

Experimental and Connectionist Perspectives on Semantic Memory Development

Samantha J. Hartley (pcp95sjh@sheffield.ac.uk)
Department of Psychology, University of Sheffield,
Western Bank, Sheffield S10 2TP, England

Tony J. Prescott (t.j-prescott@sheffield.ac.uk)
Department of Psychology, University of Sheffield.

Roderick I. Nicolson (r.nicolson@sheffield.ac.uk)
Department of Psychology, University of Sheffield.

Abstract

We describe an experimental investigation of the development of children's knowledge structures which aims to provide data for connectionist modelling. 167 children between 5 and 11 years of age completed two category fluency tasks where they were asked to produce as many names of a) animals and b) parts of the body, as they could in one minute. Similarity scores were derived based on distances between concepts in the lists produced. These were analysed using the ADDtree algorithm (Sattath & Tversky, 1977) to build structures representing the organisation of the children's knowledge of animals and body parts. The results showed that animal knowledge was generally organised in terms of environmental context/habitat, however, there was evidence for subtle changes in knowledge organisation between age groups. More pronounced changes were observed in the organisation of knowledge of body parts which gave some support to the assertion that children progress from making coarser to making finer distinctions between concepts (see Keil, 1979) and reflected the progression observed in knowledge structure development in a connectionist model of semantic memory discussed by McClelland, McNaughton and O'Reilly (1995). Our aim is to extend this work to provide data enabling connectionist modelling of semantic memory within a developmental framework.

Introduction—Connectionist Modelling of Semantic Memory Development

Recent years have seen much interest in the use of connectionist modelling to investigate the structure and function of semantic memory (e.g. Hinton & Shallice, 1991; McClelland, McNaughton & O'Reilly, 1995; McRae, de Sa & Seidenberg, 1997; Plaut, 1995; Rumelhart, 1990; Rumelhart & Todd, 1992). This work has been based on knowledge of adult semantic memory, with training sets derived from adult data, and has investigated the organisation and functioning of the established semantic system and the effects of brain damage on this system (e.g. Farah & McClelland, 1991; Hinton & Shallice, 1991; Plaut & Shallice, 1986). This is valuable, insofar as it aids

understanding of an adult's semantic system, however, it generally sheds little light on how this system develops in children. One exception is the work of McClelland et al. (1995) who draw parallels between the development of an internal knowledge structure in their model (a replication of Rumelhart & Todd, 1992) and observations from developmental psychology (see later). These parallels are drawn post hoc, however, as this model was not designed or trained with the express intention of investigating development in semantic memory. The value of connectionist modelling techniques in investigating developmental processes is well recognised (Elman et al., 1996). It is clear that the study of semantic memory with connectionist models can benefit from a developmental perspective using data from work on child memory for network training and behaviour comparison.

A developmental approach to semantic memory modelling should also improve our understanding of the effect of a child's knowledge base on learning processes and cognitive performance (see Bjorklund, 1987; Schneider, Korkel & Weinert, 1989; Schneider, Korkel & Weinert, 1990). As Grube and Hasselhorn (1996) point out, although the impact of changes in the knowledge base are well documented, little is understood about how exactly such changes exert their effect. One possibility is that quantitative increases in a child's knowledge base lead to qualitative changes in the structure of internal knowledge representations, in turn affecting the child's performance on tasks tapping this knowledge. There is much evidence for such changes between 4 and 10 years of age. Children's conceptual representations undergo shifts from being concrete to more abstract between these ages (e.g. Keil, 1989; Vygotsky, 1962). Carey (1985) observed marked changes in the inductive inferences children will make about biological kinds, while Keil (1979) demonstrated a gradual change from children making coarser to their making finer distinctions between concepts, reflected in their ability to correctly decide which terms (or features) could be attributed to certain concepts. It was to the latter progression that McClelland et al. (1995) drew parallels with the development of conceptual knowledge structures in their connectionist model. Figure 1

illustrates our own replication of this study which used the same the network architecture and training schedule as reported by McClelland et al. To aid comparison with data from developmental studies reported later in the paper, an ADDtree analysis (Sattath & Tversky, 1977) is shown here for the relationships between concept representations formed in the hidden unit weights of the network. As illustrated in the figure there is a clear progression from coarser to finer distinctions between concepts as training proceeds. This is indicated by the distinct clusters of animals and plants which are not present at 50 epochs, develop by 100 epochs, and are most pronounced at 200 epochs. By the final stages of training the conceptual structure is also further refined with birds, fish, trees, and flowers clearly differentiated.

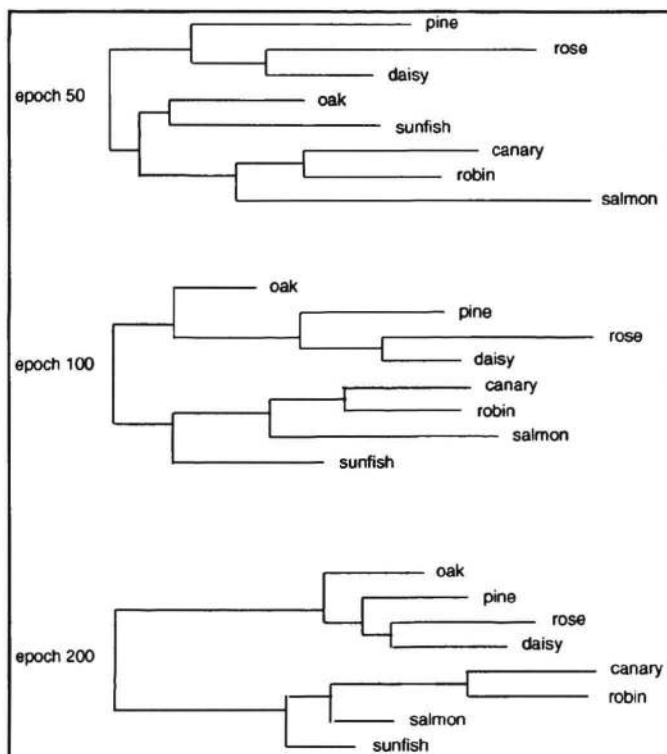


Figure 1: ADDtree analysis of hidden unit weights in a connectionist model of semantic memory. The distance between items is indicated by the length of the *horizontal* arcs in the path joining them. The length of the arc joining a cluster to the rest of the tree gives an indication of the distinctiveness of that cluster. The root is chosen to give the most balanced tree possible. See Sattath & Tversky (1977) for a full description of ADDtrees, and McClelland et al. (1995) for description of network architecture and training schedule.

The method used to judge concept similarity is quite different, however, in this connectionist model (comparing internal distributed representations) from that used in Keil's (1979) developmental study (asking children to judge whether a feature applies to a concept, e.g. 'can milk be alive?', then comparing feature attributions). This renders

the comparison between the two studies somewhat indirect. Our implementation of the model was extended to include a measure of response latency (the cascade mechanism developed by McClelland, 1979) that allowed us to investigate semantic priming effects (for details see Hartley, in prep.). These experiments demonstrated that semantic priming in the model is particularly sensitive to the type/token distributions of conceptual and featural information in the training set— inappropriate distributions generate priming effects quite different from those described in the human priming literature. The data set used for training the network by McClelland et al. and ourselves was not influenced at all by distributions obtained from either adult or child data. These observations suggest that the current model can provide only limited insight into the operation and development of child semantic memory.

Our proposal is to investigate connectionist models of semantic memory within a more explicitly developmental framework—training models on data collected from studies with children, and evaluating them by direct comparison with behavioral measures. The main study described in this paper presents the results of a category fluency study illustrating the development of children's knowledge structures for the animals and body parts categories. The results of this study are currently being extended with a series of experiments on children's concept and feature knowledge in order to provide developmental data for future connectionist modelling.

The Category Fluency Task

The category (semantic) fluency task has been used in several studies investigating the development of knowledge organisation in children (e.g. Grube & Hasselhorn, 1996; Kail & Nippold, 1984; Storm, 1980). This procedure requires participants to generate as many exemplars of a category as possible in a given length of time. It may be assumed that when a word or concept is activated (generated), it will in turn activate other words or concepts which are semantically similar or associatively related to it. Such an assumption is well supported by evidence from numerous semantic priming studies (see Neely, 1991). The order in which words are recalled can therefore be assumed to indicate the psychological proximity of the items produced. Henley (1969) provided evidence to support this assumption by demonstrating that animals of close psychological proximity are named in closer proximity in lists.

Fluency studies have generally concentrated on the category of animals, as one which is familiar to most children. In addition to increased production frequency with age, the findings have typically shown a tendency for the children's responses to be clustered based on habitat/environmental context (e.g. those which live on farms, those from tropical places etc.) and for older children to produce larger clusters. Grube and Hasselhorn (1996), and Storm (1980) observed older children to use a greater number of strong inter-item associations (e.g. lion-tiger) in their

responses. These results certainly suggest quantitative changes in the child's knowledge base, but provide limited insight into changes in knowledge structures. Storm (1980) attempted to find evidence for developmental change in the structure of children's knowledge of animals, however, she was unable to show any substantial differences in organisation of concept clusters between ages 5 to 11 years.

The present study was stimulated by the need to provide data on the development of internal knowledge structures to be compared to that obtained from existing and future connectionist models. However, it was also motivated by the curious fact that Storm's (1980) study had failed to show any progression (from coarser to finer detail) in children akin to that described by McClelland et al. (1995), or that might be expected given Keil's (1979, 1989) findings. There are several reasons why this might have been the case.

First, Storm analysed concepts which a majority of children were shown to be able to name, in a separate picture naming task, rather than those which had occurred with high frequency during the fluency task. Since picture naming involves somewhat different memorial processes (visual recognition and word recall) than concept generation this could have an influence on the pattern of results obtained. The present study therefore focuses on the most frequent responses made during the category fluency task across age groups.

Second, knowledge of animals develops early in children (Clark, 1995), so it is possible that, by the time Storm tested her subjects, a developmental progression in conceptual knowledge had already taken place, even for the youngest children. Later changes may be somewhat subtle and therefore more difficult to detect. To address this issue, some researchers have attempted to assess the conceptual knowledge of children below school age and of infants (e.g. Gelman, 1988; Mandler, 1993; Oakes, Coppage & Dingel, 1997). Although the results of these studies have been of great value, there are recognised difficulties in assessing the breadth of a child's knowledge given the communication problems that arise with children of this age. The current study therefore took a different approach, which was to investigate a second category of conceptual knowledge, that of human body parts, which may continue to develop later in childhood (Meadows, 1993).

Third, the hierarchical cluster analysis method used by Storm imposes severe constraints, such that with disjoint clusters, all intra-cluster distances are smaller than inter-cluster distances, and all inter-cluster distances are equal. Similarity measures derived from category fluency data generally violate such constraints. The ADDtree algorithm (Sattath & Tversky, 1977) provides a more flexible clustering method allowing representation of different inter and intra-cluster distances, leading to a more faithful representation of the data. ADDtree analysis has been successfully employed to investigate conceptual structure in adult semantic memory (Chan et al., 1993; Sattath & Tversky, 1977). The present paper uses the ADDtree

algorithm and demonstrates how this technique may be used to look at changes in the structure of conceptual knowledge as indicated by performance in the fluency task.

Method

Participants 167 children from two schools, one in Sheffield, UK, the other in Cheshire, UK, were involved in the study. Three groups of children participated in school years 1 (equivalent to US first grade), 3, and 5. There were 60 children in the year 1 group, mean age 6:1 years, range 5:7 to 6:6; 54 in year 3, mean age 8:4, range 7:10 to 8:10; and 53 in year 5, mean age 10:1, range 9:3 to 10:9. All of the children spoke English as their first language.

Administration of the BPVS and Category Fluency Tasks The children were tested individually at school, in a quiet area. The short form British Picture Vocabulary Scale (BPVS) (Dunn, 1982) was administered to all children to assess receptive vocabulary. Each child then completed two tasks, which will be henceforth referred to as the 'animals task' and the 'body parts task'. Half of the children in each class were asked to complete the animals task first followed by the body parts task, the other half vice versa. For the animals task, the child was asked name as many different animals as they could think of in one minute, similarly for the body parts task, they were asked to name as many different parts of the body as they could in one minute.

Obtaining Similarity Scores The fluency task provided separate lists of animals and body parts produced by each child, in production order. Target concepts were chosen for analysis (13 animals and 12 body parts) on the basis that they occurred within the top twenty most frequently produced responses for all age groups. Six matrices (3 year groups x 2 categories) of distance measures between pairs of terms were derived using the following algorithm (modified from Chan et al, 1993), to control for differences in production frequency between terms and age groups:

$$D_{ij} = \frac{\bar{n}_k \bar{T}_{ij}}{T_{ij}^2} \sum_{k=1}^{T_{ij}} \frac{d_{ijk}}{n_k}$$

Here, for terms i and j , D_{ij} is the distance recorded in the matrix for the two terms, d_{ijk} is the measured distance between the terms for participant k (proportional to the number of intervening terms), T_{ij} is the total number of subjects who responses contained the pair i and j , and n_k is the total number of responses for subject k . Further details of this algorithm are given in Chan et al. (1993) and Hartley (in prep.).

Results and Discussion

Production frequency 1 factor ANOVA revealed highly significant overall increases with age in the total number of responses given for both the animals ($F(2,164)=67.731$;

$p < 0.0001$) and body parts ($F(2,164) = 62.204$; $p < 0.0001$) tasks (see table 1). The raw BPVS scores and production frequency showed a significant positive correlation ($p < 0.0001$) suggesting that the increase in production frequency is related to increased vocabulary size. In other words, the increase in production frequency suggests that the children's conceptual knowledge base is increasing over the age range investigated.

Representing category structure ADDtree analysis was applied to the distance measures to build a representation of the category structure of the children's knowledge of animals and body parts. The ADDtree program used was that written by Corter (1982). The resulting trees are shown in figures 2 and 3. Note that the signal-to-noise ratio is higher than for the artificial data shown in figure 1 (this is indicated by the generally shorter lengths of internal arcs in the ADDtrees). This is to be expected given the nature of the task, the use of child subjects, and the larger number of concepts analysed.

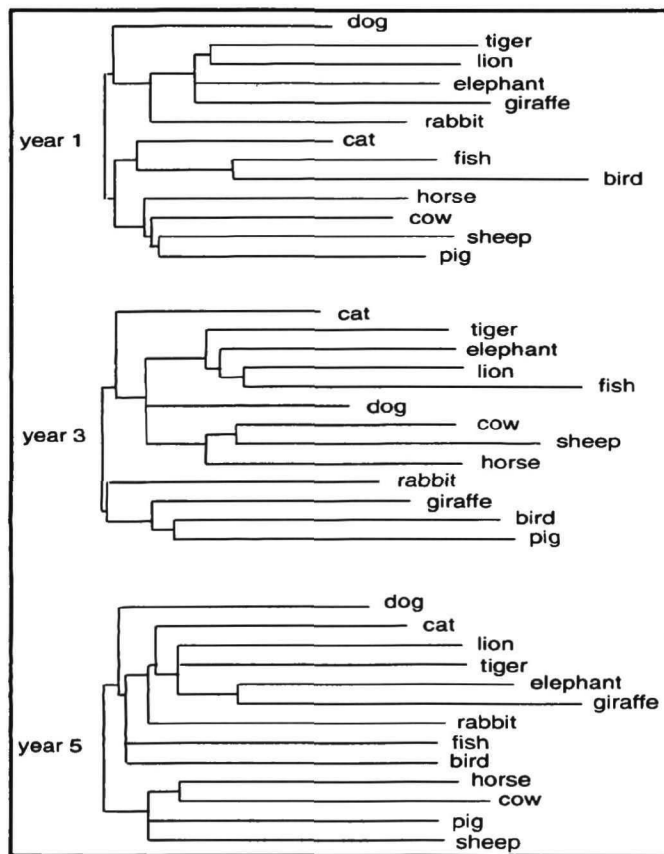


Figure 2: ADDtree showing differences in conceptual structure of animal knowledge for children in different age groups (years 1, 3, and 5 correspond to mean ages of approximately 6, 8 and 10 years).

In accordance with the results of Storm (1980), analysis of the animal data revealed major clusters of wild/exotic (tiger, lion, elephant, giraffe) and farm animals (horse, cow, sheep,

Table 1: Mean number of responses for category fluency tasks (standard deviations in brackets)

Year	Animals	Body Parts
1	9.200 (2.730)	9.450 (3.520)
3	13.037 (3.376)	12.574 (4.183)
5	16.377 (3.834)	17.981 (4.555)

pig) in years 1 and 5 and to a lesser degree in year 3. There was some evidence for change in the organisation of concepts, however, possibly reflecting the broadening of children's knowledge. For example, 'cat' is closer to the cluster of big cats (tiger and lion) within the wild animals cluster, for year 5 children, than for younger children. This may suggest a development of taxonomic associations in addition to those of habitat outlined in previous studies. In year 5, elephant and giraffe form a more distinct cluster possibly because, in addition to being exotic, they are recognised as being large animals. Analyses of adult data have often found size to be the second most significant factor after environment in determining conceptual relationships between animals (Chan et al., 1993; Henley, 1969;).

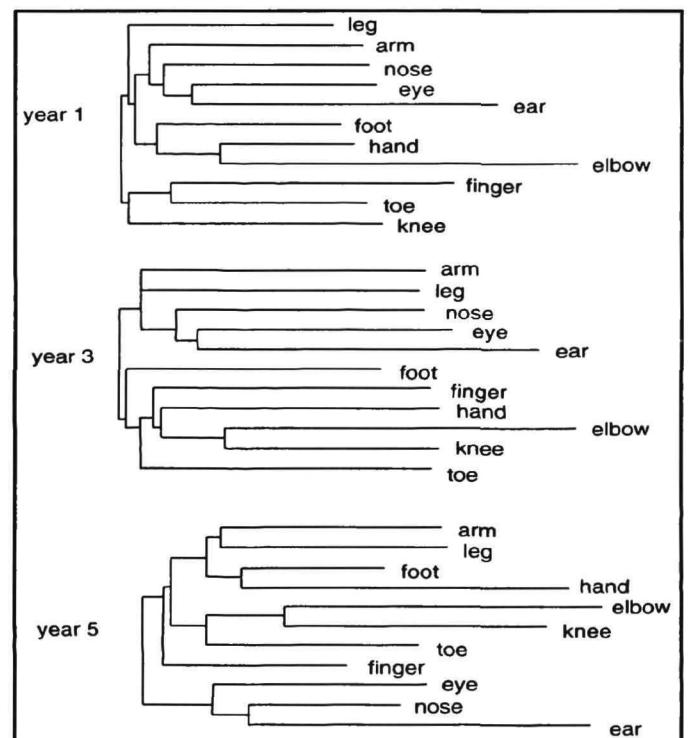


Figure 3: ADDtree showing differences in conceptual structure of body part knowledge for children in different age groups.

Progressive change in conceptual structure was perhaps more evident for body parts than for animals (see figure 3). The children in all years distinguished a cluster of sense organs (eye, ear and nose), however, this becomes

progressively more distinctive with age forming a well-defined major cluster in the year 5 group. The year 1 children showed additional sub-clusters of digits (toe, finger) and arm parts (hand, elbow), and the year 3 children a cluster of joints (knee, elbow), however, overall the year 5 group again shows the greatest differentiation of concepts with a large cluster of 'limbs and attachments' sub-dividing into clusters of limbs (arm and leg), attachments (foot and hand), and joints (elbow and knee).

General Discussion

The results provide strong support for observations in previous studies that children produce more exemplars with increasing age. Thus an important developmental trend that needs to be captured in the modelling work is the quantitative increase in conceptual knowledge with age.

With respect to the animal data the current study confirms previous reports (e.g. Corter, 1982; Grube & Hasselhorn, 1996; Kail & Nippold, 1984; Storm, 1980) that animal concepts are generally clustered in terms of habitat/environmental context. The results showed no clear evidence of a trend toward finer discrimination between concepts with age, however, some evidence of subtle developments in conceptual structures with age were found which could be followed up in future studies (for instance, the suggestion that older children may be more influenced by taxonomic category or may make finer discriminations on the basis of animal size).

This is the first study to look at children's semantic memory for body parts using the category fluency task. The results for this domain show some clustering of body parts with similar functional status in all three groups which is much better defined in year 5 where there are more pronounced clusters for sense organs, limbs, attachments, and joints. These changes provide some support for the notion of a progression from children making coarser to their making finer distinctions with age.

Bjorklund (1985) has argued that when different age groups have equal knowledge of a conceptual domain there will be no difference observed in performance on tasks which investigate knowledge structure. Since knowledge of animals develops early, it is likely that to a greater extent the knowledge of the older and younger children is similar. Although knowledge of some body parts is acquired at an early age more knowledge is acquired later during schooling (Meadows, 1993). This could explain why there is more evidence of conceptual change in the body parts data.

It is possible that using only those concepts named frequently by children in all three year groups yields results that do not do justice to the diversity of the older children's knowledge. Preliminary analyses of the top 20 responses from each year group has therefore also been carried out. These results illustrate further subtle changes within the categories of animals, and the development of structured knowledge of internal body parts (e.g. heart, lung, brain etc.). Category fluency data has also been collected from a

range of other semantic domains, and is being analysed for developmental changes in knowledge structures. We have also investigated children's explicit similarity judgements for animal concept pairs. ADDtree analysis of the results reveals an almost identical clustering pattern as that obtained from category fluency, showing our results to be robust with an independent measure of concept relatedness.

There is much evidence to support the assumption that the ordering of items produced in a category fluency task reflects internal knowledge structures (see earlier summary). It has been suggested, however, that changes in word generation behavior could reflect the development of metamemorial skills (Hasselhorn, 1992) rather than a change in the underlying knowledge. In other words, older children may use specific strategies to make their recall more efficient—e.g. they may purposefully decide to produce as many farm animals as they can, name parts of the face, etc. Bjorklund (1985) has investigated this possibility and suggests that although older children are able to use strategies to aid their recall, the use of these strategies depends upon, and is supported by, their conceptual knowledge structures. Our results from the similarity judgements study provide further support for this. It is unlikely then, that the changes in category knowledge noted for the category fluency data are due solely to the development of metamemorial skills.

The difficulties involved in interpreting empirical findings illustrate the need for models of semantic memory development. Initial comparisons between the current results and those of McClelland et al. are promising. Specifically, there are similarities between the developmental trajectories in the connectionist model and the body part data presented here (compare figures 1 and 3) both showing more distinct conceptual categories with training/age. Much more could be done, however, to bring the modelling work closer to the empirical studies. The present study constitutes part of an ongoing project which aims to provide developmental data for constructing future models. Theoretical accounts of semantic memory generally suppose that concept-relatedness derives from shared features or properties, and most modelling approaches begin with this premise. For example, in the McClelland et al. model internal concept representations are derived by training the network with concept/relation/feature-list triples (e.g. robin/isa/animal-bird-etc.). Therefore in addition to collecting concept norms and similarity judgements we are also investigating property norms and children's attributions of properties to concepts. These experiments should provide a further means for testing the hypothesis that children make finer distinctions between concepts with age, and will allow the construction of connectionist models that bear more direct comparison with developmental data.

References

- Bjorklund, D. F. (1985). The role of conceptual knowledge in the development of organisation in children's memory.

- In C. J. Brainerd (Ed.), *Basic Processes in Memory Development*. New York: Springer-Verlag.
- Bjorklund, D. F. (1987). How age changes in knowledge base contribute to the development of children's memory: An interpretative review. *Developmental Review*, *7*, 93-130.
- Carey, S. (1985). *Conceptual Change in Childhood*. Cambridge, MA: MIT Press
- Chan, A. S., Butters, N., Paulsen, J. S., Salmon, D. P., Swenson, M. R., & Maloney, L. T. (1993). An assessment of the semantic network in patients with Alzheimer's disease. *Journal of Cognitive Neuroscience*, *5*, 254-261.
- Clark, E. V. (1995). *The lexicon in acquisition*. Cambridge: Cambridge University Press.
- Cortner, J. E. (1982). Addtree/P: A Pascal program for fitting additive tree based on Sattath and Tversky's ADDTREE algorithm. *Behavior Research Methods and Instrumentation*, *14*, 353-354.
- Dunn, L.M. (1982). *The British Picture Vocabulary Scale*. Windsor: NFER-Nelson.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking Innateness: A Connectionist Perspective on Development*. MIT Press.
- Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of experimental psychology: General*, *120*, 339-357.
- Gelman, S. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology*, *20*, 65-95.
- Grube, D., & Hasselhorn, M. (1996). Children's freelisting of animal terms: developmental changes in activating categorical knowledge. *Zeitschrift fur psychologie*, *204*, 119-134.
- Hartley, S. J. (in preparation.). *The development of conceptual knowledge: connectionist and experimental insights*. Unpublished PhD thesis, University of Sheffield.
- Hasselhorn, M. (1992). Task dependancy and the role of category typicality and metamemory in the development of an organisational strategy. *Child Development*, *63*, 202-214.
- Henley, N. M. (1969). A psychological study of the semantics of animal terms. *Journal of Verbal Learning and Verbal Behaviour*, *8*, 176-184.
- Hinton, G. E., & Shallice, T. (1991). Lesioning a connectionist network: Investigations of acquired dyslexia. *Psychological Review*, *98*, 74-95.
- Kail, R., & Nippold, M. A. (1984). Unconstrained retrieval from semantic memory. *Child Development*, *55*, 944-951.
- Keil, F. C. (1979). *Semantic and Conceptual Development: An ontological perspective*. (first edition. ed.). London: Harvard University Press.
- Keil, F. C. (1989). *Concepts, kinds and cognitive development*. Cambridge, MA: MIT Press.
- Mandler, J. M. (1993). On concepts. *Cognitive development*, *8*, 141-148.
- McClelland, J. L. (1979). On the time-relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287-330.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419-457.
- McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*(2), 99-130.
- Meadows, S. (1993). *The child as thinker: The development and acquisition of cognition in childhood*. London: Routledge.
- Neely, J. H. (Ed.). (1991). *Semantic priming effects in visual word recognition: A selective review of current findings and theories*. Hillsdale, N.J: Lawrence Erlbaum Associates.
- Oakes, L. M., Coppage, D. J., & Dingel, A. (1997). By land or by sea: The role of perceptual similarity in infants' categorisation of animals. *Developmental Psychology*, *33*(3), 396-407.
- Plaut, D. C. (1995,). *Semantic and associative priming in a distributed attractor network*. Paper presented at the 17th Annual Conference of the Cognitive Science Society.
- Plaut, D. C., & Shallice, T. (1986). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, *10*(5), 377-500.
- Rumelhart, D. E. (Ed.). (1990). *Brain style computation: Learning and generalisation*. San Diego, CA: Academic Press.
- Rumelhart, D. E., & Todd, P. M. (Eds.). (1992). *Learning and connectionist representations*. Cambridge: MA: MIT Press.
- Sattath, S., & Tversky, A. (1977). Additive similarity trees. *Psychometrika*, *42*(3), 319-345.
- Schneider, W., Korkel, J., & Weinert, F. E. (1989). Domain specific knowledge and memory performance: A comparison of high and low aptitude children. *Journal of Educational Psychology*, *81*, 306-312.
- Schneider, W., Korkel, J., & Weinert, F. E. (1990). Expert knowledge, general abilities, and text processing. In W. Schneider & F. E. Weinert (Eds.), *Interactions among aptitudes, strategies and knowledge in cognitive performance*(pp. 235-251). New York: Springer.
- Storm, C. (1980). The semantic structure of animal terms: A developmental study. *International Journal of Behavioral Development*, *3*, 381-407.
- Vygotsky, L. S. (1962). *Thought and Language*. Cambridge, M.A.: MIT Press.