Competition affects word learning in a developmental robotic system*

K. E. Twomey
Centre for Research in Human Development and Learning, Department of Psychology
Lancaster University, Lancaster, UK
E-mail: k.twomey@lancaster.ac.uk
http://www.psych.lancs.ac.uk/research/crhdl/

A. F. Morse & A. Cangelosi
Centre for Robotics and Neural Systems
University of Plymouth, Plymouth, UK
E-mail: anthony.morse@plymouth.ac.uk, angelo.cangelosi@plymouth.ac.uk

J. S. Horst
School of Psychology
University of Sussex, Brighton, UK
E-mail: jessica@sussex.ac.uk

It is well-established that toddlers can correctly select a novel referent from an ambiguous array in response to a novel label, or *fast-map*. There is also a growing consensus that robust word learning requires repeated label-object encounters. However, the effect of the context in which a novel object is encountered is less clear. Horst, Scott & Pollard (2010) demonstrated that the more competitor objects present when toddlers chose a novel object from an array in response to a novel label, the harder it was for children to retain that label. Here, we present a developmental robotic replication of this study using a variant of the Epigenetic Robotic Architecture (Morse, de Greeff, Belpeame, & Cangelosi, 2010) implemented in the iCub humanoid robot (Metta *et al.*, 2010). We discuss children's word learning in terms of the real-time processes governing the iCub's performance in the word learning task.

Keywords: Developmental Robotics; Neural Networks; Word Learning; Fast-Mapping.

1. Computational insights into young children's word learning

Toddlers perform impressively when confronted with the seemingly difficult task of choosing the referent of a new word in an ambiguous environment

* This work is supported by EU FP7 project Poeticon++ (no. 288382)

(Quine, 1960). A rich empirical literature shows that children reliably map a novel label to the single novel referent in an array of otherwise known objects, a phenomenon known as *fast-mapping* (Carey & Bartlett, 1978). There is also a growing consensus that over and above this quick initial mapping, robust word learning requires repeated encounters with the label and the target object (Smith & Yu, 2008). Although we have a sound understanding of *what* children do in response to novel objects and labels, *how* they do it is less well understood.

Recently researchers have begun to address this issue, using computational methods to explore the perceptual and cognitive processes underlying empirically observed behaviour (e.g., McMurray, Horst, & Samuelson, 2012; Twomey, Horst, & Morse, 2013). Computational models of word learning simulate how children behave (e.g., pointing to the flamingo and not the dog or the fish) based on what they see and hear (e.g., one novel object, two known competitor objects and the novel word "flamingo"). Just like children, models have internal representations which change with learning. In sharp contrast to children, however, we can inspect these representations as they develop over time. Thus, by examining the mechanisms driving representational change in the model, we can build an explicit account of the mechanisms which drive cognitive development in the child (Westermann & Mareschal, 2012)

2. Target Empirical Data: The Effect of Referent Competition on Word Learning

Historically, the word learning literature has explored the biases which children appear to use when mapping novel labels to novel referents (e.g., Markman, 1990). Horst, Scott and Pollard (2010; henceforth HSP) extended these findings by exploring the relationship between the context in which words are fast-mapped and children's ability to retain those words. Specifically, they presented 36 30-month-old children with referent selection trials consisting of an array of 3D age-appropriate toys, one of which was novel and the rest of which were known to children. Critically, HSP varied between conditions the number of competitor objects seen during referent selection: trials consisted of a novel object and two, three, or four known competitors. In the two competitor condition, for example, a trial might consist of a plastic cone with multicoloured strings attached to it, a small plastic horse, and a small plastic block. In the four competitor condition, a trial might consist of the cone, the horse, the block, a spoon and a toy car. All other aspect of the design were held constant.

During the referent selection phase, children were presented with four sets of objects, across eight referent selection trials. Each novel object was presented twice and served as a target once. On each trial, children were allowed to look at the objects for three seconds before being asked to select either a known or the novel object (e.g., *known trial*: "Can you show me the horse?"; *novel trial*: "Can you show me the *fode*?"; five times per target; thus, two trials per set.) Children therefore had an equal amount of experience with each novel target during referent selection.

After referent selection children were presented with four test trials. Each of the four novel targets appeared on each test trial, and children were asked for each object in turn. If children had retained the novel label-object mappings formed during the referent selection phase, then they should pick the target object at levels greater than expected by chance (25%).

In line with existing studies, children were very good at referent selection, performing significantly above chance on both known and novel trials regardless of the number of competitors present. However, only children in the two competitor condition retained novel labels at levels greater than expected by chance, and did so significantly more reliably than children in the three and four competitor conditions (see light blue bars, Fig. 2). An analysis of reaction times during referent selection revealed that children in the two competitor condition selected novel objects marginally faster than children in the four competitor condition. Importantly, when reaction times were weighted by the number of competitors present, they did not differ between conditions. Thus, reaction times were related to the number of objects they saw, and children who retained label-object mappings after seeing the fewest objects were significantly faster to make those mappings than children who saw the most objects. The authors reasoned that the disambiguation task in the two competitor condition was less onerous in terms of processing capacity than in the other conditions. Put differently, these children only had to identify two known objects before mapping the novel label to the novel referent – a quicker and simpler undertaking than identifying three or four known objects.

The authors concluded that fast-mapping – and eventual word learning – involves paying attention to competitor objects in order to establish what the referent *is not*, as well as paying attention to the novel object, to establish what it is (see also Zosh, Brinster & Halberda, 2013 for evidence of the importance of competitor objects to the disambiguation task). On this account, fast-mapping is influenced not only by novelty, but also by knowing the names of the competitor objects; subsquent word learning is therefore the product of learning which associations are correct (e.g., novel object-*fode*), but also of learning which are wrong (e.g., cow-*fode*; see also McMurray *et al*., 2012) . The implication, therefore, is that word learning emerges from the interaction of multiple timescales of development: long-term experience (children's known vocabularies), medium-term cross-situational learning (multiples exposure to target/label mappings) and in-the-moment fast-mapping (McMurray *et al*., 2012).

3. The iCub and the Epigenetic Robotics Architecture

In summary, HSP's children did not learn novel label-object mappings when exposed to the target object in the context of more than two competitor objects during the referent selection phase. Here, we explore the processes underlying these results using a version of the Epigenetic Robotics Architecture (ERA; Morse, de Greeff *et al.*, 2010) implemented in the iCub humanoid robot (Metta *et al.*, 2010).

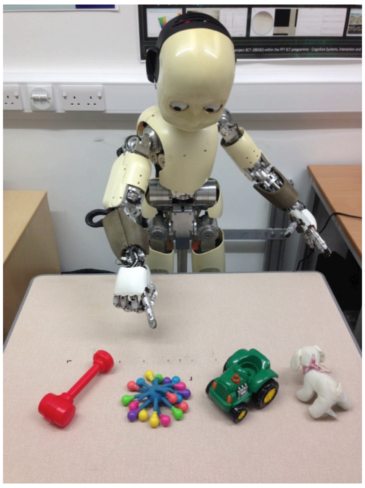

Fig. 1. iCub during referent selection

iCub's design reflects the physical proportions of a three-year-old child, with 53 degrees of freedom, and sensors, which encode a range of naturalistic perceptual input, approximating young children's perceptual environments (Fig 1). Thus, like children, iCub integrates visual, auditory, tactile and proprioceptive information to generate behaviour, for example auditory and visual information in a word learning task (although which modalities contribute to a given simulation are decided *a priori* by the modeller). Thus far, iCub has captured a diverse range of language acquisition data (e.g., affordance-based verb learning, Marocco, Cangelosi, Fischer, & Belpaeme, 2010; spatiallygrounded noun learning, Morse, Belpaeme, Cangelosi, & Smith, 2010; the effect of category structure on word learning, Twomey *et al*., 2013; affordance-based adjective learning, Yürüten *et al*., 2012).

The architecture employed in the current project was based on the ERA previously used in a replication of a category label learning experiment (Twomey *et al*., 2013). The ERA consists of a network of Self-Organising Maps (SOMs; Kohonen, 1998): connectionist networks that reorganise their internal structure based on a winner-takes-all response to input stimuli. At the end of learning, the SOM reflects the structure of the input in its own topological structure; that is, neurons that are close together in the network fire in response to perceptually similar stimuli (e.g., the colours red and pink). SOMs naturally lend themselves to categorisation of complex naturalistic stimuli such as the auditory and visual inputs generated by iCub's sensors.

The model comprises two visual SOMs which receive processed visual information from iCub's cameras. One map receives an HSV spectrogram of each object in view and so represents colour, and the other receives shape information about each object (circleness, squareness, convexity, elongation; Montesano, Lopes, Bernardino & Santos-Victor, 2008). Speech recognition, via the commercial software Dragon Dictate^{TM}, was used to provide speech-to-text input for the words, where each word dynamically activates a single unit in a label field. The visual SOMs are bidirectionally coupled to the field of label inputs via Hebbian-like links to form a dynamic spreading activation network. Objects which are primed cause iCub to look at them or to reach and point to them.

For an object in a particular region, colour information is extracted by determining the location in HSV colour space (Alvy Ray, 1978) of each pixel in that region. Ignoring the white background of the table, pixels with a saturation value greater than a threshold of 0.2 are allocated to one of 36 bins each representing 10 degrees of the 360 degree HSV hue continuum, which generates a histogram-like colour profile for each object. Each object profile is unique and based on the entire range of the colour SOM. Thus, the model takes into account differences between uniformly and multicoloured objects.

Input SOMs were initialised with random connection weights. The 20 objects used in this study were simultaneously placed in view and the SOMs trained using standard equations 1 (SOM activity) & 2 (SOM learning rule):

$$BMU = Max_i\left(1-\sqrt{\sum\left(a_j - w_{ij}\right)}\right) \tag{1}$$

$$\Delta w_{ij} = \alpha \exp\left(-\frac{dist^2}{2size}\right)\left(a_j w_{ij}\right) \tag{2}$$

The Best Matching Unit (*BMU*) is the unit with weights (*w*) most closely matching the current input vector (*a*). The weights of each node within the neighbourhood of the BMU are then modified to move closer to the current input vector, scaled according to the distance of this node from the BMU in the space of the SOM (i.e., not in terms of the input space).

The neighbourhood size and learning rate both monotonically decrease until the neighbourhood size is 1, at which point both the neighbourhood size and the learning rate of the SOM are fixed to allow learning to continue at a low rate.

The two SOMs and the label field are fully connected via Hebbian links subject to spreading activation and learning as in equations 3 (IAC Spreading Activation) & 4 (Hebb-like learning rule):

$$net_i = \sum w_{ij} a_j + \beta BMU_i \tag{3}$$

If $net_i > 0$ $\quad \Delta a_i = \left(\max - a_i\right) net_i - decay\left(a_i - rest\right)$

Else $\quad \Delta a_i = \left(a_i - \min\right) net_i - decay\left(a_i - rest\right)$

The net input to each unit in the whole network is the sum of spreading activation plus external activation if this node happens to be the BMU of a SOM or a currently active word. The experiment reported here used the parameter values: External Input Bias = 0.5; Max = 1; Min = -0.2; Decay = 0.5; Rest = -0.01.

If a_i OR $a_j > 0$: (4)

if $a_i a_j > 0$ $\quad \Delta w_{ij} = \lambda a_i a_j \left(1 - w_{ij}\right)$

else $\quad \Delta w_{ij} = \lambda a_i a_j \left(1 + w_{ij}\right)$

The Hebb-like learning rule increases the strength of a connection if both units connected by this weight are positively active or reduces the strength of the weight if only one is positively active, and scales that change according to the product of their activity and how close the weights are to 1 or -1 respectively for positive and negative weight changes. Finally each field or pool is fully connected by fixed inhibitory connections. The experiment reported here used the parameter value $\lambda = 0.005$.

Note that adaptive connections exist only between the SOMs and label field, while constant-valued (-0.8) inhibitory spreading activation connections exist within each SOM and within the label field.

4. Simulating Word Learning in a Robotic System

4.1. *Procedure*

The procedure in the robot experiment was kept as close as possible to the procedure in the empirical task. As in the empirical study, the experiment was run 12 times per condition and trial order and counterbalancing were the same. In addition, we simulated the three developmental timescales that McMurray *et al*. (20012) argue allow word learning to emerge via fast-mapping.

4.1.1. *Long-term vocabulary learning*

To simulate children's vocabularies, we taught the robot a vocabulary of "known" labels in an initial training session during which the robot was not exposed to the novel words used in the subsequent experiment. The SOMs were provided with object and label input for the 18 competitor objects it would encounter during the referent selection phase. For each individual object, the experimenter placed the object centrally in the robot's field of vision on a white surface and allowed the SOMs to settle – equivalent to allowing the children to look at the objects before providing the target label. Once the SOMs had settled (that is, once the system had formed a representation of that object), the experimenter provided the label SOM with the appropriate input, keeping the object in view (equivalent to the ostensive labelling shown to facilitate word learning in children; Axelsson, Churchley, & Horst, 2012). Each object received 20 unambiguous labelling events. Thus, the robot began the experiment with a robust known vocabulary.

4.1.2. *In-the-moment referent selection*

The robot was presented with the same referent selection and test trials as children in the empirical task, again across two, three and four competitor conditions. During referent selection the robot was presented with four sets of objects on a white tabletop via eight fast-mapping trials. Each set consisted of a novel object and two, three or four known competitors selected from the pre-trained set (see Fig. 1 for an example referent selection trial). As in HSP, object locations and trial order were pseudorandomised across trials. Thus, the same set of objects was never presented on successive trials, known/novel trials occurred no more than twice in succession and each novel label/object pair was encountered in first, second, third or last position equally often. All objects were placed in the robot's field of vision and the SOMs were allowed to settle (equivalent to the three-second pause before labelling in the empirical study). Then, the experimenter labelled the object five times with either a known (pretrained) or novel label. Following labelling, the robot moved its head to centre its field of vision on each object in turn. This caused a node in the label SOM to become activated. If the SOM activated the appropriate label node for the target object, the robot's response was scored as correct, and if not, the robot's response was scored as incorrect. For example, on a known trial, activation of the *horse* node in response to the horse stimulus would be scored correct, and activation of the *yok* label would be scored as incorrect. Each novel object was presented twice and served as a target once. Object locations and trial order were pseudorandomised as in the empirical study.

4.1.3. *Word learning*

After referent selection the robot was presented with four test trials that proceeded in an identical manner to the referent selection trials. As in the empirical study, each of the four novel targets appeared on each test trial, and each served as the target on one trial. If the model had learned and retained novel label-object mappings during the referent selection phase, then it should activate the appropriate label node in response to each novel object at levels greater than expected by chance.

4.2. ***Results***

Results from the robot task are depicted in Figure 2. In line with a previous simulation of a related fast-mapping task (Twomey *et al.*, 2013; 2014), the model successfully mapped known names to known objects and novel names to novel objects during referent selection, and did so at levels greater than expected by chance (100% correct on all known trials, all *p*s < .0001 on novel trials; all *p*s two-tailed; note that chance = 0.33, 0.25 and 0.20 in the two, three and four competitor conditions, respectively). Thus, the model captured HSP's referent selection results. At test, the model retained novel label-object mappings at levels greater than expected by chance (0.25) in the two competitor condition only (two competitor: $t(11) = 8.86$, $p < .0001$, $d = 5.34$; three competitor, $t(11) = 1.39$, *ns.*, $d = 0.84$; four competitor, $t(11) = 0.56$, *ns.*, $d = 0.34$). We submitted the model's proportion of correct choices on test trials to an ANOVA with condition as a between-subjects factor. This revealed a significant effect of

Condition, $F(1,34) = 34.19$, $p < .0001$, $\eta_p^2 = 0.50$. Planned comparisons revealed that the model made significantly more correct choices in the two competitor condition than in the three or four competitor conditions (both *p*s < .0001). Thus, the model's responding at test also captured that of the children.

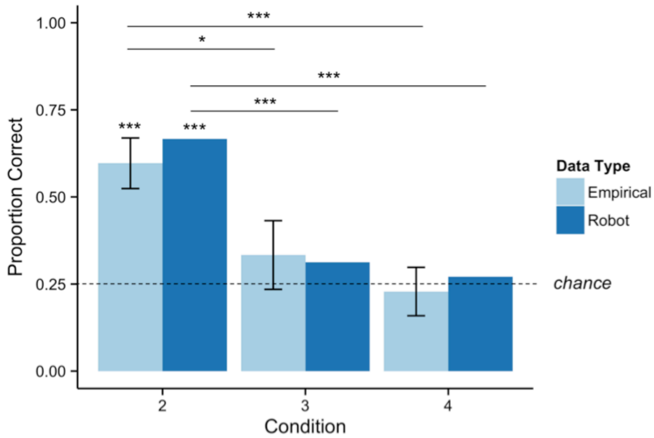

Fig. 2. Proportion correct of children's (light blue bars) and the model's correct choices (dark blue bars) at test. *** $p < .001$, ** $p < .01$, * $p < .05$.

5. Discussion

In the current study we used an embodied neural network model (Metta *et al*., 2010; Morse *et al*., 2010) to replicate 30-month-old children's behaviour in a word learning task (Horst et al., 2010). Children in the empirical task were presented with a referent selection phase in which they were asked to select novel objects in response to novel names, from an array in which all other objects were known. When tested on retention of novel label-object mappings, only children who had initially encountered novel objects alongside two competitors successfully retained those mappings; children who saw more competitors did not retain novel labels. We taught the model a known vocabulary and then presented it with a maximally similar task. The model correctly mapped known and novel names to the appropriate objects during referent selection – that is, the model exhibited the same in-the-moment fast-mapping as the children. At test, only when novel objects had initially been encountered alongside two (but not three or four) competitors did the model successfully retain label-object mappings. Again, then, the model exhibited the same word learning behaviour as the children.

This simulation supports HSP's account of the processes underlying children's fast-mapping and word learning, and demonstrates that fast-mapping involves attention to known competitor objects as well as encoding the novel object-label mapping. Specifically, fast-mapping in the model occurs because inhibition from the strong known label-known object connections prevents the formation of new known label-novel object connections, meaning that the only mapping available during referent selection is between the novel label and the novel object (see also McMurray *et al.*, 2012; Twomey *et al.*, 2013).

Importantly for the current study, however, time is key to capturing children's behaviour at test. Specifically, during referent selection the robot is provided with a label, which causes a spike of activation in the relevant label node. The robot then "looks" at each object in turn to establish whether the current label is associated with each object. After looking at all objects in the array, learning is applied. However, the activation in the label node decays over time. Thus, after the robot has encoded five objects (i.e., one target and four competitors), activation in the label node is lower than after the robot has encoded three objects (i.e., one target and two competitors). Thus, the novel label-object mappings formed in the three and four competitor conditions are weaker than the equivalent mappings formed in the two competitor condition; so weak, in fact, that the model is unable to use them to choose the correct object at test. Thus, consistent with McMurray et al. (2012), we provide an explicit, mechanistic account of this word learning behaviour as emergent from the interaction between long-term learning (its pretrained, "known" vocabulary), its cross-situational learning, and its in-the-moment fast-mapping. However, the current study adds to this account the effect of micro-level temporal dynamics as instilled by activation decay during the time taken to look from one object to another. HSP note that children spent approximately 0.5s observing each object, irrespective of condition. Taken together, these studies demonstrate that word learning is an exquisitely sensitive and flexible process.

The current study is one of the first full-scale replications of a fast-mapping experiment using an embodied robotic system, and is the first to examine the effect of context on word learning via fast-mapping. As such, it contributes in a broader sense to the emerging interdisciplinary literature in the cognitive sciences that in recent years has begun to apply mathematical and computational innovations to some of the decades-old puzzles of developmental psychology (e.g., McMurray *et al*., 2012; Schlesinger, 2009; Westermann & Mareschal, 2014), while helping us build an explicit account of the complex and subtle temporal, environmental and physiological interactions that govern word learning and cognitive development.

References

1. Alvy Ray, S. (1978). Color gamut transform pairs. *SIGGRAPH Computer Graphics*, *12*(3), 12–19. doi:10.1145/965139.807361
2. Axelsson, E. L., Churchley, K., & Horst, J. S. (2012). The right thing at the right time: why ostensive naming facilitates word learning. *Frontiers in Psychology*, 3.
3. Carey, S., & Bartlett, E. (1978). Acquiring a single new word. Papers and Reports on Child Language Development, 15, 17–29. doi:10.1037/0012-1649.38.6.967
4. Horst, J. S., Scott, E. J., & Pollard, J. P. (2010). The role of competition in word learning via referent selection. *Developmental Science, 13*(5), 706–713.
5. Kohonen, T. (1998). The self-organizing map. *Neurocomputing, 21*(1-3), 1–6.
6. Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, *14*(1), 57–77.
7. Marocco, D., Cangelosi, A., Fischer, K., & Belpaeme, T. (2010). Grounding action words in the sensorimotor interaction with the world: experiments with a simulated iCub humanoid robot. *Frontiers in Neurorobotics, 4*.
8. McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review, 119*(4), 83877. doi: 10.1037/a0029872
9. Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L., & Montesano, L. (2010). The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks*, *23*(8-9), 1125–1134.
10. Montesano, L., Lopes, M., Bernardino, A., & Santos-Victor, J. (2008). Learning object affordances: from sensory–motor coordination to imitation. *IEEE Transactions on Robotics, 24*(1), 15–26.
11. Morse, A. F., Belpaeme, T., Cangelosi, A., & Smith, L. B. (2010). Thinking with your body: Modelling spatial biases in categorization using a real humanoid robot. In S. Ohlsson & R. Catrambone (Eds.), *32nd Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society.*
12. Morse, A. F., de Greeff, J., Belpeame, T., & Cangelosi, A. (2010). Epigenetic Robotics Architecture (ERA). *Autonomous Mental*

Development, IEEE Transactions on Autonomous Mental Development, 2(4), 325–339.

13. Quine, W. V. O. (1960). *Word and object: An inquiry into the linguistic mechanisms of objective reference*. Cambridge: MIT Press.
14. Schlesinger, M. (2009). The robot as a new frontier for connectionism and dynamic systems theory. In J. P. Spencer, M. Thomas, & J. L. McClelland (Eds.), *Toward a Unified Theory of Development: Connectionism and Dynamic System Theory Re-Considered* (pp. 182–201). New York, NY: Oxford Unviersity Press.
15. Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition, 106*(3), 1558–1568.
16. Twomey, K. E., Ranson, S. L., & Horst, J. S. (2014). That's more like it: Multiple exemplars facilitate word learning. *Infant and Child Development*, *23*(2), 105–122.
17. Twomey, K. E., Horst, J. S., & Morse, A. F. (2013). An embodied model of young children's categorization and word learning. In L. J. Gogate & G. Hollich (Eds.), *Theoretical and Computational Models of Word Learning: Trends in Psychology and Artificial Intelligence* (pp. 172–196). Hershey, PA: Information Science Reference.
18. Westermann, G., & Mareschal, D. (2012). Mechanisms of developmental change in infant categorization. *Cognitive Development 27*(4), 367-382.
19. Westermann, G., & Mareschal, D. (2014). From perceptual to language-mediated categorization. *Philosophical Transactions of the Royal Society B: Biological Sciences,* 369(1634), 20120391.
20. Yürüten, O., Uyanık, K., Çalışkan, Y., Bozcuoğlu, A., Şahin, E., & Kalkan, S. (2012). Learning adjectives and nouns from affordances on the iCub humanoid robot. In T. Ziemke, C. Balkenius, & J. Hallam (Eds.), *From Animals to Animats 12* (Vol. 7426, pp. 330–340). Berlin: Springer.
21. Zosh, J. M., Brinster, M., & Halberda, J. (2013). Optimal contrast: Competition between two referents improves word learning. *Applied Developmental Science*, *17*(1), 20–28. doi:10.1080/10888691.2013.748420