent. These plots can be thought of as characterising the agent-environment coupling and describe the dynamic interaction history of the agent. The results show that very simple behaviours have different extensions

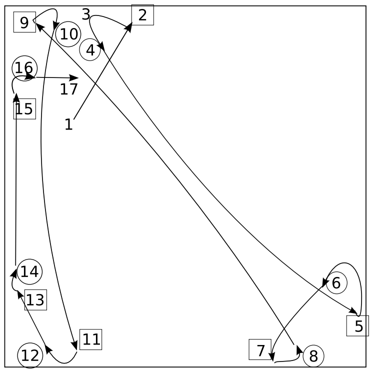

Fig. 6. *Path traversed by AIBO during a particular "explore" experiment.* View is overhead of the 2m x 2m arena. Numbers are waypoints marking changes in behaviour. *end-of-turn* waypoints are circled, *end-of-walk* waypoints are enclosed in a square. 2⟶4 was a multiple turn consisting of a turn from 2⟶3 briefly back towards 2 and finally to 4.

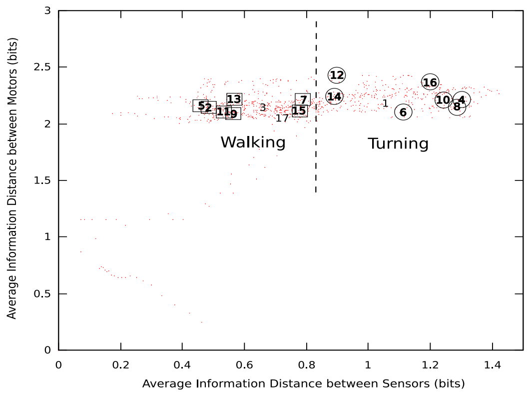

Fig. 7. *Positions on sensor-motor AID phase-plot of waypoints of Fig. 6.* Note: phase-plot itself not shown for clarity. *End-of-turn* positions (4,6,8,10,12,14,16) appear on right side, *end-of-walk* positions (2,5,7,9,11,13,15) appear on left side.

on these plots and these can be identified and distinguished from others. Behaviours were characterised by centre-of-gravity and transition vectors in sensor-motor AID phase-space. This would suggest that an automatic method could be developed for categorising behaviours at a gross scale and identifying behaviours based on past experience.

Robustness of the method was indicted by the qualitatively similar plots produced using different parameter values and repeated runs. Also, reprocessing the data using a higher visual resolution (108 visual sensors rather than 27) produced qualitatively similar results.

VI. Future Directions

This work could be developed in five principal directions: 1) To use the autobiographic history of the agent described by the AID phase-plots to predict and adaptively modify future behaviour. This might be done by learning how to move along trajectories between regions charactersing different types of agent-environment interaction in the higher dimensional metric space of information sources. 2) To use the method to bootstrap exploration and development based on the robot's sensorimotor experience allowing movement toward, between and away from known attractors in sensor-motor AID phase-space. 3) To explore the granularity of the information in the AID phase-plots by varying τ and Q as well as introducing adaptive grouping among sensors and motors (increasing the dimension of the AID Phase-plots). 4) To introduce adaptive weighting of the various inputs to emphasise those important to the agent and its particular embodiment. 5) To take advantage of a developmental contraints to deal with the complexity of the environment and interactions.

Acknowledgement

We would like to thank Lars Olsson and Alexander Klyubin for their help throughout this project.

References

[1] C. L. Nehaniv, K. Dautenhahn, and M. J. Loomes, "Constructive biology and approaches to temporal grounding in post-reactive robotics," in *Sensor Fusion and Decentralized Control in Robotics Systems II (September 19-20, 1999, Boston, Massachusetts), Proceedings of SPIE*, G. T. McKee and P. Schenker, Eds., vol. 3839, 1999, pp. 156–167.

[2] K. Dautenhahn, "Embodied cognition in animals and artifacts," in *Proc. AAAI FS Embodied Cognition and Action*. AAAI Press, 1996, pp. 27–32, technical report FS-96-02.

[3] T. Cover and J. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[4] J. Crutchfield, "Information and its metric," in *Nonlinear Structures in Physical Systems - Pattern Formation, Chaos and Waves*, L. Lam and H. Morris, Eds. New York: Springer-Verlag, 1990, pp. 119–130.

[5] 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*. IEEE Computer Society, 2004, pp. 153–160.

[6] 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*, ser. Lecture Notes in Computer Science 2114, O. Kaynak, E. Alpaydin, E. Öja, and L. Xu, Eds. Berlin: Springer, 2003, pp. 496–503.

[7] 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, 2004, pp. 512–520.

[8] L. Smith and E. Thelen, *A Dynamic Systems Approach to Development*. MIT Press, 1993.

[9] K. Dautenhahn and T. Christaller, "Remembering, rehearsal and empathy - towards a social and embodied cognitive psychology for artifacts," in *Two sciences of the mind. Readings in cognitive science and consciousness*, S. O'Nuallain and P. M. Kevitt, Eds. John Benjamins North America Inc., 1996, pp. 257–282.

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

[11] K. Dautenhahn, B. Ogden, and T. Quick, "From embodied to socially embedded agents - implications for interaction-aware robots." *Cognitive Systems Research*, vol. 3, no. 3, pp. 397–428, 2002.

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

[13] 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. Vienna: Austrian Society for Cybernetic Studies, 1998, pp. 651–656.

[14] 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*. MIT Press, 2002, pp. 345–349.

[15] D. Touretzky and E. Tira-Thompson, "Tekkotsu," *Carnegie Mellon University, Online Open Source, http://www.tekkotsu.org*, 2004.