All Things Considered: Dynamic Field Theory Captures Effect of Categories on Children's Word Learning

Katherine E. Twomey[1] **& Jessica S. Horst**[1]
(k.twomey@sussex.ac.uk, jessica@sussex.ac.uk)
[1]School of Psychology, University Of Sussex, Falmer, BN1 9QH, UK

Abstract. Recent research demonstrates that both real-time variability in perceptual input and task demands influence young children's word learning and categorisation. The current study extends these findings by testing both children and a dynamic field theory (DFT) computational model in a category labelling task. Specifically, children and the model were introduced to multiple category members that were either moderately or highly variable. Both children and the model were better able to learn category labels when the individual category members were moderately variable. Overall, these findings have implications for both our understanding of children's categorisation and the use of computational models to investigate cognition more generally.

1 WORD LEARNING AND CATEGORISATION

In order to understand the world, children must learn to label and categorise objects in their environments; they do so astonishingly quickly [1]. The complexity of learning a single new word is well-documented [2]: children must not only parse the speech stream into individual words but also determine the meaning of a word from a seemingly infinite array of possible referents [3]. Children's ability to rapidly link a novel label to a novel object is known as fast mapping [4; 5; 6], however, as demonstrated by Horst & Samuelson [7], fast mapping is only one part of the word learning process. To have truly learned a word, children must be able to use that word after a delay or in a new context [8].

By the time children begin to learn words, they are already experienced categorisers. Each new word they encounter refers not just to a single object, but to a category of objects [9; 10]. For example, when a child learns that their family collie is called a "dog", she may also learn that their neighbours' poodle is a "dog", that her cuddly toy is a "dog" [11], and so on. Research in domains as diverse as motor development [12], phonological acquisition [13], and visual categorisation [14] has demonstrated that multiple and variable experiences facilitate learning [15; 16]. Further, variability among category members has also been shown to affect categorisation; that is, categorisation is facilitated by experience with multiple exemplars [17].

However, how variability among category members influences category label learning remains unclear. Recent research demonstrated that 30-month-old children exposed to multiple category members (exemplars) were significantly more likely to retain the category label after a 5-minute delay than children exposed to a single category member multiple times [18]. These data suggest that experience with multiple exemplars facilitates word learning. However, in this case the category members only varied in one feature (colour). The current research extends these findings both empirically and computationally with highly variable categories to further understand how categorisation influences word learning.

2 SUPPORTING EMPIRICAL DATA

2.1 Method

2.1.1 Participants

Twenty-four typically-developing, monolingual English-speaking 30-month-old children participated. 12 children were randomly assigned to the *narrow* condition, and 12 to the *variable* condition.

2.1.2 Stimuli

Known stimuli for all conditions consisted of 18 objects likely to be known to 30-month-old children (e.g., a toy chicken or a toy bike). Novel stimuli consisted of nine novel exemplars from three categories (examples are depicted in Figure 1). For children in the *narrow* condition, novel exemplars were moderately variable and differed only in colour. For children in the *variable* condition, novel stimuli were highly variable and differed in colour, shape and texture. For extension trials atypical exemplars from the novel categories were used. On the extension trials the same stimuli were used for both conditions.

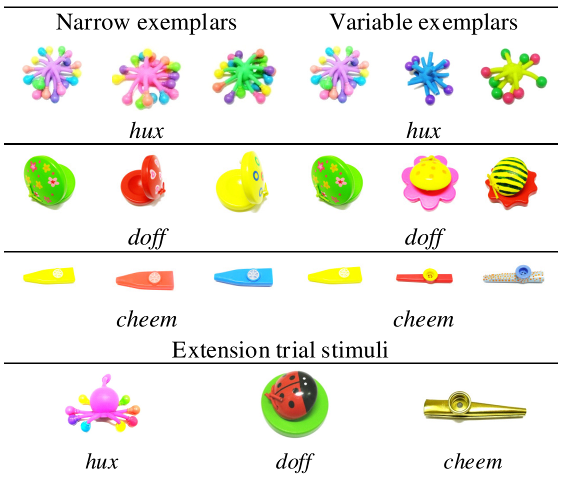

Figure 1. Novel stimuli used in the experiment

2.1.3 Procedure and design

The experiment consisted of three phases: referent selection (18 trials), retention (three trials) and extension (three trials). An example referent selection trial is depicted in Figure 2. On each referent selection trial children saw an array of three objects (two known, one novel) and were asked to get either the novel or one of the known objects (e.g., “can you get the *hux*?”). Overall, children received nine known name trials and nine novel name trials. Children received three trials per novel category (e.g., *hux*). Across trials, children saw novel categories with either *narrow* or *variable* exemplars.

.

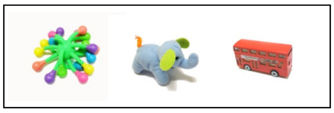

“Can you get the *hux*?”

Figure 2. Example referent selection trial

After a 5-minute break the test phase began. On each of the three retention test trials children saw an array of three objects (one from each of the just-encountered novel categories) and were asked to get each of the objects across trials (for an example, see Figure 3). Extension trials immediately followed and were identical to retention trials except that the atypical exemplars were used.

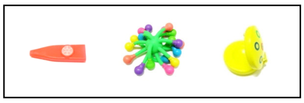

“Can you get the *hux*?”

Figure 3. Example retention trial

2.2 Results

2.2.1 Referent selection

Results are depicted in the left panel of Figure 4. All children were very good at referent selection. Children in both conditions chose the target object at significantly greater than chance levels on both known name trials (.33, all *p*s two-tailed, $t(11) = 10.51$, $p < .0001$, $d = 3.05$ and $t(11) = 17.42$, $p < .0001$, $d = 5.05$, respectively) and novel name trials ($t(11) = 5.95$, $p < .0001$, $d = 1.73$ and $t(11) = 15.58$, $p < .0001$, $d = 4.52$, respectively). Unpaired t-tests revealed no difference between conditions for either known or novel referent selection (known: $t(22) = -0.30$, *ns*; novel: $t(22) = -0.63$, *ns*.) Thus, whether children saw *narrow* or *variable* exemplars had no effect on referent selection.

2.2.2 Test trials

Results are depicted in the right panel of Figure 4. Data for test trials were submitted to a repeated measures ANOVA with Trial Type (retention, extension) as the repeated measure and Stimui (narrow, variable) as a between-subjects factor, The ANOVA revealed a significant interaction between Trial Type and Stimuli, $F(1, 22) = 7.86$, $p = .01$. To unpack this interaction, planned one-tailed *t*-tests against chance were performed. Only children in the *narrow* condition retained novel labels at levels significantly greater than chance, $t(11) = 4.73$, $p < .001$, $d = 1.38$. Importantly, this replicates Horst et al.’s [18] finding: experience with a category of objects clearly facilitates children’s ability to retain labels. A planned, unpaired *t*-test revealed a significant difference between conditions, $t(22) = 2.84$, $p < .01$, $d = 1.22$. In contrast, only children in the *variable* condition extended the novel labels at levels greater than chance, $t(11) = 2.60$, $p < .05$, $d = 0.76$. Thus, encountering a variable category facilitates children’s ability to extend labels to new category members [15].

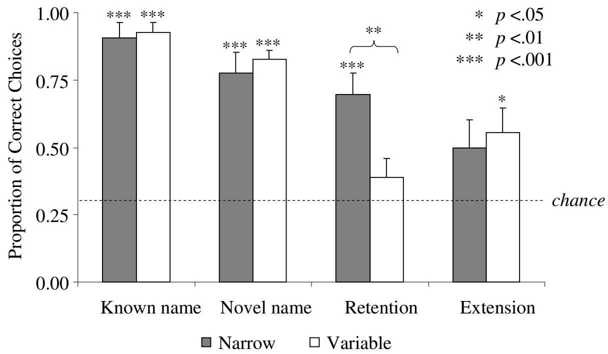

Figure 4. Experimental results

2.3 Discussion

Only children in the *narrow* condition retained novel category labels; however, these children did not extend this newly-learned label to a completely novel atypical category member. In contrast, children in the *variable* condition did not retain the novel labels but were nonetheless able to extend novel category labels. We explored this surprising result by simulating the task using a dynamic field theory model.

3 WORD LEARNING IN-THE-MOMENT

Dynamic Field Theory (DFT) is a formal instantiation of Dynamic Systems Theory (DST) [19] which has been successfully implemented to model children’s decision-making processes in various motor and perceptual tasks [20; 21] as well as larger-scale robotic systems [22]. According to DST, behaviour is self-organising in the moment and is thus inextricably linked to real-time input, as well as just-past experience and longer-term learning history [23]. DST has been applied in many domains to explain hitherto puzzling phenomena; for example, the sudden disappearance of young children’s stepping reflex [24], perseverative reaching in A-not-B tasks [25] and variable development of goal-directed reaching [12]. More recently, DST has been formalised in the DFT [26], a dynamic neural field framework in which self-sustaining, stable peaks of activation reflect self-organised behaviours. Critically, the DFT allows us to examine the interplay of multiple timescales underlying children’s in-the-moment choices in experimental settings.

The goal of this simulation is to investigate whether small changes in stimuli in word learning tasks can give rise to better retention and extension of novel category labels. DFT models have successfully captured experimental data from looking tasks

[27] dimensional change card-sorting tasks [28] and novel noun generalisation tasks [29]. The current simulation adapts Faubel & Schöner's [22] feature binding DFT model of object recognition to a word learning context. If the simulation reflects the experimental data, this suggests that the apparently complex learning processes driving word learning may, in fact, depend on the simple, bottom-up, dynamic associative mechanisms that underlie DFT models.

3.1 The current simulation

3.1.1 Architecture

DFT models consist of continuous, topologically functional neural fields in which spreading activation governed by local excitation/global inhibition [30] generates localised, self-sustaining peaks of activation [31]. The current simulation, depicted in Figure 5, consists of two 2-dimensional dynamic neural fields; specifically, a perceptual layer coupled reciprocally to a memory layer. Activation in the perceptual layer is generated by input along the *label* and *object* dimensions. and is captured by the general equation below:

$$\begin{aligned}&\tau\dot{u}_{o,l}(x,t) = -u_{o,l}(x,t) + h + S_{o,l}(x,t) \\ &+ \int w(x-x')\sigma(u(x',t))dx'\end{aligned} \qquad (1)$$

where $\dot{u}_{o,l}(x,t)$ is the rate of change of activation level across the object (o) and label (l) dimensions at location x, as a function of time (t) mediated by the timescale of the dynamics, τ. Current activation in the perceptual layer, $-u_{o,l}(x,t)$, receives external, experimenter-defined input, $S_{o,l}(x,t)$. Activation in the perceptual and memory layers is subject to excitatory and inhibitory interaction defined by a Gaussian kernel with weight w, and width σ. The resting level of the system is defined by h <0.

Units of representation are peaks of activation. The formation of a self-sustaining peak at any point in the perceptual layer represents a mapping between input along the object dimension and the label dimension. Activation from these peaks spreads to the memory layer, leaving a corresponding, slow-decaying memory trace. Activation in the memory trace acts as short-term memory, by feeding activation back to the perceptual layer, thus facilitating subsequent object-label mappings.

3.1.2 Stimuli and procedure

Known object stimuli were presented as inputs along the object dimension (length = 531 neurons) at intervals of at least 20 neurons. Novel object stimuli were presented at intervals of at least 20 neurons to their nearest known neighbour, with spacing between novel stimuli varying according to condition (see below). On every trial, each object stimulus was separated from its nearest neighbour by at least 75 neurons. Similarly, label stimuli were presented as inputs regularly spaced along the label dimension (length = 22 neurons). In the current model a single neuron on the label dimension was arbitrarily assigned to a single label. However, the model is sufficiently flexible for future work to explore further effects of categorisation, such as phonetic similarity of labels, or the global/basic distinction [32].

Variability in object inputs to the model reflects the variability in category structure encountered by children during the experiment. Specifically, the model is either presented with *narrow* category exemplars, in which novel object input is presented at the central category exemplar and two nearby locations, or with *variable* category exemplars, in which novel object input is presented at the central category exemplar and two more distant locations. For example, *narrow* stimuli might consist of input at locations 114, 115 and 116 along the *object* dimension, while *variable* stimuli might consist of input at positions 109, 115 and 121 along the *object* dimension.

Like the children, the model is presented with 18 referent selection, three retention and three extension trials, using dimensional cueing on each trial to distribute the presentation of stimuli and object labelling over time.

At the beginning of each referent selection trial, the model is presented with "known" cues located at the intersection between object and label for the two known objects, generating two stable peaks, and a "novel" cue at a specific location along the object dimension but generic along the label dimension (see Panel A of Figure 5). Thus, input for novel objects could correspond to any label.

Next, the model is presented with a ridge of input along the label dimension. This new label input intersects with either the existing "known" or "novel" object cues (see Panel B of Figure 5). Formation of a peak at any point in the perceptual layer is considered to reflect the model's choice of object in response to a given label; that is, when a peak is formed the model has fast mapped a label to an object. Note that both correct and incorrect choices are included in the analysis.

Object cues for test trials consist of three generic ridges of activation at the previously encountered novel object locations along the object dimension. The model then receives label input as during referent selection. The three subsequent extension trials are identical to retention trials except that the initial novel object cues are given at locations close to but not identical to the previously locations. Thus, during extension trials the model associates novel labels with completely new novel objects.

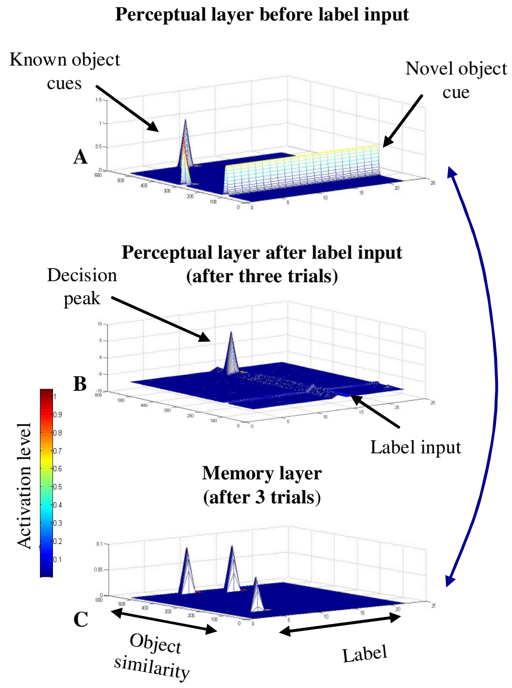

Figure 5. Architecture of the DFT model

3.2 Results

Simulation data are depicted in Figure 6. The model is very accurate on referent selection trials, both with narrow and variable categories. Like the children in our experiment, when the model is presented with *narrow* categories it correctly associates previously-encountered novel category members with previously-encountered novel labels on retention trials and does not associate completely novel, atypical exemplars with previously-encountered labels on extension trials. In contrast, like the children, when the model is presented with *variable* categories it does not associate previously-encountered novel category members with previously-encountered labels on retention trials and does associate completely novel atypical exemplars with previously-encountered labels on extension trials.. Thus, preliminary simulation data reflect children's behaviour in the word learning task, even reproducing the counterintuitive result in the *variable* condition.

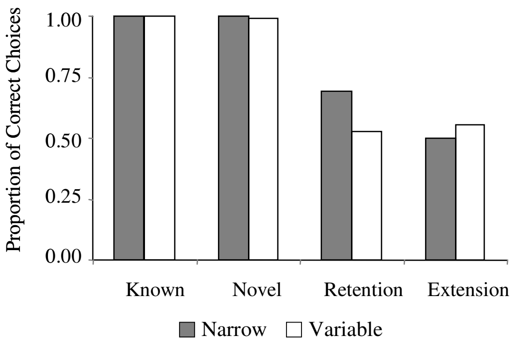

Figure 6. Simulation results

4 DISCUSSION

We have demonstrated both experimentally and computationally that word learning is susceptible to task effects; that is, small changes in stimuli during a fast-mapping task can dramatically influence retention and extension of novel labels. For example, when children encounter wide within-category variability, they do not show evidence of retaining a label for this category, despite being able to extend this label to a completely novel category member. A dynamic field simulation captures this phenomenon by repeated association of different perceptual input over time, generating a remarkably similar pattern of results.

This model offers considerable opportunity for further investigation of the interplay between category variation and word learning. For example, when a child sees an object, she is aware of its colour, shape and the visual components of its texture. In the current model, however, visual input is simplified and schematised: all visual input is collapsed across an overall "perceptual similarity" metric and presented to a single perceptual layer. The addition of further layers representing, for example, colour, shape and texture, allowing the separation of colour, shape and material inputs (cf. [22]), represents an important step towards understanding what constitutes "variability" for children learning to categorise. Comparable extensions of the model, for example taking into account motor feedback, and potential hybridisation with other connectionist architectures more commonly used in computer vision (for example, Self-Organising Maps, [33]), also offer opportunities for its deployment in an embodied agent.

These results have implications for our broader understanding of cognitive development. First, we have extended the DFT to reliably simulate children's fast mapping and word learning behaviour. Second, simulation data suggest that absence of evidence for a behaviour in one context does not imply that the behaviour will not be seen in a different context. Further, as DFT models are simple, associationist, spreading-activation networks, the present data lend further weight to the growing body of evidence suggesting that cognition develops in a bottom-up manner via associations learned from statistical regularities in the input, without recourse to innate learning mechanisms [34]. Taken together, the present data suggest a productive future direction for the integration of psychological and computational research in cognitive development.

REFERENCES

[1] L. Fenson, P.S. Dale, J.S. Reznick, E. Bates, D.J. Thal, and S.J. Pethick, Variability in Early Communicative Development. Monographs of the Society for Research in Child Development 59 (1994).
[2] L.K. Samuelson, Statistical regularities in vocabulary guide language acquisition in connectionist models and 15-20-month-olds. Developmental Psychology 38 (2002) 1016-1037.
[3] W.V.O. Quine, Word and Object, MIT Press, Cambridge: MA, 1960.
[4] S. Carey, and E. Bartlett, Acquiring a single new word. Papers and Reports on Child Language Development 15 (1978) 17-29.
[5] T.H. Heibeck, and E.M. Markman, Word Learning in Children - an Examination of Fast Mapping. Child Development 58 (1987) 1021-1034.
[6] L. Gershkoff-Stowe, and E.R. Hahn, Fast mapping skills in the developing lexicon. Journal of Speech Language and Hearing Research 50 (2007) 682-697.
[7] J.S. Horst, and L.K. Samuelson, Fast mapping but poor retention by 24-month-old infants. Infancy 13 (2008) 128-157.
[8] E. Mather, and K. Plunkett, Learning words over time: The role of stimulus repetition in mutual exclusivity. Infancy 14 (2009) 60-76.
[9] R.M. Golinkoff, C.B. Mervis, and K. Hirsh-Pasek, Early object labels - the case for a developmental lexical principles framework. Journal of Child Language 21 (1994) 125-155.
[10] L.K. Samuelson, and L.B. Smith, Early noun vocabularies: do ontology, category structure and syntax correspond? Cognition 73 (1999) 1-33.
[11] L.B. Smith, S.S. Jones, B. Landau, L. Gershkoff-Stowe, and L. Samuelson, Object name learning provides on-the-job training for attention. Psychological Science 13 (2002) 13-19.
[12] E. Thelen, and J.P. Spencer, Postural control during reaching in young infants: A dynamic systems approach. Neuroscience and Biobehavioral Reviews 22 (1998) 507-514.
[13] G.C. Rost, and B. McMurray, Speaker variability augments phonological processing in early word learning. Developmental Science 12 (2009) 339-349.
[14] B.A. Younger, and L.B. Cohen, Infant perception of correlations among attributes. Child Development 54 (1983) 858-867.
[15] P.C. Quinn, P.D. Eimas, and S.L. Rosenkrantz, Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. Perception 22 (1993) 463-475.
[16] J.S. Horst, L.M. Oakes, and K.L. Madole, What does it look like and what can it do? Category structure influences how infants categorize. Child Development 76 (2005) 614-631.
[17] K. A. Kovack-Lesch & L. M. Oakes, Hold your horses: How exposure to different items influences infant categorization. Journal of Experimental Child Psychology 98 (2007), 69-93.

[18] J.S. Horst, K.E. Twomey, and S.L. Ranson, That's more like it: More exemplars facilitate word learning. (in prep.).
[19] E. Thelen, and L.B. Smith, A Dynamic Systems Approach to the Development of Cognition and Action, MIT Press, Cambridge, Mass., 1996.
[20] L.K. Samuelson, and J.S. Horst, Confronting complexity: insights from the details of behavior over multiple timescales. Developmental Science 11 (2008) 209-215.
[21] V.R. Simmering, A.R. Schutte, and J.P. Spencer, Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. Brain Research 1202 (2008) 68-86.
[22] C. Faubel, and G. Schöner, Learning to recognize objects on the fly: A neurally based dynamic field approach. Neural Networks 21 (2008) 562-576.
[23] J.P. Spencer, M. Clearfield, D. Corbetta, B. Ulrich, P. Buchanan, and G. Schöner, Moving toward a grand theory of development: In memory of Esther Thelen. Child Development 77 (2006) 1521-1538.
[24] E. Thelen, and B.D. Ulrich, Hidden skills - a dynamic-systems analysis of treadmill stepping during the 1st year. Monographs of the Society for Research in Child Development 56 (1991) R1-&.
[25] L.B. Smith, E. Thelen, R. Titzer, and D. McLin, Knowing in the context of acting: The task dynamics of the A-not-B error. Psychological Review 106 (1999) 235-260.
[26] J.P. Spencer, and G. Schöner, Bridging the representational gap in the dynamic systems approach to development. Developmental Science 6 (2003) 392-412.
[27] S. Perone, J. Spencer, and G. Schöner, A dynamic field theory of visual recognition in infant looking tasks., Cognitive Science Society, 2007.
[28] A.T. Buss, and J.P. Spencer, The emergence of rule-use: a DNF model of the DCCS. 30th Annual Meeting of the Cognitive Science Society, Washington, DC, 2008, July.
[29] L.K. Samuelson, A.R. Schutte, and J.S. Horst, The dynamic nature of knowledge: Insights from a dynamic field model of children's novel noun generalization. Cognition 110 (2009) 322-345.
[30] S.I. Amari, Dynamics of pattern formation in lateral-inhibition type neural fields. Biological Cybernetics 27 (1977) 77-87.
[31] J.P. Spencer, V.R. Simmering, A.R. Schutte, and G. Schöner, What does theoretical neuroscience have to offer the study of behavioural development? Insights from a dynamic field theory of spatial cognition. in: J.M. Plumert, and J.P. Spencer, (Eds.), the Emerging Spatial Mind, Oxford Universoty Press, New York, 2007.
[32] K. Plunkett, J. F. Hu, and L. B. Cohen, "Labels can override perceptual categories in early infancy," Cognition, vol. 106, pp. 665-681, Feb 2008.
[33] T. Kohonen, The Self-Organizing Map, a possible model of brain maps, Brain and Values (1998) 207-236.
[34] J.P. Spencer, M. Thomas, and J. McClelland, Toward A Unified Theory of Development: Connectionism and Dynamic Systems Theory Reconsidered, Oxford University Press, New York, NY, 2009.