"And what is a Seasnake?" Modelling the Acquisition of Concept Prototypes in a Developmental Framework

Paul Baxter, Joachim de Greeff, Rachel Wood, Tony Belpaeme
Centre for Robotics and Neural Systems
Plymouth University, U.K.
Email: paul.baxter@plymouth.ac.uk

***Abstract*—The use of concepts is regarded as fundamental to human-level cognition, but there remain a number of open questions as to the functions of, and structures supporting, this competence. In the context of autonomous cognitive systems, the processes by which such concept functionality could be acquired would be particularly useful - indeed, a developmental account would obviate the need for the provision of explicit information or training supervision by the system designer. This paper seeks to explore this issue by applying a set of principles within a developmental framework that support a wide range of cognitive competencies to this problem, thus seeking to embed the development of concepts into a wider framework of cognitive processing. Comparison with a benchmark conceptual modelling system indicates that the proposed approach can account for a number of features, namely concept-based classification, and its extension to prototype-like functionality.**

I. Introduction

The ability to deal with conceptual knowledge is a fundamental requirement for human cognition, and so is likewise a necessary competence for autonomous synthetic agents [11]. In addition to supporting categorisation, one fundamental characteristic is regarded to be the formation of *prototypes*, from which typicality can be assessed [13]. This perspective has largely displaced the notion that concepts can be defined through solely logical descriptions, as it was shown that people readily judge an instance of a concept as being more or less typical of that concept prototype, where this judgement is based on a similarity measurement, rather than a property checklist-like matching procedure.

In this paper, the question of how to account for this aspect of concept utility within a developmental framework is explored. As a model supported by behavioural data on concept competence, a Conceptual Spaces (CS) system is used as a benchmark of human performance. Whilst a good predictive model, the CS system is rather static in structure, and it is unclear how conceptual learning over time can be accounted for. In response to this, a system set within a developmental framework and inspired by a neuropsychological perspective on memory is applied: the Distributed, Associative and Interactive Memory (DAIM) model. An account within a developmental framework such as that provided by DAIM would be desirable not just for the purposes of autonomous operation, but also in terms of providing an account of how such competencies may arise, and how they may interact with other aspects of cognition.

After an introduction to the computational models used to explore the issues of concept prototype acquisition in this paper (section II), the Zoo dataset used for this exploration is described (section IIIA). The experimental procedure and results obtained from the two systems are subsequently presented (section IIIB&C), and then discussed in the context of concept development in autonomous synthetic systems (sections IV and V).

II. Model Architectures

A. Distributed Associative Interactive Memory (DAIM) Model

The DAIM model operates on a set of functional principles derived from the operation of memory within biological systems, embedded within the context of a wider cognitive system [4], [14]. These principles are as follows [14]: (1) memory as being fundamentally associative; (2) memory, rather than being a passive storage device, is an active component in cognition through activation dynamics; (3) memory as having a distributed structure; and finally (4) activation-based priming as subserved by the first three points. The DAIM system has been implemented so as to embody each of these principles in a computational architecture.

Assuming that this system is embedded within a wider agent cognitive system with multiple sensory and motor *modalities*, associations may be seen to form based on the experiences of the agent between *units* of processing in these modalities (i.e. a localist representation scheme) in a hebbian manner, which subsequently form the substrate for activation dynamics. Prior experience as encoded in associative networks, i.e. memory, thus plays an active role in the generation of ongoing behaviour through the mechanism of *priming*, which is the reactivation of modality-specific localist representations on the basis of existing associations. These principles (or variations thereon) have been used to provide candidate mechanisms for a wide range of cognitive phenomena, from visual recognition and analogies [1], [9], to episodic memory, language development and social interaction [4].

1) Mechanisms: The computational implementation of the DAIM model is based on an extension to an Interactive Activation and Competition (IAC) model of face learning

[5], and uses an explicit encoding for associations: i.e. an association is encoded as an object (in the context of Object-Oriented Programming), following the approach taken in [2]. This model differs from standard IAC models (such as [10]), and their learning extensions (e.g. [5]) in four main respects. Firstly, rather than committing to defining a hub of connectivity, DAIM allows all pools of property units (i.e. modalities) to link to other units in any other modality: i.e. point-to-point connectivity. Secondly, weights are updated incrementally at run time, rather than as a batch process only when certain activation stability criteria are met. Thirdly, there is the capacity to create new associative links at run-time, rather than only enabling the adaptation of a structure initialised *a priori*. Finally, in contrast to standard IAC implementations, in the DAIM model used in this study, mutual inhibition between the units of a modality are not implemented. The consequences of this decision are discussed in section IV-3.

There are two main mechanisms present in DAIM: activation spread, and weight update. Activation is taken to be a scalar in the range $[a_{min}, a_{max}]$ where $a_{min} = -0.2$ and $a_{max} = 1.0$. A resting activation level is defined, which is the steady-state activation level of a unit in the absence of stimulation: $a_{rest} = -0.1$. Similarly, weights (of associative links) are taken to be scalars in the range $[w_{min}, w_{max}]$ where $w_{min} = -1.0$ and $w_{max} = 1.0$; the initial weight of associative links upon creation is defined as $w_{init} = 0.2$. A new associative link is created between two *units* in different *modalities*, *iff* they are the most active units in their respective modalities, and such a link does not already exist. The net activation input to each unit is derived as follows, where ext_i is the activation derived from an 'external' source (input to a modality); $\xi_g = 0.6$ is a parameter controlling the proportion of externally derived activation used; int_i is the activation derived from within DAIM (activation from other modality units); and $\zeta_g = 0.3$ controls the influence of int_i:

$$net_i = (\xi_g \times ext_i) + (\zeta_g \times int_i) \tag{1}$$

This is based on the derivation of the activation spread from within the DAIM system (int_i), which is calculated as follows (on every time-step), where w_{ij} is the weight of an associative link linking unit i to unit j in another modality, and out_j is the activation of the linked unit j:

$$int_i = \sum_j w_{ij} \times out_j \tag{2}$$

The calculated net_i effectively encodes the net effect of all inputs to each individual unit, on each time-step, with the result being 'excitatory' if $net_i > 0$, and 'inhibitory' if $net_i < 0$. On this basis, the change in activation level of each unit may be updated (Δa_i), where $\delta_g = 0.08$ determines the proportional decay in activation level:

$$\begin{aligned} If(net_i > 0) : \Delta a_i &= net_i\,(a_{max} - a_i) - \delta_g\,(a_i - a_{rest}) \\ else : \Delta a_i &= net_i\,(a_i - a_{min}) - \delta_g\,(a_i - a_{rest}) \end{aligned} \tag{3}$$

The bounded weight update mechanism is based on that derived by Burton et al [5], basing update magnitude on activation magnitudes of the linked units, where $\lambda_g = 0.01$ is the learning rate:

$$\begin{aligned} If(a_i a_j > 0.0) : \quad & \Delta w_{ij} = \lambda_g a_i a_j\,(1 - w_{ij}) \\ else : \quad & \Delta w_{ij} = \lambda_g a_i a_j\,(1 + w_{ij}) \end{aligned} \tag{4}$$

This weight update mechanism operates under the additional condition that $if\,((a_i < 0) \cap (a_j < 0)) \;:\; \Delta w_{ij} = 0$, to counteract the effect of the negative activation resting value (a_{rest}) gradually increasing all weights.

2) What makes DAIM a developmental architecture?: As described above, the DAIM model is essentially an active associative substrate upon which activation dynamics operate. This substrate is subject to adaptation (e.g. associative link weight update) over the course of interaction of the system as a whole with its environment: as such, adaptation and activation dynamics are inherently inter-dependent. It is reasonable to ask then whether DAIM can be classed as being a developmental system, rather than 'merely' a learning system. We contend that it is subject to a developmental process, since a fundamental feature of operation is the creation of the associative substrate itself based on experience (i.e. through the creation of new associative links), in addition to its subsequent adaptation, which may be regarded as learning [3].

B. The Conceptual Spaces (CS) Model

A Conceptual Space (CS) consists of a geometrical representation in vector space along various quality dimensions [8]. This perspective on concept representation is consistent with accounts of human behaviour (e.g. [13]); a CS model is therefore used as a benchmark system against which the performance of DAIM can be assessed. A CS is a collection of domains (like colour, shape, or tone), where a domain is postulated as a collection of inseparable sensory-based quality dimensions with a metric (similar to a DAIM modality). For instance, to express a point in the colour domain using an RGB encoding, the different quality dimensions *red*, *green*, and *blue* are all necessary to express a certain colour and are therefore inseparable. In its simplest form, a concept can be represented as a point in the conceptual space, where the coordinates of the point determine the features of the concept.

Crucial to modelling concepts in a CS is the ability to take a distance measurement. For each of the dimensions involved, a suitable metric to calculate distance between coordinates on this dimension must be defined. For the case of numerous dimensions the Euclidean distance is typically the most appropriate. For example, for colour, a normalised RGB space can be defined, such that the Euclidean distance between any two points in this space can be calculated. The metric can be augmented with a weight (w) to allow certain dimensions to be more prominently expressed than others.

Within a CS the learning of prototypes can be modelled by exposing the model to instances with associated labels. For example, various shades of red could be presented, each with the label 'Red': the prototype in this case could correspond to the point mean of all instances. After this learning phase, the model is able to classify new examples as belonging to some known class, and specify how typical the example is, based on a distance measurement between it and the various possible classes.

1) Mechanisms: The notion of prototypes comes naturally to conceptual space modelling, as the distance metric functions naturally as a notion of typicality. Distance d_{xy} between a prototype x and and an example y takes the general form:

$$d_{xy} = \left(\sum_{i=1}^{N} w_i |x_i - y_i|^r \right)^{\frac{1}{r}} \tag{5}$$

where r denotes the type of metric with $r = 2$ for the Euclidean distance and $w = 1.0$ the weight of the dimension. To do justice to psychological evidence of how people tend to rate concepts [12], we can convert the distance into a similarity measurement:

$$s_{ij} = e^{-cd_{ij}} \tag{6}$$

where similarity s between i and j is computed as an exponentially decaying function of distance, where $c = 1.0$ is a sensitivity parameter.

III. Modelling Conceptual Prototypes

The fundamental task to which the two systems (DAIM and CS) are applied is the learning of animal classes from presented instances of animals, and applying this learned information to determine the class of a novel set of animal instances. While classification accuracy is a useful metric of performance, of greater interest to this study is the ability of DAIM to model prototype-based similarity judgements, with respect to the CS benchmark.

A. The Zoo dataset

The data set that is the subject of this study is the Zoo animal data set from the UCI Machine Learning Repository [7], which is a database of 100 named animals with 17 properties. The majority of these properties are binary, such as 'has hair', 'is aquatic', or 'lays eggs'. The other two properties are categorical (animal class, which takes one of seven values: Mammal, Bird, Reptile, Fish, Amphibian, Insect, or Invertebrate), and scalar (number of legs; 0, 2, 4, 5, 6, 8). The distribution of animal instances over the seven classes is shown in table I, as is the distribution of the training vs. probe instances. It can be seen that the distribution of instances over the classes is very uneven, leading to potential difficulties: this issue is explored further in section IV. However, one advantage of using a dataset such as this is that it enables an intuitive assessment of the system behaviours, in addition to the quantitative results that may be obtained.

To provide an illustration of the relationship between the animal classes and the probe instances used in this study (i.e. the structure of the dataset), a Principal Components Analysis (PCA) was conducted on the dataset (figure 1). The two primary components were identified, which together account for around 59% of the variability: as such, the relative

TABLE I
Class distribution in the Zoo animal dataset, in the training data, and in the probe data.

Class	Dataset	Training	Probe
Mammal	41	39	2
Bird	20	16	4
Reptile	5	4	1
Fish	13	12	1
Amphibian	3	3	0
Insect	8	8	0
Invertebrate	10	10	0
Total	100	92	8

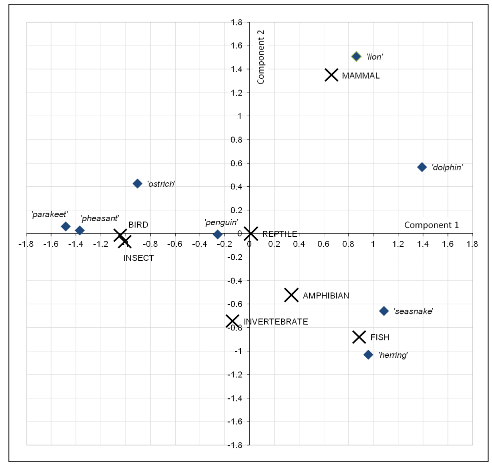

Fig. 1. The first two components of a Principal Component Analysis dimension reduction (from sixteen to two): the seven animal type categories are shown (crosses), as are the eight instances used for testing (diamonds).

locations of the classes and instances should be taken as illustrative, and not as a representation of the optimal similarity ratings to be achieved by the CS and DAIM models.

B. Experimental Procedure

In this study, both models were trained on a subset of the zoo dataset. Eight animal instances were reserved for probe trials, in which the respective models were tested on the ability to identify the animal class, given the 16 instance properties. The remaining 92 instances were used as training data. The distribution of animal classes used in the training and probe sets are shown in table I. The choice of instances to use as probes was based on the utility in illustrating the functionalities of classification and similarity to prototype judgements that are the subject of this study. There are two stages to the experiment: in the first stage, the two systems are trained, and in the second stage, the probes are used to examine the behaviour of the system. In this second stage, further learning and adaptation is disabled to enable a proper assessment of performance.

The training of the CS model involves presenting the properties of all training instances to the system sequentially: there is no order effect. A conceptual space is set up based on the animal 'class' property such that each presented animal instance is projected to a point in this space so that a distance (and hence typicality) measurement can be determined between the instances and prototypes. It should be noted that the property 'number of legs' is normalised prior to training, to maintain equivalent ranges across all properties. Learning in the CS model is thus supervised, as the subject of the classification (animal class) is provided with the instance to be learned. These explicit prototype representations enable the classification of novel stimuli: during the probe phase, the probe instance properties are presented to the CS system, projected to the animal class conceptual space, and similarity measures to the prototypes present derived (see equation 6).

For the DAIM model in this study, each of the animal instance properties constitutes a *modality*, with the features of each property (i.e. true/false, name, number of legs, animal class) constituting the modality *units* (please refer to section II-A for the relation to the theory). The training procedure takes into account the temporal and iterative nature of the learning mechanisms (specifically weight update), in an unsupervised manner. As such, it takes the form of a sequence of instances to learn: the properties of an instance are presented to the system for 5 time-steps (during which time associative links are created and updated), followed by a period of 20 time-steps in which there is no input to the system (to ensure all activation decays before the next instance presentation). This presentation occurs as follows: an activation value of 1.0 is applied to all of the modality units corresponding to the properties of the animal instance; all other units receive no activation (i.e. 0.0). For the probe trials, the properties for each of the probe instances are presented to the system for 10 time-steps, followed by 60 time-steps of no input. The length of the probe stimulation is sufficient for the activations on the animal class properties to reach a steady state (see figure 2). These steady state activation levels are then normalised: the resultant values are used as the basis of the results reported below.

Given that the learning mechanism in DAIM is incremental (as a result of the associative link creation and weight update mechanisms), the order of data presentation during learning influences the learned information, and hence the behaviour of the system. To assess the effects of varying presentation orders on the ability of the system to correctly classify novel instances, the order of the training set was randomised: 10 training sets were derived in this way. Each resultant version of the DAIM model may thus be regarded as having a different experience in the environment, albeit with common statistical relationships for all versions. The probe phase was the same for each instance of the DAIM model and the weight update mechanism was disabled to ensure that comparisons could be made across the probe animal instances on a common basis.

C. Results

The first assessment that can be made is the classification accuracy of the CS and DAIM models given novel animal instances. A comparison of four animal instance probes provides an illustration of various characteristics of this classification process: for each case, a classification may be determined by a winner-takes-all strategy (i.e. the class with the highest

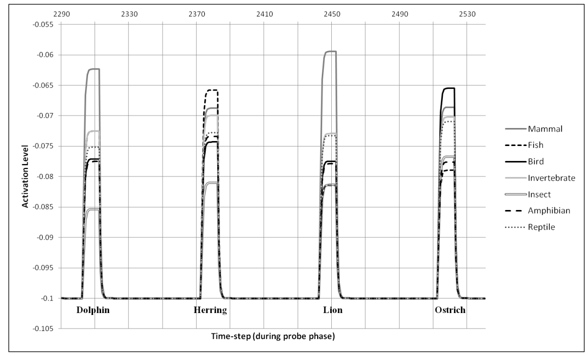

Fig. 2. Example activation profiles for animal classes during the probe phase of one DAIM run, showing four probes (the properties of Dolphin, Herring, Lion and Ostrich). 'Resting' activation level is -0.1; rising activity is due to activation spread on the substrate of created associations between properties. Steady-state activation levels (constant for at least 3 time-steps) normalised for subsequent analysis.

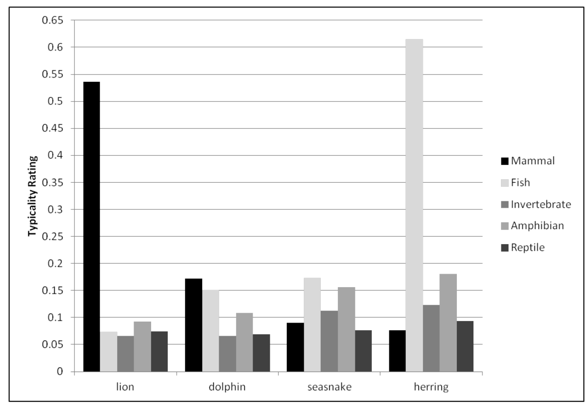

Fig. 3. Conceptual Space Model classification results for four animals. Five of the seven available animal type categories are shown, for the purpose of clarity. All animals are classified correctly, except 'Seasnake' (Reptile): this is classed as a fish, although the rating for amphibian is similarly high.

typicality rating or activation level). For the CS model (figure 3) the cases of 'lion' and 'herring' are clear, in that the correct class has a far greater typicality rating than the other classes. Similarly, 'dolphin' is correctly classified as a mammal, although the typicality rating for fish is comparable. Finally, 'seasnake' is incorrectly classified, with fish and amphibian being identified to a similar extent instead.

For the DAIM model, the results show the aggregated results for the 10 ordered training datasets (figure 4). Firstly, it can be noted that there is a high level of consistency of resuts across the 10 training set orders (as evidenced by the small 95% CI's), indicating that while there is an order effect, it does not disrupt classification accuracy. In accordance with the CS results, 'lion', 'herring' and 'dolphin' are correctly classified. 'Seasnake' is not classified correctly, with similar activation values for mammal, fish and invertebrate. This is quantitatively the most divergent result in comparison with those derived from the CS model, although qualitatively, the fact that 'seasnake' is misidentified, with a number of possible (indeed intuitively plausible) alternatives present in both cases.

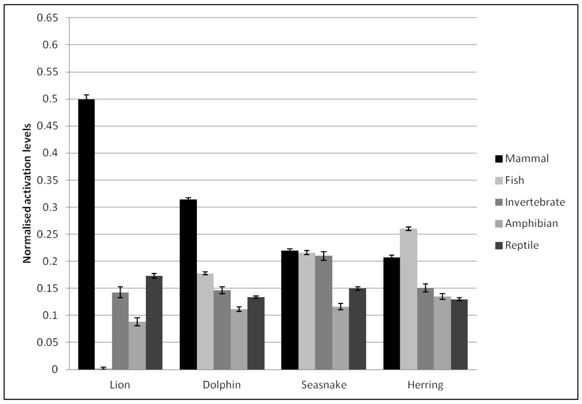

Fig. 4. DAIM Model classification results for four animals. The same five categories are as in figure 3. The results show the mean of ten ordered datasets (see main text for details): the error bars are 95% confidence intervals. As for the CS model, all animals are classified correctly, except 'Seasnake': in this case, there is ambiguity between mammal, fish, and invertebrate.

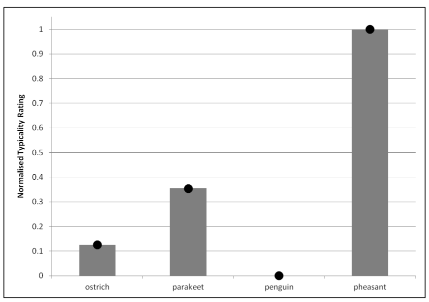

Fig. 5. Conceptual Space Model results comparing the relative typicality of four birds. All four were correctly classified as birds; this shows that 'parakeet' and 'pheasant' are more typical birds than 'penguin' and 'ostrich'.

Additionally, it can be noted that the relative magnitude of the activation of mammal for 'seasnake' and 'herring' is higher than in the CS model.

The second assessment that can be made is the degree to which the respective systems are able to determine the typicality of novel input animal instances to a prototype. Assuming that classification is correctly performed, the question here is whether the two models can produce a measure of how close a presented animal instance is to a learned concept prototype. This measure of typicality is explicitly implemented in the CS model (cf. eq 5 and 6), but not for DAIM, where all information is maintained in a distributed state: any such assessment must be made on the basis of relative activation levels (e.g. figure 2). It is thus necessary to demonstrate that this is not only possible, but that DAIM generates results consistent with the empirically supported CS benchmark.

To illustrate this, a comparison of four birds may be conducted. Each bird ('ostrich', 'parakeet', 'penguin' and 'pheasant') is presented as a probe in the second phase of the simulations. For both the CS and DAIM models, all of these are classified correctly. For the CS model, typicality ratings are calculated, and normalised across the four values (figure 5) to enable a comparison with the DAIM results. It can be seen that in this case, 'pheasant' is identified as the most typical bird,

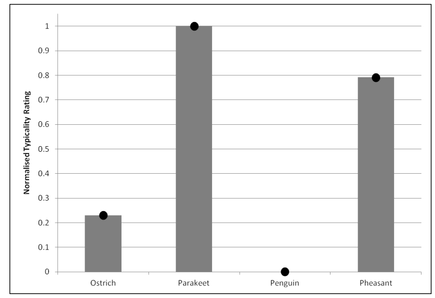

Fig. 6. DAIM Model results for the relative typicality of four birds. Results show the mean of 10 ordered datasets (see main text for details). All four were correctly classified as birds; as with the CS model result, 'parakeet' and 'pheasant' are classed as more typical birds than 'penguin' and 'ostrich'.

and that there is a clear difference between the pair 'parakeet'-'pheasant' (being fairly typical), and 'ostrich'-'penguin' (being relatively atypical). This same qualitative pattern and division of the two pairs can be seen in the DAIM results (figure 6), even though 'parakeet' is identified as the most typical bird instance. The DAIM results are derived from the activation levels as described above, with a similar normalisation process used to enable comparison with the CS model results.

IV. Perspectives

The described results show that there is similarity between the behaviours of the CS and DAIM systems for classification of the novel animal instances: indeed, for the 8 probes, classification accuracy (based on a winner-takes-all selection procedure) is the same for both (7/8, 87.5%). Additionally, there is a qualitative similarity of the DAIM system performance in typicality ratings to those derived by the CS model, despite the fact that the means of calculating and assessing typicality and similarity differ fundamentally. In a comparison of the bird typicality ratings from the two models, while the actual order of ratings differs (with 'parakeet' being most typical of a bird for DAIM, and 'pheasant' for CS), the indication that these are more typical than either 'penguin' or 'ostrich' is clear (and thus matching intuition). There are a number of potential sources for the differences, including the exponential-based calculation of similarity for CS, and the order-dependent effects for DAIM. However, that such strong similarities exist for both classification and prototype-based similarity assessments provides support for the notion that the mechanisms that DAIM makes use of can account for these two fundamental features of concept functionality. In addition to these observations, a number of other issues merit further consideration:

1) Uneven distribution of data across categories - the case of the Seasnake: The case of the mis-classified 'seasnake' raises a number of questions. That both CS and DAIM fail to classify it correctly may be an indication that there is some inherent ambiguity resulting from the dataset itself. Indeed, if the distribution of animal classes is considered (table I), there are very few examples of reptiles in comparison to mammals for example. In this case, it is instructive to consider the manner in which the mis-classifications were made: for both CS and DAIM, they reflect to some degree the overlapping properties of the seasnake with animals from other classes, notably fish (cf. figure 1). The inclusion of mammal into this consideration for the DAIM results may reflect the influence of dominance of mammals in the dataset on the incremental nature of the weight update mechanism. Nevertheless, there is a clear indication in this behaviour that even with mistaken classifications, there is the possibility for the outcome to be of utility in further processing, by, for example, providing a set of hypotheses that can be used as the basis for further disambiguation actions.

2) Robustness to varying experience: The results of the classification task for the DAIM model (figure 4) indicate that the classification performance is robust to presentation order within the learning phase. In the context of developmental systems, the importance of a trajectory based on experience is typically emphasised; this result supports the notion that even with unique experiences, there is the capacity for the statistical regularities of the environments shared by agents to lead to robust conceptual categories for those agents, in support of coordination through inter-agent interaction. Indeed, such an extension to the present work is desirable (e.g. [6]). It should be noted though that in this study, the full data set was used (minus the probe instances, but these remained constant), and only the order of the training set was altered. As such, the statistics of the training set as a whole remained unchanged. An n-way cross validation could therefore be applied to resolve this issue: this is not reported in the present study due to the focus on the prototype-based functionalities of DAIM, which was best served by the analysis of specific cases.

3) The role of mutual inhibition: A notable difference between the DAIM instantiation used in this study and standard IAC models is the lack of a mutual inhibition mechanism. In addition to preventing runaway activation levels (due to mutual excitation), mutual inhibition can also increase the efficacy of classification in that in that slight advantages tend to lead to clear discrimination. With the version of DAIM used in this study however, the utility of no such mutual inhibition is emphasised, enabling similarity judgements to be made. A trade-off between excitatory and inhibitory mechanisms is thus apparent: while such a trade-off is well known in natural neural systems, in this study it is seen that allowing activation to persist in multiple units without direct competition may be beneficial for processing in a wider cognitive context, particularly where a clear classification fails in the first instance. This aspect of integrated cognitive processing presents itself as a promising avenue for further investigation.

4) The developmental context of DAIM: The DAIM model has been described above as being a system that has a developmental trajectory in terms of increasing competencies with increasing interactions with its environment. As part of a wider cognitive system, the principles upon which it operates enable DAIM to bias, influence and/or modulate ongoing system behaviours using the mechanism of priming. By demonstrating that DAIM can account for (at least some central aspects of) conceptual functionality, there is a potential reduction in reliance on explicit symbolic constructs. This provides a fundamental mechanistic integration of conceptual competencies and their development within wider cognitive processing (e.g. [11]), within a developmental framework.

V. Conclusion

The results obtained indicate that the DAIM model can reproduce two central functional properties of concepts: categorisation and prototype-based similarity assessments. It does so using a distributed representation scheme operating on the principles of association and activation dynamics, which are consistent with, and are being used to account for, a wide range of other cognitive competencies. This enables concept development and functionality to be explicitly viewed in the context of a wider cognitive architecture, and, given the DAIM operating principles, within a developmental framework. There are many questions that remain open in this regard, but this study has provided evidence in support of the integration of conceptual competencies within the developmental memory-centred perspective on cognition.

Acknowledgment

This work is funded by the EU FP7 ALIZ-E project (grant 248116).

References

[1] M. Bar, The proactive brain: using analogies and associations to generate predictions. Trends in cognitive sciences, vol. 11, no. 7, pp. 280-289, 2007.

[2] P. Baxter, *Foundations of a constructivist memory-based approach to cognitive robotics*, PhD. Thesis, University of Reading, U.K., 2010.

[3] P. Baxter and W. Browne, *Memory as the substrate of cognition: a developmental cognitive robotics perspective*, 10th International Conference on Epigenetic Robotics, pp. 19-26, 2010.

[4] P. Baxter, R.Wood, A. Morse, and T. Belpaeme, *Memory-Centred Architectures: Perspectives on Human-level Cognitive Competencies*, AAAI Fall 2011 Symposium on Cognitive Systems, pp. 26-33, 2011.

[5] A. M. Burton, *Learning new faces in an interactive activation and competition model*, Visual Cognition, vol. 1, no. 2, pp. 313-348, 1994.

[6] J. De Greeff, F. Delaunay, and T. Belpaeme, *Human-Robot Interaction in Concept Acquisition: a computational model*, International Conference on Development and Learning, pp. 1–6, 2009.

[7] A. Frank and A. Asuncion, *UCI Machine Learning Repository*, http://archive.ics.uci.edu/ml (*accessed 28/06/2012*), Irvine, CA: University of California, School of Information and Computer Science, 2010.

[8] P. Gardenfors, *Conceptual Spaces: The Geometry of Thought*, Cambridge, MA: the MIT Press, 2000.

[9] R. Leech, D. Mareschal, and R. P. Cooper, *Analogy as relational priming: a developmental and computational perspective on the origins of a complex cognitive skill*, Behavioral and Brain Sciences, vol. 31, no. 4, pp. 357-78; discussion 378-414, 2008.

[10] J. L. McClelland and D. E. Rumelhart, *An Interactive Activation Model of Context Effects in Letter Perception: Part 1, an account of basic findings*, Psychological Review, vol. 88, no. 5, pp. 375–407, 1981.

[11] A. Morse, J. De Greeff, T. Belpaeme, and A. Cangelosi, *Epigenetic Robotics Architecture (ERA)*, IEEE Transactions on Autonomous Mental Development, vol. 2, no. 4, pp. 325–339, 2010.

[12] R. M. Nosofsky, *Attention, similarity, and the identification/categorization relationship*, Journal of Experimental Psychology-General, vol. 115, no. 1, pp. 39-57, 1986.

[13] E. H. Rosch, *Natural categories*, Cognitive Psychology, vol. 4, no. 3, pp. 328-350, 1973.

[14] R. Wood, P. Baxter, and T. Belpaeme, *A review of long-term memory in natural and synthetic systems*, Adaptive Behavior, vol. 20, no. 2, pp. 81-103, 2012.