A PDP Account of Transitions in Conceptual Development

Robert J. Powers, David C. Plaut

Department of Psychology, Carnegie Mellon University

powers@cmu.edu
plaut@cmu.edu

Abstract

As children gain experience with the world, the organization of their conceptual knowledge becomes increasingly complex, as reflected by the successive emergence of sensitivity to different types of similarity over the course of development. Though this phenomenon has been well-studied, it is often explained by reference to innate, domain-specific mechanisms that are stipulated to come on-line at specific ages. We present a Parallel Distributed Processing (PDP) model that learns from the structure of its environment and exhibits transitions in the relative salience of perceptual, thematic, and taxonomic similarity, as observed empirically, without any built-in knowledge or changes to the learning mechanism.

Background

The Role of Similarity: Empirical Evidence

Things can be similar in different ways. Sensitivity to these different types of similarity varies with age, reflecting underlying conceptual change. Three types of similarity are important:

Type	Depends on	Age range	Studies
perceptual	surface-level features	infants to 3-years	[4] [11]
thematic	co-occurrence and complementary features	2- to 5-years	[3] [16]
taxonomic	shared/inherited abstract features	5-years to adult	[16] [9]

Theoretical Accounts

- *Knowledge-based*: The traditional perspective is that perceptual information alone is not sufficient to develop adult-like categories, and that some knowledge must be innate [1] [5] [6].
- *Similarity-based*: A family of alternatives to knowledge-based theories. The central claim is that perceptual input is rich, and a learning mechanism sensitive to statistical regularities can learn higher order, taxonomic structure [10] [13] [15]. PDP models are mechanistic descriptions belonging to this family.

Limitations of Existing PDP Models

- Autoencoders have been used [12], but they cannot learn structure beyond that present in the input
- More advanced models usually train on explicit semantic features, which is not how children learn in most cases [8] [13]
- These and related models usually focus on one aspect of the problem, usually either ignoring the role of perceptual features [14] [8] or of thematic co-occurrence [12] [13]

Present Work

We offer an extension of previous models that is similar in structure to an autoencoder network, but that is capable of abstraction beyond input similarity. With such a model, is perceptual information alone enough to account for the observed developmental pattern without the need for built-in knowledge?

Methods

Artificial Environment

- *Building event structure*: The goal was to train the model on perceptual representations of events. Each event consisted of two objects bearing some relation, and the whole event could be both viewed and described with auditory labels.
  - *Objects*: Thirty nouns were chosen from a feature norms database [7]. Features were narrowed down to those tagged as "visual" in [7]. Thematic and taxonomic categories were chosen by the experimenter (see Figure 1).
  - *Relations*: Five relations ("drives", "wears", "pets", "eats", and "is inside of") were chosen to bind objects thematically and taxonomically. Each relation was given a visual representation by pseudo-randomly generating a 20-dimensional bit vector.
  - *Labels*: Each object and relation was assigned a label by pseudo-randomly generating a 20-dimensional bit vector.

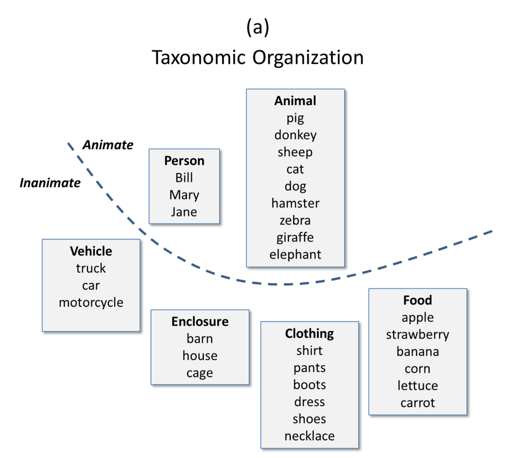

Figure 1: *(a) Objects sorted by taxonomic category. (b) Objects sorted into thematic groups.*

Network Architecture

The model was a simple recurrent network [2]

- *Input/Output*:
  - Input and output layers both divided into visual and auditory subsets.
  - Modality-specific processing constraints
    - ∗ Visual "scene" spatially organized
    - ∗ Auditory "sentence" temporally organized

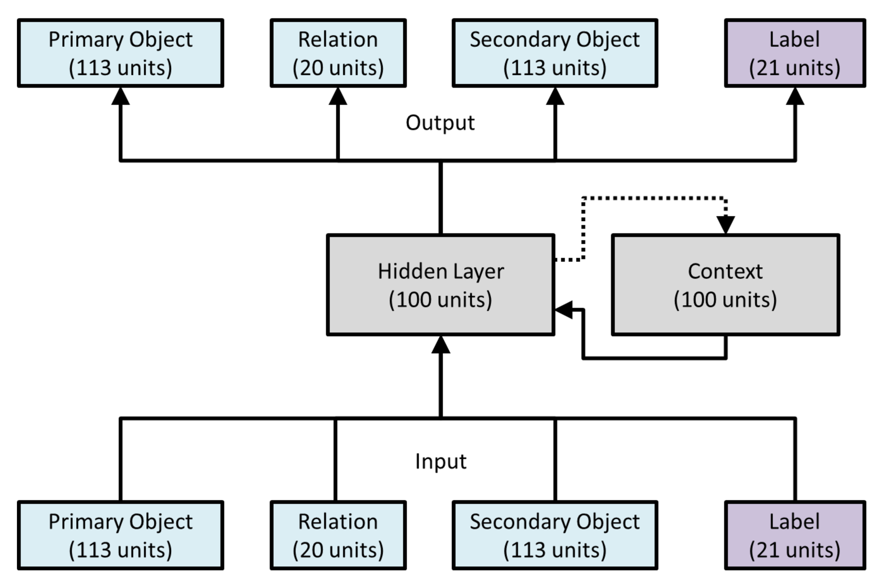

Figure 2: *Network diagram. Input and output layers are divided by modality; here, visual units are shown in light blue, and auditory units in lavender.*

Training Procedure

Error was accumulated over the entire training set before adjusting the weights via back-propagation. Each pass through the training set constituted one epoch.

- *Task*: The network had to reconstruct the full event for each trial, given only partial perceptual input.
- *Trial types*:
  - Unimodal visual: One object missing, relation missing, or both.
  - Unimodal auditory: Last word of the input missing
  - Bimodal: The union of the unimodal conditions.
- *Active/Passive*: The network was trained on each event in both active and passive forms. The "primary object" slot always matched the first object label.
- *Distortion*:
  - Objects: Gaussian noise added to each unit (mean=0, sd=0.05)
  - Relations and labels: Prior to each epoch, a new exemplar close to the original prototype was chosen by changing features with p=0.05.

Testing Procedure

The network was tested after every 10 epochs by presenting each visual object pattern in the "primary object" slot for a single tick. Activation values for all hidden and output units were recorded in response to each object.

Results & Discussion

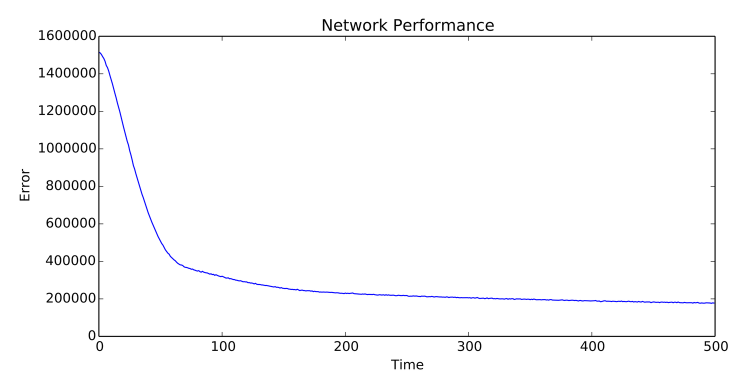

Figure 3: *Cross-entropy error over time, summed over the entire training set.*

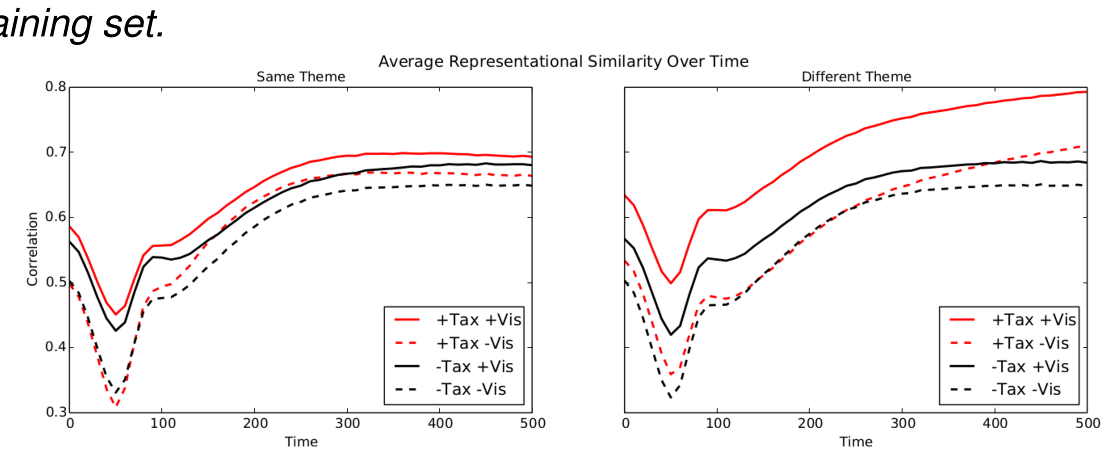

Figure 4: *Pairwise correlations between object representations (concatenated hidden and output vectors) over the course of training. Color indicates taxonomic similarity (red means similar). Line type indicates visual similarity (solid means similar). The left plot shows within-theme correlation, while the right shows across-theme.*

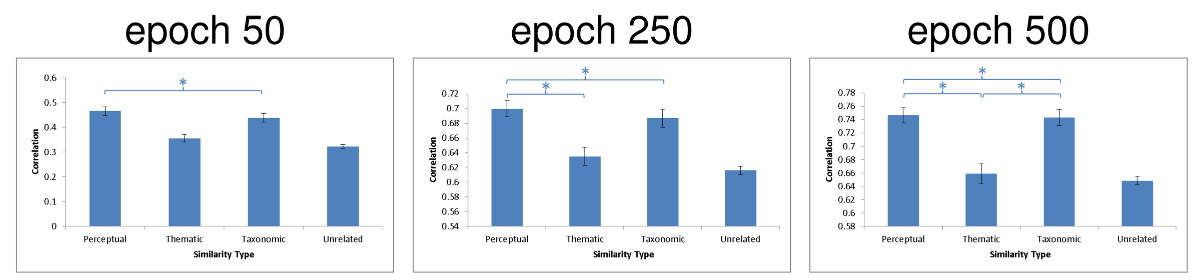

Figure 5: *Correlations with significant interactions shown for early, middle, and late training, $\alpha = 0.05$. Error bars display 95% confidence intervals.*

This network replicates some findings from empirical studies with infants and children.

- Perceptually similar objects start out with more similar representations. Since within-category perceptual similarity tends to be higher than between-category similarity, taxonomic clustering depends largely on perceptual similarity.
- Later, representations for visually dissimilar, taxonomically related objects get an "associative boost" from thematic co-occurrence.
- Finally, the network becomes able to represent pure taxonomic relations as similar, benefiting from bootstrapping of perceptual and thematic relations.

Conclusions

- With the appropriate pressure, a simple network can abstract higher-order structure, such as taxonomic relatedness, from low-level input
- No innate knowledge structures or explicit category representations are required to account for developmental data.

References

[1] Susan Carey and Elizabeth Spelke. Domain-specific knowledge and conceptual change. *Mapping the mind: Domain specificity in cognition and culture*, pages 169–200, 1994.

[2] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.

[3] Zachary Estes, Sabrina Golonka, and Lara L Jones. 8 thematic thinking: The apprehension and consequences of thematic relations. *Psychology of Learning and Motivation-Advances in Research and Theory*, 54:249, 2011.

[4] Anna V Fisher. Processing of perceptual information is more robust than processing of conceptual information in preschool-age children: Evidence from costs of switching. *Cognition*, 119(2):253–264, 2011.

[5] Susan A Gelman and Ellen M Markman. Categories and induction in young children. *Cognition*, 23(3):183–209, 1986.

[6] Frank C Keil. *Concepts, kinds, and cognitive development*. mit Press, 1992.

[7] Ken McRae, George S Cree, Mark S Seidenberg, and Chris McNorgan. Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, 37(4):547–559, 2005.

[8] Daniel Mirman, Grant M Walker, and Kristen M Graziano. A tale of two semantic systems: taxonomic and thematic knowledge. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, pages 2211–2216. Cognitive Science Society, 2011.

[9] Gregory L Murphy. Causes of taxonomic sorting by adults: A test of the thematic-to-taxonomic shift. *Psychonomic Bulletin & Review*, 8(4):834–839, 2001.

[10] Paul C Quinn and Peter D Eimas. The emergence of category representations during infancy: Are separate perceptual and conceptual processes required? *Journal of Cognition and development*, 1(1):55–61, 2000.

[11] Paul C Quinn, Peter D Eimas, and Stacey L Rosenkrantz. Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, 22:463–463, 1993.

[12] Paul C Quinn and Mark H Johnson. The emergence of perceptual category representations in young infants: A connectionist analysis. *Journal of Experimental Child Psychology*, 66(2):236–263, 1997.

[13] Timothy T Rogers and James L McClelland. *Semantic cognition: A parallel distributed processing approach*. MIT press, 2004.

[14] Anna C Schapiro, Timothy T Rogers, Natalia I Cordova, Nicholas B Turk-Browne, and Matthew M Botvinick. Neural representations of events arise from temporal community structure. *Nature Neuroscience*, 16(4):486–492, 2013.

[15] Vladimir M Sloutsky and Anna V Fisher. The development of categorization. *Psychology of learning and motivation*, 54:141–166, 2011.

[16] Sandra S Smiley and Ann L Brown. Conceptual preference for thematic or taxonomic relations: A nonmonotonic age trend from preschool to old age. *Journal of Experimental Child Psychology*, 28(2):249–257, 1979.

Acknowledgments

Special thanks to Anna Fisher, Charles Kemp, David Rakison, & Layla Unger for their helpful feedback.
The project described was supported by Award Number T32GM081760 from the National Institute of General Medical Sciences. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of General Medical Sciences or the National Institutes of Health.

CogSci 2015, 37th Annual Cognitive Science Society Meeting, July 23-25, Pasadena, CA, USA