

When a Word is Worth a Thousand Pictures: A Connectionist Account of the Percept to Label Shift in Children's Reasoning

Jonathan Loose and Denis Mareschal

Department of Psychology
University of Exeter, Perry Rd, Exeter,
EX4 5QR, UK.

{ j.j.loose,d.mareschal}@exeter.ac.uk

Abstract

We present a connectionist model of children's developing reliance on object labels as opposed to superficial appearance when making inductive inferences. The model learns to infer a fact about an object based on the object's label (and not percept) even though that fact *has never been previously associated with* the label. The shift in reliance from perceptual to label information is found to depend on: (a) the presence of a pre-linguistic ability to categorize perceptual information, and (b) the greater variability of percepts than labels. The model predicts that children will shift their inductive basis at different ages depending on the perceptual variability of the test categories. This prediction is discussed with respect to studies of children's induction and with particular reference to conflicting results reported in the literature concerning the onset of label use.

Introduction

This paper presents a connectionist model of the child's developing reliance on object labels as opposed to appearance when making inductive inferences. Early studies of categorization and induction in young children have suggested that judgements regarding an object's category membership, or the likelihood of its sharing a property with another object, are made on a different basis depending on the child's age. Younger children apply a new fact to perceptually similar objects, whilst older children or adults utilise more profound conceptual information. This account of the perceptual/conceptual shift can be found throughout Piaget's work. The younger child is perceptually bound, and only after entering a subsequent stage of development can the child utilise abstract, categorical information (Inhelder & Piaget, 1964).

Since Piaget, this view of the perceptual/conceptual shift has been undermined as younger and younger children have been shown to behave in ways that Piaget would not have expected. Counter examples to Piaget are the induction studies discussed below (Gelman & Markman, 1986; Carey, 1985; Keil, 1989). These authors have shown that children as young as three and a half years old are able to make inferences that Piaget would have considered to be characteristic of much older children.

The fact that children's induction behavior changes earlier than Piaget expected has been used as support for the view that conceptual development is fundamentally a formal process of the growth of theory-like structures (Gelman &

Markman, 1986). Thus, development is taken to consist of changes in symbolic structures, and not the emergence of the *ability to use* such structures (Carey, 1985; Murphy & Medin, 1985). Whilst concepts are able to develop, the ability to modify abstract, symbolic representations is seen as a prerequisite throughout development.

Although Gelman and Markman (1986) found that young children rely on object labels as opposed to perceptual information when making inductions about natural kinds, studies with children *younger* their subjects are more equivocal. For example, McCarrell and Callanan's two year old subjects found *perceptual* information to be a more robust basis for induction than object labels (McCarrell & Callanan, 1995). Other induction studies using different (though still natural kind) stimuli have found a shift in the basis of induction across development, but that the onset of this shift depends on the concepts under investigation (Keil, 1989).

To understand these diverse findings, we must know why younger children are sometimes observed as being more perceptually reliant. Either they have difficulty comprehending label information, or (more likely) they simply value perceptual information more than verbal information in certain contexts and with certain concepts (Freeman & Sera, 1996). Unfortunately this account does not answer the question of how and why a child's bias might *change*. Thus, we are left with two issues. First, how might we account for the changing emphasis on perceptual and symbolic information? And secondly, must such an account presuppose an ability to manipulate abstract symbolic information?

We developed a connectionist model to address the issue of transition. The model makes spontaneous inferences, and changes the basis of its inferences over the course of learning from a reliance on perceptual information to a reliance on object labels.

The rest of this paper unfolds as follows. First, we present in greater detail the inductive reasoning paradigm used to explore children's conceptual knowledge. This paradigm will form the basis of the training regime used to model the perceptual / label shift. Next, we present the model architecture and training regime. The model's performance is then reported in terms of its behavior when (a) the percept and label information do not conflict, and (b) the percept and label information are in conflict. The next section explores how the variability of the perceptual category impacts on the

onset of a percept to label shift. The effect of multiple categories is also addressed. Finally, the results are discussed in terms of the implications for theories of children's concept acquisition.

The Inductive Reasoning Paradigm

The model presented in this paper focuses on an experimental paradigm often used in exploring the development of inductive inference (Gelman & Markman, 1986; Freeman & Sera, 1996). The paradigm is best explained with reference to Gelman and Markman (1986).

Gelman and Markman presented four year olds and adults with pairs of pictures of natural kind objects. Underneath each picture was written a fact about the object. The example given is the presentation of a flamingo, and a bat. Written under the picture of the flamingo was, "This bird's heart has a right aortic arch," whilst under the bat was written "this bat's heart has a left aortic arch." Having thus labeled the objects and provided a fact about them, a third picture is presented. The third picture is of an object which is perceptually similar to one of the objects already presented, but shares its label with the other. Thus, in the flamingo / bat example, the final picture is of a blackbird, labelled 'bird'. Perceptual and label information are in conflict as predictors of the fact. The child must tell the experimenter which fact is true of the final object (blackbird), and in so doing reveal whether he or she is relying on perceptual similarity or shared identity in making such inferences.

The Model

A successful model of this paradigm should begin by demonstrating a reliance on perceptual information then, at the end of training, demonstrate a reliance on label information instead. We start from the assumption that the child is able to perceptually categorise objects before being able to label them. It is now clear that even very young pre-linguistic children (infants) have impressive perceptual categorisation abilities (Quinn & Eimas, 1996). To capture this, we begin our simulations at a point where perceptual categorization abilities are already present (see Quinn & Johnson, 1996 for an example of how these early abilities can be modelled in connectionist networks). In contrast to many contemporary accounts of children's inductive reasoning (Gelman & Markman, 1986; Carey, 1985), our model does not presuppose any prior abstract or symbolic processing system.

Figure 1 shows a schematic outline of the model. It consists of a fully connected back propagation network with a single hidden layer. The input and output layers have been split into 3 banks representing perceptual information, label information, and a fact associated with an object. Note that figures 2-4 present simulation results using a diagrammatic form equivalent to figure 1, but without a representation of the hidden units. A similar approach has been used to model early language acquisition. (Chauvin, 1989; Plunkett, Marchman, Moller & Stransby, 1992). This is a model of cognitive processing and we would not necessarily equate the input to raw perceptual uptake and the output to observable behavior.




Figure 1 Schema of the model architecture

There are 48 units in the percept and prototype banks, and 16 units in each of the label and fact banks. The number of hidden units used depends on the number of patterns to be learned, and is kept as low as possible. In the initial example, only a single hidden unit is required. Learning rate and momentum are set to 0.25 and 0.9 respectively. Unit activation varies within the range 0—1 and does not decay. Weights are unbounded, and are initialised to random values in the range ± 0.5 .

The perceptual prototypes were defined as binary vectors, constrained to be greater than a certain fixed distance apart (i.e. minimum hamming distance = 36) with 50% of all values set to '1'. The percept inputs were generated on the fly by adding Gaussian noise (mean = 0, variance = 0.6) to each component of a prototype vector.

Different phases of training were used to model the different steps in the inductive reasoning paradigm described previously. First, the network is trained *only* on the percept to prototype association. This phase of training stops when the network is able to produce appropriate prototypes regularly with an error less than a fixed criterion (SSE = 0.6). The training is meant to capture the fact that the children in these studies arrive with the ability to categorise complex perceptual information. Note that during this phase, there is no training on either the label or fact pathways. The label and fact inputs are set to 0.5. Since activation ranges from 0-1, a value of 0.5 provides the network with no information.

The second phase involves training the network on the percept-and-label to prototype-and-label association. During this phase there is no training on the fact pathway. This phase of training is meant to capture the fact that older children have had more opportunity to learn which labels go with which percepts. Thus minimal training in this phase

(≈ 1000 epochs) represents a child early in development, whilst a much longer period of training ($\approx 30,000$ epochs) represents a child later in development.

The third phase involves training the percept-and-fact to prototype-and-fact associations, but *not* the label. This time, it is the label inputs that are set to 0.5 (no information). Weight changes do not occur in that pathway. This is meant to capture the experiment itself, in which the child, coming to the task with the ability to perceptually categorize plus some ability with label information, is taught to associate a fact with a particular object. The amount of training required at this point is significantly less than that required to model the development of percept/label associations (typically ≈ 100 epochs).

Note that the networks are *never* trained on the label to fact association. This association emerges by virtue of learning about percepts + labels, and percepts + facts *separately*. The network's inductive ability (i.e. its ability to produce a fact when presented with perceptual and label information) is probed by presenting both percept and label inputs and observing the fact response at the output. Since the networks have always experienced the fact in association with the percept and never in association with the label, one might expect the fact "inductive inference" to be driven by the percept input. As will be shown in the results, this is not always the case. We report on 3 types of testing. The first involves presenting the network with a new percept only. The second involves presenting the network with a label only. These tests are intended to model the first statements of the experimental paradigm in which a new exemplar of a familiar category (e.g. bird) is presented to the child. This tests whether the network is able to make "inductive inferences" based on either the percept or the label. The third test involves presenting the network with conflicting percept and label input and observing the fact produced at the output. This test provides a measure of the basis of the "inference"—i.e. a percept or label driven "inference". This test corresponds to the third statement in the paradigm used to evaluate the inferential basis of children.

Results

Percept only Inferences

Figure 2 shows the network's performance when presented with perceptual information alone. The left side of the figure shows performance early in development whereas the right side shows performance late in development. Each of the four sections of the figure shows the activation states of the inputs and outputs when a particular percept is presented to the network. The top row corresponds to performance when percept (A) is presented and the bottom row corresponds to performance when percept (B) is presented. As in figure 1, unit activations are represented by small squares. Dark squares correspond to high activations and light squares correspond to low activation. When no input is presented the squares are an intermediate shade,

signifying an activation of 0.5. In this condition neither the label nor the fact inputs are presented to the network.

Figure 1 demonstrates that both early and late in development, the network can produce the appropriate label (α) and fact (α) when presented only with percept (A), as well as label (β) and fact (β) when presented only with percept (B). The networks successfully infer label and fact information from percepts only.

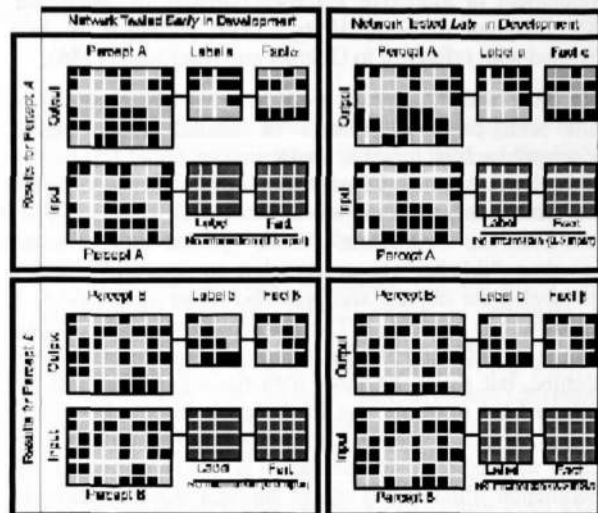


Figure 2 Input/Output mappings demonstrating correct induction performance when given percept information alone. Throughout development, a perceptual input leads to recognition of the appropriate prototype, label and fact.

Label Only Inferences

Figure 3 shows the network's performance when presented with labels only at the inputs. The top row shows performance when presented with label (a) and the bottom row when presented with label (b). Both early and late in development, the network produces the appropriate prototype and fact from the label only. Presentation of label (b) results in prototype (B) and fact (β) while presentation of label (a) results in prototype (A) and fact (α).

The fact that presenting any known percept *or* label produces the appropriate prototype, label and fact output reveals how the model is working. The model only has a *single* internal representation that is required to encode *multiple* mappings. It does not have separate representations of what an object looks like, what it is called, and what is true about it. When learning to categorize perceptually, the model develops a representation of distinct objects on the basis of perceptual information. When subsequently trained on label and fact information, this new information is incorporated into the existing internal representation. The result is an internal representation that is constrained by both the perceptual and label information.




Figure 3 Input / Output mappings demonstrating spontaneous induction when given label information. Throughout development, a label input leads to recognition of the appropriate prototype, label and fact.

Percept and Label Conflict Inferences

Figure 4 shows the crucial conflicting percept and label condition. The top of the figure shows percept (A) presented alongside label (b) while the bottom of the figure shows percept (B) presented alongside label (a). The principle guiding the network's induction in the conflict condition shifts over the course of development from a reliance on perceptual information to a reliance on label information.

Early in development, presentation of percept (A) and label (b) results in the production of prototype (A), label (a) and fact (α). The presentation of percept (B) and label (a) results in the production of prototype (B), label (b) and fact (β). In both cases, the network is relying on the percept information only to infer the object category, label, and fact. The reader may wish to compare with figures 2 and 3 to verify what the expected (A—a— α) and (B—b— β) percept-label-fact associations are.

A very different thing happens later in development. The basis on which the network makes an inference has shifted. When presented with percept (A) and label (b), the network now produces prototype (B), label (b) and fact (β). Similarly, when presented with percept (B) and label (a) the network produces prototype (A), label (a) and fact (α). Again, the reader may wish to verify this against performance depicted in figures 2 and 3.

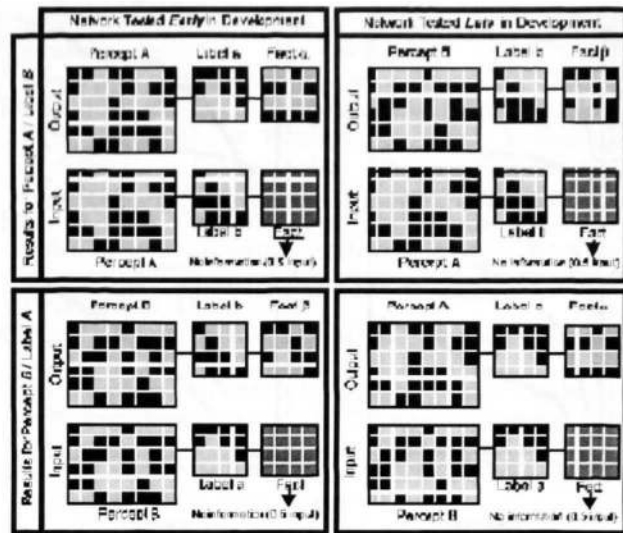


Figure 4 Input / Output mappings demonstrating a shift in the basis of induction when presented with conflicting percept / label information. Early in development the output is determined by the perceptual input, but later on, the output is determined by the category label.

This performance can be explained as a combination of two factors. First, the single hidden unit bank used to map multiple sets of information. This fact implies that a single internal representation will be developed for any prototype/label/fact combination, and all outputs will be triggered by the same internal (hidden unit) representation. Second, perceptual input rarely corresponds perfectly to the prototype, whilst labels are invariant. As the network seeks to reduce error, it will inevitably discover that the best trigger of the appropriate internal representation is the label input. Thus the weights between the perceptual input and the hidden representation should gradually reduce in magnitude whilst the weights between the label input and the hidden representation should gradually increase during development.

The Role of Perceptual Variability

Variability across percepts seems to lead to a shift in reliance from percept to label when performing induction. We constructed a smaller network to test the role of variable percepts more precisely. This network had 12 units for the percept / prototype representation, and 2 for each of the label and fact representations. The network showed the same shift in behaviour as the larger network when given the same level of noise (Noise taken from a Gaussian distribution, $s^2=0.6$).




Figure 5 Plots of average absolute weights from percept input to hidden layer (thin), and label input to hidden layer (thick) averaged over 5 simulations

The results of the previous section suggest that a reliance on percept or label depends on which set of units most effectively triggers the appropriate internal representation. The reliance on percept vs. label can be measured by comparing the average absolute size of the weights between the hidden layer and the percept inputs with the average absolute size of the weights between the hidden layer and the label inputs. Figure 5 shows this as a function of epochs for different levels of variance.

The two prototypes used in these simulations were the inverses of each other (i.e. hamming distance between patterns = 12). It is worth noting that the average of the label weights grows to exceed the average of the perceptual weights for variances exceeding 0.5. At this point, the level of variability in percepts means that some will actually be more similar to the opposite prototype. The shift is not driven by percepts that are slightly different from the prototype, but by percepts that are more similar to *another* prototype than to their own category prototype. A second

point to note is that as the variance level increases, so the earlier the network shifts to label inferences. This is less obvious at very high variability levels, since as noise increases the number of epochs required for learning perceptual categorization becomes less predictable and also tends to increase. Hence, the model makes the strong prediction that a shift from perceptual reliance to label reliance will occur earlier for perceptually similar objects with different labels.

The Effect of Multiple Prototypes

This section briefly focuses on the effects of *multiple* prototype/label mappings. Although the Gelman and Markman studies focused on the use of two percepts and facts, in the real world children have knowledge of many more. All of these could interact with performance in the simplified inductive reasoning task. To explore possible interaction effects, we trained the full scale model (64 percept units, 16 label units, 16 fact units) with 10 prototypes each associated with its own fact and label. To accommodate this number of patterns, 4 hidden units were used. The prototype patterns are all greater than a certain fixed (hamming) distance apart. Two important findings emerge from this work. First, the shift from percept reliance to object label reliance in induction does not occur at the same point in training for all objects. Second, objects which are perceptually more similar tend to shift earlier. The reason for this is that the network will tend to shift patterns that reduce more of the error earlier in training. The implication of this for child development research is that perceptually more similar patterns shift at a younger age.

Discussion

We present a connectionist model of the development of a reliance on label information as opposed to perceptual information when making inductive inferences. This model performs the shift without any ability to manipulate formal systems as is the suggestion of the 'concepts as theories' view. Evidence that younger children rely on labels is often taken as support for this view, however we show that the development of label reliance may be seen as a product of a confusing perceptual environment, as opposed to an innate need for structure. We would not want to deny that such structures do develop. Rather, we suggest that they emerge from simple systems such as the one described here. This work may be viewed within the context of the growing body of work utilising connectionist principles to develop a new understanding of ideas critical to developmental psychology - particularly what it means for a behavior to be either learned or innate (e.g., Elman, Bates, Karmiloff-Smith, Johnson, Parisi, & Plunkett, 1996). Our model makes architectural assumptions which may reflect innate structure - however such structure would be of a very different kind from what might be proposed outside a connectionist framework.

Our model predicts that differently labelled objects that are nevertheless perceptually confusable will promote induction on the basis of label earlier than objects that are

easier to distinguish perceptually. This prediction may shed some light on the contrasting results of different induction paradigms. The original studies by Gelman et al (Gelman & Markman, 1986; Gelman & Markman, 1987) did not find significant evidence of a percept / label shift in subjects as young as 3.5 years old. Given their stimuli this result is not surprising. Their studies rely on perceptually confusable natural kind stimuli (for example, blackbirds and bats). Our model predicts that these are the very cases that are most likely to promote reasoning on the basis of labels early on.

Some studies do not rely on perceptually confusable examples. Keil (1989) asked children older than Gelman's subjects to say whether a Skunk that had been surgically transformed to look exactly like a Raccoon was in fact a Raccoon or not. In these studies, Keil found that young children were *still* reliant on perceptual information in making inductions, and only later did they realise that the truly salient characteristic for induction was the animal's label. This is what our model would predict—the basis of induction will take longer to shift when perceptual categorization is more clear cut.

However, other studies involving still other stimuli do not report the same findings (Keil, 1989). Even young children rely on object kind (label) as opposed to perceptual similarity when presented with a toy dog and a real dog. Note, however, that whilst in our study all stimuli were presented equally often to the network, this is not true of children's learning. We might propose that toy dogs are *precisely* the kinds of objects that young children would have had a great deal of experience of - and thus the basis of induction with such objects would be more likely to have shifted. Our work with multiple prototypes demonstrates that the basis of induction for different objects will shift at different times. An alternative simple account of this is to say that the stuffed dogs and real dogs are perceptually similar and, therefore, will be differentiated by label use early in development. Again, this is compatible with the behaviour of our model.

The question arises as to whether this model might also be applicable to novice/expert shifts in adult learning (e.g. Slotta, Chi, & Joram, 1995). The data suggests that this may be the case. In future, however, we intend to extend the model to cover a broader spectrum of the child data in detail. We expect the analogy between adult learning and child development to be strained by such progress.

The principles embodied in this model are inadequate to fully explain the richness of children's induction. However, the model demonstrates that simple mechanisms can cover broad sections of data without requiring complex, structured internal representations, and provides support for the view that the genesis of conceptual thought need not require a fully developed representational system.

References

Carey, S. (1985) *Conceptual Change in Childhood*. Cambridge, MA, MIT Press.

- Chauvin, Y. (1989) Toward a connectionist model of symbolic emergence. In *Proceedings of the 11th Conference of the Cognitive Science Society*, p. 580-587.
- Elman, J.L., Bates, E.A., Karmiloff-Smith, A., Johnson, M.H., Parisi, D., & Plunkett, K. (1996); *Rethinking Innateness: Connectionism in a developmental framework*. Cambridge, MA., MIT Press.
- Freeman, K. & Sera, M. (1996) Reliance on visual and verbal information across ontological kinds: What do children know about animals and machines? *Cognitive Development*, **11**, 315-341.
- Gelman, S. & Markman, E. (1986) Categories and Induction in Young Children. *Cognition*, **23**, 183-209.
- Gelman, S & Markman, E. (1987) Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*, **57**, 1532-1541.
- Inhelder, B. & Piaget, J. (1964) *The Early Growth of Logic in the Child*. New York, Norton.
- Keil, F (1989) *Concepts, Kinds and Cognitive Development*. Cambridge, Cambridge, MA, MIT Press.
- McCarrell, N. & Callanan, M. (1995) Form-function correspondences in children's inference. *Child Development*, **66**(2), 532-546.
- Murphy, G. & Medin, D (1985) The Role of Theories in Conceptual Coherence, *Psychological Review*, **92**, 289-316.
- Plunkett, K., Sinha, C., Moller, M., and Strandsby, O. (1992) Symbol grounding or the emergence of symbols? Vocabulary growth in children and a connectionist net. *Connection Science*, **4**:293-312.
- Quinn, P. and Eimas, P. (1996) Perceptual organization and categorization in young infants. In Rovee-Collier, C. and L.P.Lipsitt, editors, *Advances in Infancy Research*, volume 10, p. 1-36. Norwoods, NJ, Ablex.
- Quinn, P. and Johnson, M. (1996) The emergence of perceptual category representations during early development: A connectionist analysis. In Cottrell, G., editor, *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, p. 638-643, Hillsdale, NJ:LEA.
- Slotta, J.D., Chi, M.T.H., Joram, E. (1995) Assessing students' misclassifications of physics concepts: An ontological basis for conceptual change. *Cognition and Instruction*, **13**(3), 373-400.