Towards a Model of Grounded Concept Formation*

Stefan Wrobel
GMD (German Natl. Research Center for Computer Science)
F3.XPS, Pf. 1240, W-5205 St. Augustin 1, Fed. Rep. Germany
wrobel@gmdzi.gmd.de

Abstract

In most research on concept formation within machine learning and cognitive psychology, the features from which concepts are built are assumed to be provided as elementary vocabulary. In this paper, we argue that this is an unnecessarily limited paradigm within which to examine concept formation. Based on evidence from psychology and machine learning, we contend that a principled account of the origin of features can only be given with a *grounded* model of concept formation, i.e., with a model that incorporates direct access to the world via sensors and manipulators. We discuss the domain of process control as a suitable framework for research into such models, and present a first approach to the problem of developing elementary vocabularies from perceptual sensor data.

1 Introduction

The goal of human concept formation is to arrive at a conceptual system that partitions the encountered objects and events in a way that enables us to effectively deal with our environment. As such, the task of concept formation can be split up into two subtasks, concept *aggregation* and concept *characterization.* Concept aggregation is the decision about which entities are to be grouped together into a concept, concept characterization means finding an intensional description of the proposed concept based on its extension. The latter task is often called *concept learning* (from examples).

Consequently, what separates concept formation from concept learning is the difficult question of deciding which objects to aggregate, i.e., how to carve up the world into different concepts. Most operational models of concept formation, most notably *conceptual clustering* systems [Kolodner, 1983; Lebowitz, 1987; Fisher, 1987; Gennari *et al.*, 1989] base their answer to this question on an assumption that was nicely formulated by Rosch and her colleagues [1976; 1978]: The "correlated feature" view assumes that features in the world occur as clusters, and that the best concepts are those that maximize intra-concept correlations, and minimize inter-concept correlations, i.e. reflect the presumed cluster structure of the world. There is a significant amount of empirical evidence that people detect and use feature correlations [Medin *et al.*, 1982; Younger and Cohen, 1984], and the existing clustering systems based on the correlated feature principle and extensions thereof[1] have been very successful, both in modeling psychological data [Fisher, 1987], and in solving application problems.

None of the existing concept formation models, however, examine the fundamental question of where the building blocks of correlations, namely features, come from. Invariably, features are assumed to be provided as elementary vocabulary.

In this paper, we argue that this is an unnecessarily limited paradigm within which to examine concept formation, and present possible approaches to the problem. We contend that a principled account of the origin of features can only be given with a *grounded* model of concept formation, i.e., with a model that incorporates direct access to the world via sensors and manipulators. Even though we present first steps towards such a grounded model, the primary goal of the paper is to pull together different lines of argumentation within psychology and machine learning in order to establish a solid foundation for research into grounded models.

The paper is organized as follows. In section 2, we pose the problem of features in more detail, and discuss proposed solutions from the fields of machine learning and neural networks. In section 3, we consider the possibility that features are simply "hard-wired" into the human perceptual system as innate capabilities. Based on those discussions, we then propose a framework in which the feature formation problem can be examined (section 4), and present first steps towards a model of grounded concept formation (section 5). Related work is discussed in section 6. Section 7 contains our conclusions.

*This work was partially supported by the European Community under ESPRIT contract No. P2154 "Machine Learning Toolbox".

[1] See [Wrobel, 1991] for a more detailed discussion of various concept formation approaches.

2 The problem of features

In the introduction, we have already identified the central assumption that underlies most concept formation research: the assumption that features are elementary units of description that are given and can be used as building blocks of concepts. This assumption can be seen as an instance of a basic axiom of the information processing paradigm of cognitive science and artificial intelligence, in that it assumes that there is information in the world that an intelligent agent can absorb. Without repeating the philosophical discussion about this point (cf. [Wittgenstein, 1953]), suffice it to say that this axiom is very hotly debated, and that it is being contrasted by the constructivist view that there is no information in the world (not even data), and that all structure or information is created inside the intelligent observer. In other words: Nothing in the world is independent of human observation.

Even though we do not share the extreme constructivist position, we believe that any concept formation process relies on the filtered perception/interpretation of the world that the observer imposes. Returning to features, this means that any concept formation model relying on features must include an account of their creation. Otherwise, we would have only replaced the concept formation problem by the equally hard feature formation problem.

2.1 The new-term problem in machine learning

Within machine learning, the problem with features has been recognized as the new-term problem (or the problem of *constructive induction)* in recent years: considerable effort is required to make sure that the input representation of a learning problem is stated with the "right" features that enable the learning program to find the desired target concept. J.R. Quinlan [1983], for example, has reported spending several weeks on designing a representation for learning a chess endgame concept ("lost n-ply").

Various solutions to the new-term problem have been proposed, many based on the heuristic that a new feature is to be introduced whenever the existing features no longer allow the formulation of a concept description that separates the positive from the negative examples. Utgoff's STABB [1986] and Schlimmer s STAGGER [1987] use this strategy for a learning-frorn-examples task; the BLIP system [Wrobel, 1988] has used it within observational learning. While other strategies for performing new-term construction have been used (eg. [Muggleton and Buntine, 1988; Wirth, 1989, Rouveirol and Puget, 1990]), all approaches construct their new terms by defining them in terms of already existing terms. Thus, they all share the same fundamental limitation — their new features can never leave the "closure of existing symbols", i.e., can never distinguish situations or objects that were indistinguishable with the original feature set[2]. In a sense, they are mere *abbreviations* that allow more concise concept descriptions.

[2]This is true even for recursive definitions.

2.2 Symbol grounding and Connectionism

Any true solution to the new feature problem requires access to the "real world" that the symbolic system is trying to model. It is insufficient to define the meaning of symbols by a semantics that is externally ascribed to them (eg. by the system constructor); instead, their meaning must be *intrinsically* grounded in the world — this is the problem of *grounding meaning* [Cottrell *et* a/., 1990], or more specifically, of *grounding symbols* [Hamad, to appear].

It is important to emphasize that the problem of grounding symbols is different from the symbolic/subsymbolic/nonsymbolic debate that is currently fought out within AI and related disciplines. On first sight, nonsymbolic approaches seem the right answer if symbols are the problem. In distributed connectionist approaches *(neural networks,* cf. [Rumelhart *et* a/., 1986]), for example, there are no symbols, i.e., no single physical tokens with an assignable meaning. Instead, information is represented in a network of interconnected nodes by the weights attached to the connections, and by the activation patterns of the nodes. Those activations are the result of propagating the activation of a set of input nodes along the weighted connections, modifying them accordingly. Meaning cannot be attributed to any single node or weight, only to distributed patterns of weights or activation. Thus there are no symbols — and no symbol grounding problem?

Closer inspection reveals that the crucial point is not the absence of symbols, but the role of the input nodes. In many applications of neural networks, the input nodes can be seen as simply encoding a set of (preselected) input features. In those cases, there still is a "grounding" problem without symbols: the function computed by the network is still expressed only in terms of those input features. As an example, we can look at recent research by Lee, Flowers, and Dyer [1990] on building symbols in a connectionist system. They feed encodings of propositions into a connectionist network, and manage to form hidden layer patterns that can be interpreted as symbolic representations of objects. The symbols seen as bit patterns have the desirable property of encoding their own meaning in terms of their relation to other objects, but they are still not grounded in the world — they can be no finer grained than the input propositions they were built from. Thus, the important point is not the presence or absence of symbols, but whether the learning system is actually *grounded* in its environment.

3 Features as innate structures

Even if the world does not come prepackaged into features, an important issue to consider is whether elementary features are perhaps innate structures that have developed evolutionarily. To answer this question, a large body of research exists within developmental psychology[3]. Even though there is no unanimously ac-

[3]We have used, among others, the reviews by Gibson and Spelke [1983] about perception, Mandler [1983] about representation, Sigel [1983] about concepts, and by Oerter and Montada [1982].

cepted position on some of the issues, and some of the data are contradictory, it is safe to say that in the domain of perception, humans are innately endowed with an important array of capabilities. Most notable, for our purposes here, is the ability to perceive the world in terms of objects and events from very early on: "Some mechanisms for detecting invariants are present at birth [.. .]." [Gibson and Spelke, 1983, p. 3].

Under suitable conditions, even a newborn infant will reach for a visible object. At 2 to 4 weeks, infants showed avoidant behavior (retracting their heads and interposing their hands) when confronted with an approaching ("looming") object. At one month of age, infants reliably turn their heads towards a target if that target is introduced not too far away from their line of sight. This behavior can be shown both for visual and auditory targets. At the age of 3 months, infants generally swipe at objects, and at 4 1/2 months, they begin to systematically reach for them. At 3 1/2 months, infants attend to the rigidity, and at 6 months, to the weight of objects as indicated by their anticipatory muscle tension. These findings demonstrate that even newborns are capable of perceiving objects and events in their environment and reacting to them, and that attributes such as weight and rigidity are being used relatively early.

Thus, from the outset, there seems to be a precoordinated system of perception and action in humans. This is also reflected in various proposals of representational systems developed eg. by Piaget [1977], Werner [Werner, 1948; Sigel, 1983], and Bruner [Bruner, 1973; Oerter and Montada, 1982]: they all include a *sensorimotor* level as the basis of the representation system (Bruner uses the term *enactive* representation). Above the sensorimotor level, we find a perceptual level, where experience is represented by a selective organization of percepts and images (Bruner's *iconic* level), and a symbolic level.

Nonetheless, it is also clear that the perceptual level is not cast in stone. For certain perceptual categories, within-category differences look much smaller than between-category differences even when they are of the same size physically. Instances of this *perceptual categorization* effect have been observed empirically in color and phoneme perception [Harnad, 1987]. Those effects are dependent on the categories available in the native language of a person, and can therefore not be present at birth. Furthermore, they can be modified by acquiring new categories.

In summary, it seems to reasonable to assume that in humans at least, the input to concept formation processes is at the level of sensing and acting, but not in the form of uninterpreted physical stimuli, but in a form that innately separates out objects and events. Later development then improves the precision of this process to the extent that the perceptual system can benefit from additional acquired categories and features.

4 A Research Framework

A grounded model of concept formation should therefore incorporate some degree of built-in structure on the perceptual level, but also mechanisms to change and/or augment the elementary features that were initially provided. Given the emphasis that psychological theories place on actions, events [Nelson, 1983], and goal-oriented behavior [Barsalou, 1983] (see [Wrobel, 1991] for a more detailed discussion), such a model should be developed in a framework that allows these aspects to be incorporated as well. Instead of looking at the concept formation task on isolated objects, it must be examined in the context of an agent that is *acting* in its enviroment according to a certain set of goals.

The obvious choice for such a framework would be to use a robot in a real environment, equipped with vision, multi-joint manipulators, and the ability to move around. Even though this setting is very attractive for machine learning from a theoretical as well as application point of view (see eg. Mitchell et. al.'s [1989] proposal), we believe it is unnecessarily complex for studying the issues we are interested in. Robot vision and motion control are difficult problems the solution of which is unrelated to the concept and feature formation issues that we want to examine: The discussion of innate capabilities in section 3 showed that at the sensorimotor level, a lot of precoordinated schemata exist that probably cannot be acquired by the kind of general (ontogenetic) mechanisms we are interested in here.

If one removes the vision and motion aspects of a robot, one is left with a setting where an agent inspects its environment with a number of simple sensors, has a number of simple effectors, some problem solving goals, and means of assessing its own success or failure. A concrete instantiation of this situation can be found in *process control,* a domain where AI techniques are beginning to be applied [Rowan, 1989]. There, a control system takes multiple-process sensor inputs, analyzes the data, makes decisions about the operating conditions of the process, and adjusts certain control parameters. Thus, in this concrete setting, the agents sensors are the process measurement devices, its effectors are the parameters that can be set (eg. valves), and its problem solving goal is to keep the process operating optimally, which can be measured by certain key parameters.

Within this general scenario, we will use the following assumptions:

- All sensor values are real numbers, and there are minimum and maximum values.
- All effectors are binary parameters, i.e., an action consists of an assignment of values to all effectors. This is also sufficient to model command selection by using the convention that an effector command is executed whenever its assigned value is non-zero.
- There is an fixed internal feedback function ("reward center", cf. [Whitehead and Ballard, 1990]) that computes a goal satisfaction indicator based on the current sensor values.

Is this too simple a scenario? We believe not:

- there is an environment that provides a continuous flow of input in which "events" can happen,
- the agent has goals and means of achieving those goals,

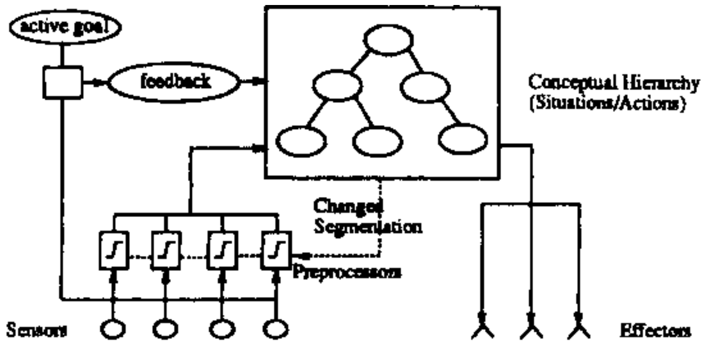

Figure 1: Sketch of a grounded concept formation model

- real-valued sensors are sufficient to model categorical perception effects, eg. for color.

Nonetheless, the models developed in this setting and the predictions they make will have to be carefully evaluated with respect to their generality; the fact that human concept formation takes place in a three-dimensional environment may require qualitatively different methods.

5 Towards a model of grounded concept formation

In this section, we present a preliminary model of a concept-forming agent that could be operating in the above "process-control" framework.

The model takes a deliberate top-down approach to the problem of grounding symbols, i.e., it tries to augment the existing successful approaches to concept formation with a grounding in the world, instead of replacing them by a totally new mechanism. This contrasts with bottom-up approaches that promise one general explanatory mechanism (cf. [Schnepf, 1990; Kugler *et al.,* 1990]), but cannot offer the capabilities of existing symbolic models (yet).

As the heart of the agent, we use a hierarchy of *probabilistic* concepts [Smith and Medin, 1981] that has been used in similar form in the clustering systems COBWEB [Fisher, 1987] or UNIMEM [Lebowitz, 1987]. The concept hierarchy is used not to cluster objects, but situations of the world in order to predict an appropriate reaction in terms of effector values. Each situation is described by a nominal (symbolic) value for each of the sensors that is computed from the real-valued "raw" sensor data by a preprocessing element[4].

The preprocessing elements segment the real-numbered value range of their sensors into a number of disjoint intervals that are each represented by a symbolic value; those values constitute the system's set of features — its elementary symbolic vocabulary. The initial vocabulary could either be a maximally coarse, i.e., binary segmentation, or a specific pre-determined segmentation that the agent could have inherited through evolution-like processes. This opens up the interesting possibility of using genetic algorithms [DeJong, 1988] to study the development of elementary feature sets.

Without specifying the details of classification and action selection (cf. [Carlson *et al.*, 1990] or [Haider, 1991] for possible approaches), we now want to briefly sketch the basic idea of a demand-driven feature refinement strategy. Whenever the system has classified a situation into its conceptual hierarchy, it executes the action that is associated with the chosen concept, and expects a positive feedback in the next sampling interval. If a negative feedback is observed instead, the system must modify its concept hierarchy by splitting the chosen concept in such a way that the erroneous action is no longer predicted. This corresponds to splitting an overly general rule into several more specific rules.

Splitting a concept node, however, is possible only if the existing concept is not maximally special yet, i.e., if the chosen concept is not a completely specified leaf of the hierarchy. Otherwise, a refinement of the existing features is necessary: the segmentation rules in the preprocessors are changed so that the segments corresponding to the attribute values of the current situation are split in half.

The reverse process, i.e., a coarsening of the available vocabulary, can be performed by analyzing the use of attribute values in the concept hierarchy, and combining the segments of those neighboring values that represent noninformative distinctions[5].

Based on the above or similar processes of vocabulary refinement and coarsening, the system develops a symbolic vocabulary that is exactly matched to its task.

6 Related Work

The approach to grounded concept formation presented here contrast with other proposals eg. by Harnad [1987] and Cottrell *et ai* [1990] that use connectionist networks to learn associations between pictorial stimuli and symbols. Our model cannot at present handle one- or two-dimensional sensor input, and cannot model the categorical perception effects that Hamad's model exhibits. On the other hand, both connectionist models require hundreds of passes through a training set to learn the required associations, and thus need input that cannot possibly come from an actual environment. Furthermore, neither one of the other two models address the question of where the categories or symbols that the network is to learn come from, and do not attempt to embed the grounding process into the context of goal-oriented activity of an agent.

In its emphasis on a complete model of an agent, the

[4] The agent's sensors sample their associated signal according to a certain sampling interval. For simplicity, one may assume the sampling interval as fixed and appropriate to the rate of change in the environment, or one could use an adaptive mechanism that increases (decreases) the sampling rate whenever a certain number of subsequent samples have been different (the same).

[5] A distinction is *non-informative* if the conditional probability of a (predicted) effector value given an attribute value is identical or nearly identical for both values. In a probabilistic concept hierarchy, this can be computed from the attribute probabilities stored with the concepts.

model presented here is similar to existing cognitive architectures like SOAR [Laird *et al,* 1986], none of which address the grounding issue, however. In turn, to be able to concentrate on symbol grounding, our model is much simpler in its treatment of goal-oriented activity and does not tackle the problem of planning at all; important questions with respect to action selection are left open. Langley *et al.* [1989] present a complete model of an intelligent robot with sensing, planning, and action that uses a very similar clustering approach as presented here, but does not specifically address the issue of symbol generation or destruction.

The process control research framework is of course not new; a great body of research exists in traditional control theory and also in AI. Specifically, our framework is very similar to the *reinforcement learning* paradigm used eg. in [Kaelbling, 1989] and [Whitehead and Ballard, 1990]. The emphasis there, however, is on action selection and not on issues of representation development.

Finally, it is important to point out the difference between our approach and learning systems that directly use numerical attributes in their concept hierarchies, such as UNIMEM [Lebowitz, 1987], CLASSIT [Gennari *et al.,* 1989], and many of the decision tree systems of the ID3/CART family. While these systems do perform implicit segmentation of numeric attributes locally in the nodes of their hierarchies, they do not acquire an explicit, global symbolic vocabulary. Most importantly, those methods store and carry along all of the original sensor data for each example. This does not correspond well with the available psychological evidence, and contrasts sharply with the *strongly incremental* approach taken here which assumes that most of the original sensor data are discarded right at the preprocessing stage, and are *not* available later on. The resulting suppression of irrelevant data right at the sensor level might also lead to conceptual structures that are simpler and more efficient than those obtained without a presegmentation stage.

7 Conclusion

In this paper, we have argued for a more encompassing framework in which to study concept formation processes. We have detailed why the issue of grounding meaning is important for a full understanding of concept formation, and why it needs to be studied in the context of a complete agent model with goal-oriented activity and true access to the world.

We have proposed a simple scenario in which we believe the problem of grounding meaning can be studied. This framework has important limitations that should be emphasized again: it does not include motion of the agent in a three-dimensional world, which may well be a key ingredient of (human) concept formation, and in its present form allows only simple real-valued or nominal sensors. At the same time, it is these very restrictions that allow an examination of the issue of symbol grounding and representation development without having to solve all existing problems of AI first.

Finally, we have sketched a model of an agent that might be able to learn to operate successfully in the simple world we have set up for it. The model is top-down in nature, and builds on existing symbolic concept formation work. It contains a demand-driven technique for introducing truly new symbols based on sensory input, and a complementary mechanism for discarding unnecessary distinctions. It has important shortcomings, among others the inability to deal with higher-dimensional input sensors.

In the future, we want to elaborate on those parts of the model that are still sketchy, and test its ability to develop a vocabulary in an actual simple process control task. The complexity of environments this model will be capable of learning will depend to a large degree on the intelligence of the yet unspecified action modifier module. Improvements in the scope of the model will have to wait until we fully understand concept formation and the origin of symbols in this simple scenario.

Acknowledgements

The author wishes to acknowledge the helpful comments of C. Haider, A. Horz, C. Lischka, Dr. K. Morik, and U. Schnepf that helped shape this paper.

References

[Barsalou, 1983] L. W. Barsalou. Ad hoc categories. *Memory & Cognttion,* 11.211 - 227, 1983.

[Bruner, 1973] J.S. Bruner. The course of cognitive growth. In J. M. Anglin, ed., *Beyond the information given: studies in the psychology of knowing.* Norton, New York, 1973.

[Carlson *et al.,* 1990] Brian Carlson, Jerry Weinberg, and Doug Fisher. Search control, utility, and concept induction. In *Proc. 1th Intern. Conf. on Machine Learning,* pp. 85 - 92, Los Altos, CA, 1990. Morgan Kaufman.

[Cottrell *et al.,* 1990] G arrison W. Cottrell, Brian Bartell, and Christopher Haupt. Grounding meaning in perception. In H. Marburger, ed., *Proc. GWA1-90, 14th* German *Workshop on Artif. Intelligence,* pp. 307 - 321, Berlin, 1990. Springer.

[DeJong, 1988] Kenneth DeJong. Learning with genetic algorithms: An overview. *Machine Learning,* 3(2/3): 121 - 138, 1988.

[Fisher, 1987] Douglas H. Fisher. Knowledge acquisition via incremental conceptual clustering. Technical Report 87-22, Department of Information and Computer Science, Univ. of California Irvine, Irvine, CA, 1987. Doctoral Dissertation.

[Gennari *et al.,* 1989] John H. Gennari, Pat Langley, and Doug Fisher. Models of incremental concept formation. *Artif. Intell.,* 40:11-61, 1989.

[Gibson and Spelke, 1983] Eleanor J. Gibson and Elizabeth S. Spelke. The development of perception. In Paul H. Mussen, ed., *Handbook of Child Psychology,* volume III, chapter 1, pp. 1 - 76. John Wiley &: Sons, New York, fourth edition, 1983.

[Haider, 1991] Christian Haider. Lernen in LeGo. *Kl - Kunstliche Jntelligenz,* 5(1): 18 - 21, 1991.

[Hamad, 1987] Stevan Harnad. Category induction and representation. In Stevan Harnad, ed., *Categorical Perception,* chapter 19, pp. 535 - 565. Cambridge University Press, 1987.

[Hamad, to appear] Stevan Hamad. The symbol grounding problem. *Physica D,* to appear. (CNLS Conf. on Emergent Computation, Los Alamos, May 1989).

[Kaclbling, 1989] Leslie P. Kaelbling. A formal framework for learning in embedded systems. In *Proc. 6th Intern. Workshop on Machine Learning,* pp. 350 - 353, Los Altos, CA, 1989. Morgan Kaufman.

[Kolodner, 1983] Janet L. Kolod ner. Reconstructive memory: A computer model. *Cognitive Science,* 7:281 - 328, 1983.

[Kugler *et al,* 1990] Peter N. Kugler, Robert E. Shaw, Kim J. Vicente, and Jeffrey Kinsella-Shaw. Inquiry into intentional systems I: Issues in ecological physics. Technical Report 30/1990, Research Group on Mind and Brain (ZiF), Univ. Bielefeld, Bielefeld, W. Germany, 1990.

[Laird *et al,* 1986] J. E. Laird, P. S. Rosenbloom, and A. Newell. Chunking in SOAR: the anatomy of a general learning mechanism. *Machine Learning,* 1(1):J1 - 46, 1986.

[Langley *et al,* 1989] P. Langley, K. Thompson, W. Iba, J.11. Gennari, and J.A. Allen. An integrated cognitive architecture for autonomous agents. UCI-ICS technical report, University of California, Irvine, CA, 1989.

[Lebowitz, 1987] Michael Lebowitz. Experiments with incremental concept formation: UNIMEM. *Machine Learning,* 2:103 - 138, 1987.

[Lee *et aL,* 1990] Geunbae Lee, Margot Flowers, and Michael G. Dyer. Learning distributed representations of conceptual knowledge and their application to script-based story processing. *Connection Science,* 4(2):313 - 346, 1990.

[Mandler, 1983] Jean M. Mandler. Representation. In Paul H. Mussen, ed., *Handbook of Child Psychology,* volume III, chapter 7, pp. 420 - 494. John Wiley & Sons, New York, fourth edition, 1983.

[Medin *et aL,* 1982] D. L. Medin, M. W. Altom, S. M. Edelson, and D. Freko. Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 8:37 50, 1982.

[Mitchell *et aL,* 1989] Tom M. Mitchell, Matthew T. Mason, and Alan D. Christiansen. Toward a learning robot. Computer Science Department CMU-CS-89-106, Carnegie Mellon University, Pittsburgh, PA 15213, 1989.

[Muggleton and Bun tine, 1988] Stephen Muggleton and Wray Buntine. Machine invention of first-order predicates by inverting resolution. In *Proc. 5th Intern. Conf. on Machine Learning,* Los Altos, CA, 1988. Morgan Kaufman.

[Nelson, 1983] Katherine Nelson. The derivation of concepts and categories from event representations. In Ellin Kofsky Scholnick, ed., *New Trends in Conceptual Representation: Challenges to Piaget's Theory?,* chapter 6, pp. 129 - 149. Lawrence Erlbaum, Hillsdale, NJ, 1983.

[Oerter and Montada, 1982] Rolf Oerter and Leo Montada, eds. *Entwicklungspsychologie.* U&S Psychologie. Urban & Schwarzenberg, Munchen, Wien, Baltimore, 1982.

[Piaget, 1977] Jean Piaget. *The development of thought.* Viking Penguin, New York, 1977.

[Quinlan, 1983] J. Ross Quinlan. Learning efficient classification procedures and their application to chess end games. In R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, eds., *Machine Learning — An Artif. Intell. Approach,* pp. 463 482. Tioga, Palo Alto, CA, 1983.

[Rosch *et aL,* 1976] E. Rosch, C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boycs-Braem. Basic objects in natural categories. *Cognitive Psychology,* 8:382 - 439, 1976.

[Rosch, 1978] E. Rosch. Principles of categorization. In E. Rosch and B. B. Lloyd, eds., *Cognition and Categorization,* pp. 27 - 48. Erlbaum, Hillsdale, NJ, 1978.

[Rouveirol and Puget, 1990] Celine Rouveirol and Jean Francois Puget. Beyond inversion of resolution. In B.W. Porter and R.J. Mooney, eds., *Proc. 1th Intern. Conf. on Machine Learning,* pp. 122 130, Palo Alto, CA, 1990. Morgan Kaufmann.

[Rowan, 1989] Duncan R. Rowan. On-line expert systems in process industries. *A J Expert,* pp. 30 - 38, 1989.

[Rumelhart *et aL,* 1986] D.E. Rumelhart, J. L. McClelland, and the PDP Research Group, eds. *Parallel Distributed Processing: Explorations in the Micro structure of Cognition.* MIT Press, Cambridge, MA, 1986.

[Schlimmer, 1987] Jeffrey C. Schlimmer. Incremental adjustment of representations for learning. In *Proc. 4th Intern. Workshop on Machine Learning,* pp. 79 90, Irvine, CA, 1987.

[Schnepf, 1990] Uwe Schnepf. Robot ethology: A proposal for the research into intelligent autonomous systems. GMD Arbeitspap. 454, GMD, POB 1240, 5205 St. Augustin 1, FRG, 1990.

[Sigel, 1983] Irving E. Sigel Is the concept of the *concept* still elusive or What do we know about concept development? In Ellin Kofsky Scholnick, ed., *New Trends in Conceptual Representation: Challenges to Piaget's Theory?,* chapter 12, pp. 239 273. Lawrence Erlbaum, Hillsdale, NJ, 1983.

[Smith and Medin, 1981] Edward E. Smith and Douglas L. Medin. *Categories and Concepts.* Harvard University Press, London, England, 1981.

[Utgoff, 1986] Paul E. UtgofT. Shift of bia* for inductive concept learning. In R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, eds., *Machine Learning An Artif. Intell. Approach,* volume II, pp. 107 - 148. Morgan Kaufman, Los Altos, CA, 1986.

[Werner, 1948] H.Werner. *Comparative psychology of mental development.* Follett, Chicago, rev. edition, 1948.

[Whitehead and Ballard, 1990] Steven D. Whitehead and Dana H. Ballard. Active perception and reinforcement learning. In *Proc. 7th Intern. Conf. on Machine Learning,* pp. 179 - 188, Los Altos, CA, 1990. Morgan Kaufman.

[Wirth, 1989] Ruediger Wirth. Completing logic programs by inverse resolution. In K. Morik, ed., *Proc. 4th Europ. Working Sess. on Learning,* pp. 239 - 250, London/San Mateo, CA, 1989. Pitman/Morgan Kaufmann.

[Wittgenstein, 1953] L.Wittgenstein. *Philosphical investigations.* Macmillan, New York, 1953.

[Wrobel, 1988] S. Wrobel. Automatic representation adjustment in an observational discovery system. In D. Sleeman. ed., *Proc. of the 3rd Europ. Working Sesston on Learning,* pp. 253 262, London, 1988. Pitman.

[Wrobel, 1991] Stefan Wrobel. Concepts and concept formation: Fundamental issues. GMD Arbeitspap., GMD, POB 1240, 5205 St. Augustin 1, FRG, 1991. To appear.

[Younger and Cohen, 1984] B. A. Younger and L. B. Cohen. Infant perception of correlations among attributes. *Child Development,* 54:858 - 867, 1984.