

Discrimination Nets, Production Systems and Semantic Networks: Elements of a Unified Framework

Fernand Gobet
ESRC Centre for Research in Development,
Instruction and Training
Department of Psychology
University of Nottingham
Nottingham NG7 2RD
England
frg@psyc.nott.ac.uk

A number of formalisms have been used in cognitive science to account for cognition in general and learning in particular. While this variety denotes a healthy state of theoretical development, it somewhat hampers communication between researchers championing different approaches and makes comparison between theories difficult. In addition, it has the consequence that researchers tend to study cognitive phenomena best suited to their favorite formalism. It is therefore desirable to propose frameworks which span traditional formalisms.

In this paper, we pursue two goals: first, to show how three (symbolic) formalisms widely used in theorizing about and in simulating human cognition—discrimination nets, semantic networks and production systems—may be used in a single, conceptually unified framework; and second to show how this framework can be used to develop a comprehensive theory of learning. Within this theory, learning is construed as (a) developing perceptual and conceptual discrimination nets, (b) adding semantic links, and (c) creating productions.

We start by giving a brief description of each of these formalisms; we then describe a theoretical framework that incorporates the three formalisms, and show how these may coexist. Throughout this description, examples from chess, a highly studied field of expertise and a classical object of study in cognitive science, will be provided. These examples will illustrate how the framework can be worked out into a more detailed cognitive theory. Finally, we draw some theoretical consequences of the framework proposed here.

Three Formalisms Widely Used in Symbolic Models

Discrimination Nets and Trees

Discrimination nets have been used, among other things, to model decision making, concept formation and recognition processes. The EPAM model, one of the most developed models that uses this formalism [Feigenbaum & Simon 1984; Richman, Staszewski & Simon 1995] simulates human behavior in various learning and perceptual tasks. The basic idea behind EPAM's discrimination tree is that, during perception, an object is sorted through a sequence of tests, each related to some feature of the object. When the description of the object mismatches the internal representation (the *image*) it has been sorted to, a new test is added in the tree, test that relates to the mismatched feature. When the object is sorted to an internal representation that under-represents it, new features are added to the image by chunking.

Semantic Networks

Semantic networks are labeled graphs, where both nodes and links may be labeled. Nodes represent various entities (objects, situations, concepts) in a domain, and links represent the relations between these entities. It is assumed that information about an object may be accessed by following the relevant links. Many semantic network models assume that some kind of activation spreads along the links leaving a node. Theories using the semantic network formalism have been widely used—with sometimes important differences between them—both in psychology [e.g. Anderson & Bower 1973; Rumelhart, Lindsay & Norman 1972] and in computer science [see Findler 1979].

Production Systems

Production systems represent knowledge as condition-action pairs. When the condition of a production is met, its action is executed. Various conflict resolution rules are specified in the case where several productions apply. In the extreme, as in Soar [Newell 1990], all productions are executed in parallel. Production systems vary in their assumption about the limitations of working memory, where matching is taking place. In cognitive psychology, production systems include, among others, Soar [Newell 1990], CAPS [Just & Carpenter 1987], and ACT* [Anderson 1983]. Note that the latter uses two of the formalisms discussed here: productions to encode procedural knowledge and semantic networks to encode declarative knowledge. The connection between procedural and declarative knowledge takes place through working memory, which is both the activated part of declarative knowledge and the place where the matching of conditions occurs.

Comparison between the three formalisms

At an abstract level, the three formalisms have the same power, in that they are computationally equivalent to a universal Turing machine. In particular, each formalism may simulate the two others. However, they differ in the way they represent information, which results in differences in the efficiency with which information is processed. For example, discrimination trees are efficient in sorting an object into its internal representation rapidly, but they are awkward when the goal is to represent a complex network of knowledge, with cross-indexing [Feigenbaum & Simon 1984]. For the latter purpose, semantic networks are more efficient.

Interestingly, these three formalisms have been traditionally used to model three different kinds of behavioral data in psychology. Discrimination nets have been mainly used to model recognition processes; semantic networks to model declarative knowledge; and production systems to model procedural knowledge. This tradition suggests that there exist differences in the cognitive mechanisms used to process these types of information.

We think that using these three formalisms together in a unified framework offers a promising approach to the study of cognition, because this allows the combination of three approaches that have each been shown to be successful at modeling complementary aspects of the cognitive system. By definition, however, a framework is too unspecified to pretend to be a theory, let alone a model. We will therefore have to flesh out our framework with a few additional assumptions, in order to present (the outline of) a workable theory. In order to make the exposition more concrete, we will illustrate our theory with empirical data. We have chosen chess, because the three aspects of cognition we are interested in—pattern recognition, declarative knowledge and procedural knowledge—play a key role in this domain of expertise.

The Composite Framework

In this section, we present a theoretical framework where the formalisms of discrimination nets, production systems and semantic networks are tightly interconnected. Three basic ideas underlie this framework: (a) a difference is made between the index and the text of semantic long-term memory (LTM) [Simon 1989, p. xii]; (b) nodes used by the production system and semantic memory mechanisms are acquired through the discrimination processes; and (c) the same node can both belong to a production and be semantically linked. In addition, an important assumption of the theory derived from this framework is that the information used to build up a new link (i.e. nodes being linked and possibly the node labeling the link) has to be stored in short-term memory (STM) during link creation.

The Discrimination Net Component

We use an EPAM-like discrimination net that acts as an index to semantic LTM, where productions and schemas are stored. Since recognition should happen in real time in order to fulfill the requirements of the environment, the discrimination net has a simple structure and processes information rapidly. In comparison with EPAM, there are two extensions. The first addition is the presence of *redundant links*, which connect nodes having similar information. When two nodes containing the same information, but which has been tested in a different order, are placed into STM, a redundant link connects these two nodes [see Fig. 1]. The redundant link is used in future traversing of the net. Redundant links have been implemented in CHREST, a model of chessplayers' memory and perception [Gobet 1993a, 1993b; de Groot & Gobet 1996]. We are now exploring the idea of relaxing the matching constraints, by allowing a link to be created when only a partial matching occurs

(say, at least 80% of the features of each node). Such partial matching may be useful in modeling concept and prototype formation.

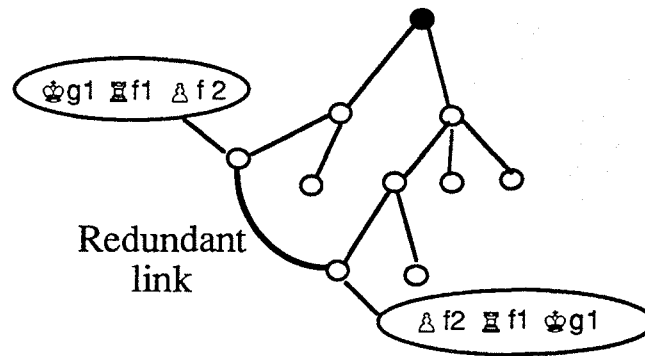


Figure 1: Example of a redundant link. The black circle represents the root node, empty circles nodes, and ellipses node images.

The second addition is the presence of several nets based on stimuli types, possibly processing in parallel. For example, a net may discriminate auditory stimuli, and another visual stimuli. At the conceptual level, a net may discriminate animals, and another vegetables. These nets may share some of their nodes. For example, the title of a movie may be accessed both through some graphic scene (visual net) or through the soundtrack (auditory net).

With these two additions, we have a discrimination *net* (a given node may be accessed by several parent nodes), while EPAM is normally implemented as a discrimination *tree* (a given node may be accessed by a unique parent node).

Illustration

Using a discrimination tree formalism along with additional assumptions about the time to access a node and about STM capacity, [Simon & Gilmarin 1973] have simulated some aspects of chess expertise (recall performance of an average player, types of chunks recalled). The same mechanisms and assumptions are incorporated in CHREST, which simulates a wider range of data from chessplayers' memory and perception, including their eye movements in a recall task. In both models, the discrimination net is used to sort patterns of pieces.

An important aspect of chessplayers' knowledge is that it is richly indexed. Players can access a known position when looking at the configuration of pieces on the board, when told the name of a opening variation, or when given the names of two players and the date and place of the game. We propose that the ease with which knowledge is accessed is made possible partly by the presence of several discrimination nets leading to a given node.

The Production System Component

In our framework, productions consist in connecting two nodes from different nets. The first node acts as condition, the second as action [see Fig. 2]. A condition for the construction of a *production link* is that both nodes are simultaneously stored in STM.

Illustration

[Chase & Simon 1973] have proposed that the recognition of chess piece patterns allows accessing information about good moves. This idea has been embodied in CHUMP [Gobet & Jansen 1994], a variant of CHREST. CHUMP implements two types of knowledge through discrimination nets. The first type of knowledge relates to patterns of pieces. The second type of knowledge relates to moves and sequences of moves. During learning, patterns of pieces are associated with potential moves. During the performance phase, patterns of pieces act as conditions and moves as actions.

CHUMP covers only a small part of the knowledge that is presumably encoded as productions in chess experts. It is plausible that they have other productions where the conditions consist in nodes containing information such as positional and tactical features, types of opening, etc., while the nodes denoting actions belong to nets encoding information such as moves, plans, heuristics, tactical tricks, etc.

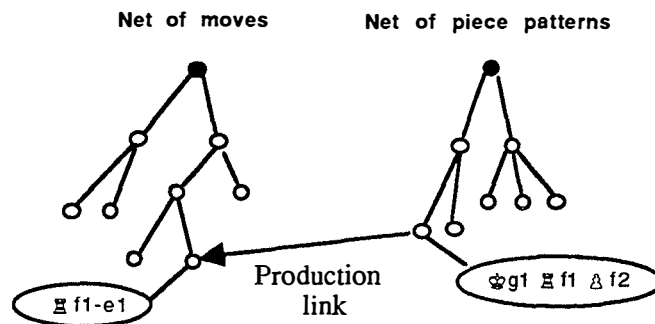


Figure 2: Example of a production link, where a pattern of pieces suggest a move. The black circles represent root nodes, empty circles other nodes, and ellipses node images.

The Semantic Network Component

Three discrimination nets are used in the construction of *semantic links* similar to the links used in earlier semantic memory models. The first two nodes are the nodes involved in some kind of relation, and the third node *labels* this relation. (We may assume that, by default, unlabeled links are redundant links when the relation occurs within a net, and are production links when the relation occurs between nets.) As with the two previous mechanisms, the (three) nodes have to be simultaneously in STM for a semantic link to be created.

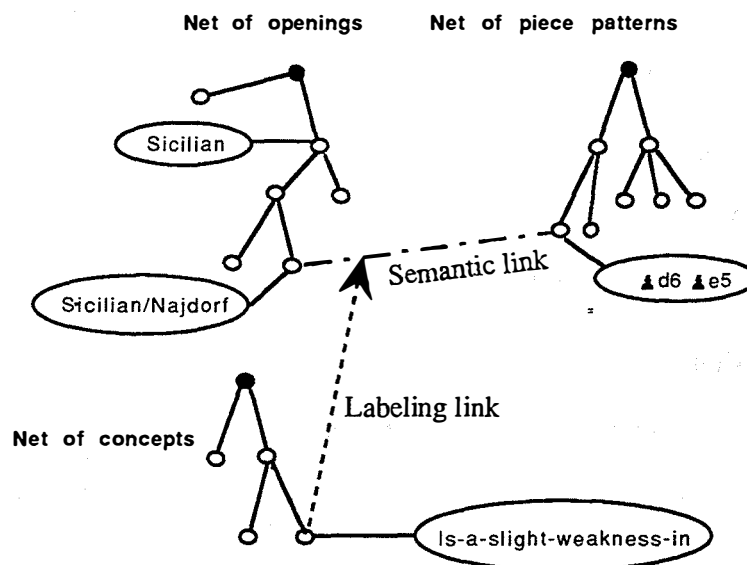


Figure 3: Example of a semantic link. The black circle represents root nodes, empty circles other nodes, and ellipses node images.

Illustration

Assume the following statement: "The pawns d6-e5 are slightly weak pawns in the Najdorf variation of the Sicilian defense." In order to represent this statement, a link is created between the node <d6 e5> from the net of piece patterns and the node <Najdorf variation> from the opening net, and this link is labeled by the node <is-a-slight-weakness-in> from the chess concept net [see Fig. 3].

We may call a cluster of semantic links a schema. As was proposed by [Chase & Ericsson 1982] and [Richman et al. 1995] with their concept of *retrieval structure* and by [Gobet & Simon in press-a] with the notion of *templates*, some schemas may be used to augment STM capacity and thus enhance immediate storage capacity. In our framework, this means that some schemas have slots into which values may be rapidly placed.

In summary, nodes can be connected with three types of links: (a) redundant links (within nets); (b) production links (between two nodes from different nets); (c) semantic links (a third node labels the link between two nodes from different nets).

Learning and Development : Consequences of the Theory

We assume that each net starts with a few biologically rooted primitives, and that learning occurs by chunking and by connecting nodes within and between the different nets. Then learning consists in (a) the growing of several nets; (b) the growing of redundant nodes within these nets; (c) the growing of connections (productions and semantic memory links) between these nets, leading possibly to (d) the growing of templates.

This complex learning explains the ten years needed to reach top-level expertise more convincingly than the creation of a simple net as in [Chase & Simon 1973]. Several nets need to be constructed, with many connections within them (redundancy) and between them (production and semantic links).

It may be the case that the development of some nets is conditional on the development of other nets. In the chess domain, for example, it seems plausible that novices will focus on abstract, non-located relations between pieces, like "Knight attacks Piece that defends Pawn," and grow a discrimination net related to such features. Such abstract knowledge is general, but the price of it is that it takes time to "interpret" and to apply in a given position. With higher levels of expertise, and building on previous nets, players will tie such abstract, declarative knowledge to specific instances, and grow a net with location-bound patterns of pieces; this compiled, variable-free, knowledge is, of course, limited in its application, but is highly reliable and can be accessed very rapidly [Gobet & Simon in press-b].

The information encoded at early stages of learning is important in later stages and a net built up in an inefficient way may have negative consequences later. For example, suppose that a player tends to systematically consider Queen's sacrifices when studying positions in practice sessions. According to the theory, this will lead, over the years, to a plethora of links between patterns of pieces and moves including a Queen sacrifice. Later on, such an examining of Queen's sacrifices may impair performance, even if this player focuses on more important aspects later on in the analysis, because important resources are wasted on the consideration of (in most cases) irrelevant moves.

This suggests that the early phases of chess training are very important, in preventing the development of "bad habits," and explains why it is desirable to have a trainer early on who is able to direct attention to important aspects of chess. It may be possible to overcome deficient learning, either by reorganizing the information in the net or by creating a new net with more adequate information. Both are probably quite difficult and time consuming. Thus, a further hypothesis is that the long plateaus that some players experience in their chess development may be due to such net reorganization or creation.

Implications for Problem Solving

Our framework, although directed mostly to modeling perception and memory processes, has some interesting consequences for the theory of problem solving. Refining the account of [Simon 1989], we propose that a cognitive system copes with problem situations in the following way. There is first a recognition phase. If traversing the net leads to a node in LTM, one or more of the following may occur. The node may lead, through a production link, to an action being performed. In this case, the action modifies either the external environment, or the internal environment (working memory). In the case where the node has semantic links, some of the links may be accessed to reach relevant features. The retrieval of information then changes the content of working memory. If a template is retrieved, either through pattern recognition or by following semantic links, it will allow both the access of a rich network of knowledge and also the circumvention of STM limitations for further processing.

By increasing the number of nodes in LTM and the links between them, an expert is therefore able (a) to reach solutions rapidly, by pattern recognition; (b) to categorize problems more efficiently, thus allowing swift access to additional information; and, as a consequence, (c) to reduce the need for search. In summary, in accordance with what is known about experts' behavior, the mechanisms embodied in our theory make the search space *denser*, in the sense that attention is directed only to that portion of the original search space likely to

contain the solution [Newell & Simon 1972]. Note that the presence of an efficient discrimination net resolves the *paradox of expertise*: How is it that, despite the large amount of knowledge experts possess, they are able to access the relevant information so rapidly? (For an alternative view, see [Ericsson & Staszewski 1989].)

Conclusion

This paper has presented a theoretical framework unifying the formalisms of discrimination nets, production systems and semantic networks. It has been argued that the unified formalism captures three aspects of cognition that are typically not treated together in cognitive modeling: pattern discrimination, semantic memory, and procedural knowledge. For each of these aspects, different learning mechanisms have been hypothesized. Although the examples and the discussion of some of the consequences of the framework have been drawn mainly from the domain of chess expertise, the framework can be applied to other domains of cognition as well. Obviously, this framework needs additional assumptions (such as parameters specifying the rate of learning, and, possibly, the rate of decay in memory) before one has a fully-fledged psychological theory. As has already been mentioned, we have made progress in this direction, with some aspects of this framework already implemented in a computer program.

References

- [Anderson 1983]. Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- [Anderson & Bower 1973]. Anderson, J. R. & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- [Chase & Ericsson 1982]. Chase, W. G., & Ericsson, K. A. (1982). Skill and working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 16). New York: Academic Press.
- [Chase & Simon 1973]. Chase, W. G., & Simon, H. A. (1973b). The mind's eye in chess. In W. G. Chase (Ed.) *Visual information processing*. New York: Academic Press.
- [de Groot & Gobet 1996]. de Groot, A. D. & Gobet, F. (1996). *Perception and Memory in Chess: Studies in the heuristics of the professional eye*. Assen: Van Gorcum.
- [Ericsson & Staszewski 1989]. Ericsson, K. A., & Staszewski, J. J. (1989). Skilled memory and expertise: mechanisms of exceptional performance. In D. Klahr. & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon*. Hillsdale, NJ: Erlbaum.
- [Feigenbaum & Simon 1984]. Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, 8, 305-336.
- [Findler 1979]. Findler, N. V. (Ed.). (1979). *Associative networks: The representation and use of knowledge by computers*. New York: Academic Press.
- [Gobet 1993a]. Gobet, F. (1993a). *Les mémoires d'un joueur d'échecs* [The memories of a chess player]. Fribourg: Editions Universitaires.
- [Gobet 1993b]. Gobet, F. (1993b). A computer model of chess memory. *Proceedings of 15th Annual Meeting of the Cognitive Science Society*, (pp. 463-468).
- [Gobet & Jansen 1994]. Gobet, F. & Jansen, P. (1994). Towards a chess program based on a model of human memory. In H. J. van den Herik, I. S. Herschberg, & J. W. Uiterwijk (Eds.), *Advances in Computer Chess 7*. Maastricht: University of Limburg Press.
- [Gobet & Simon in press-a]. Gobet, F. & Simon, H. A. (in press-a). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*.
- [Gobet & Simon in press-b]. Gobet, F. & Simon, H. A. (in press-b). Recall of random and distorted positions: Implications for the theory of expertise. *Memory & Cognition*.
- [Just & Carpenter 1987]. Just, M. A. & Carpenter, P. A. (1987). *The psychology of reading and language comprehension*. Newton, MA: Allyn and Bacon, Inc.
- [Newell 1990]. Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- [Newell & Simon 1972]. Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- [Richman, Staszewski & Simon 1995]. Richman, H. B., Staszewski, J., & Simon, H. A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review*, 102, 305-330.
- [Rumelhart, Lindsay & Norman 1972]. Rumelhart, D. E., Lindsay, P., & Norman, D. A. (1972). A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New-York: Academic Press.
- [Simon 1989]. Simon, H. A. (1989). *Models of thought*. (Volume II). New Haven: Yale University Press.
- [Simon & Gilmarin 1973]. Simon, H. A., & Gilmarin, K. J. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, 5, 29-46.