# Emergent Letter Perception: Implementing the Role Hypothesis

**Gary McGraw**[†*] and **Douglas R. Hofstadter**[†]
[†]Center for Research on Concepts and Cognition
510 North Fess Street, Bloomington, IN 47405
`http://www.cogsci.indiana.edu`
&
[*]Reliable Software Technologies Corporation
21515 Ridgetop Circle, Suite 250, Sterling, VA 20166
`http://www.rstcorp.com`
`gem@rstcorp.com` `dughof@cogsci.indiana.edu`

## Abstract

Empirical psychological experimentation (very briefly reviewed here) has provided evidence of top-down conceptual constraints on letter perception. The role hypothesis suggests that these conceptual constraints take the form of structural subcomponents (roles) and relations between subcomponents (r-roles). In this paper, we present a fully-implemented computer model based on the role hypothesis of letter recognition. The emergent model of letter perception discussed below offers a cogent explanation of human letter-perception data — especially with regard to error-making. The model goes beyond simple categorization by parsing a letterform into its constituent parts. As it runs, the model dynamically builds (and destroys) a context-sensitive internal representation of the letter that it is perceiving. The representation emerges as by-product of a parallel exploration of possible categories. The model is able to successfully recognize (*i.e.*, conceptually parse) many diverse letters at the extremes of their categories.

## The role hypothesis

Results from a series of previously reported psychological experiments in human letter recognition provide empirical evidence for the existence of conceptual-level representations of letter-parts that we call *roles* (McGraw et al., 1994; McGraw, 1995). The *role hypothesis* of human letter recognition jibes with other psychological theories of perception that posit higher-level relational structure, including work by Palmer (1977), Treisman and Gelade (1980), Hock et al. (1988), Biederman (1987), and Sanocki (1986). In the role hypothesis, the conceptual components of a letter representation are not explicit shapes *per se* but are ideas about what acceptable bounds for letter-part shapes are, how far such shapes can be stretched before they lose their interpretation, and how they interact with other roles to form a complete object.

Letterforms, or *physical* instances of letters, are made up of parts that correspond to the conceptual roles of the *mental* level. Little work on machine-based letter-recognition systems has considered intermediate-level parts (sometimes called "high-level features" in the literature), let alone collections or groups of such parts (Mori et al., 1984; Gaillat & Berthod, 1979). One important exception to this trend is the work done some two decades ago by Barry Blesser's research group[1] (Blesser et al., 1973). The Blesser group's approach to machine letter recognition was, like ours, strongly based on the psychology of human letter perception (Naus & Shillman, 1976). The idea was to describe letters not in terms of their physical attributes, but in terms of more general descriptions of their underlying representations. Rich also briefly mentions an idea for a part-based model of letter recognition in her introductory AI text (Rich, 1983). The model she sketches is similar in spirit to the implemented work of Sanocki (1986; 1991).

## Letter concepts, roles, and parts

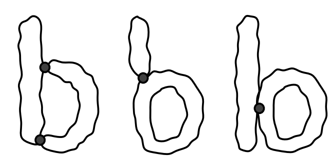

Figure 1: Three common conceptualizations of the letter 'b', featuring two roles apiece.

Figure 1 is a graphical representation of three common ways one can break the abstract concept of letter 'b' into conceptual pieces according to the role hypothesis. Roles (the wiggly outlines) and r-roles that include relationships between roles (the black dots) make up the internal structure of a letter category and together define a particular *conceptualization* of a letter. Category membership at the whole-letter level is partially determined by category membership at the lower level of roles. (We say "partially" because the interaction between roles also matters. For example, a graphic shape might have a strong exemplar of *post* to the left of a strong exemplar of *loop* and yet the way they interact might still make them look more like '**lo**' than '**b**'.)

Figure 2 shows how actual letterforms (*e.g.*, shapes on a page) are comprised of *parts* that fill a letter-conceptualization's roles. During perception, representations of such parts are formed under top-down pressure from roles and are sensitive to context. As stated by Palmer (1978) [p. 96], "components [or parts] enter into relationships with other components, resulting in larger structural units whose importance supersedes that of [their] constituents." We hold that most of the

[1] Blesser's group included Shillman, Cox, Naus, Kuklinski, Ventura, and Eden.

"importance" attributed to the emerging parts stems directly from their role-filling ability. In other words, the way in which a part fills a role directly determines its "goodness". Experimental evidence reported in (McGraw, 1995) supports this claim by showing that human subjects prefer parts that correspond to natural roles over parts that fit the Gestalt criteria described by Palmer (1978).

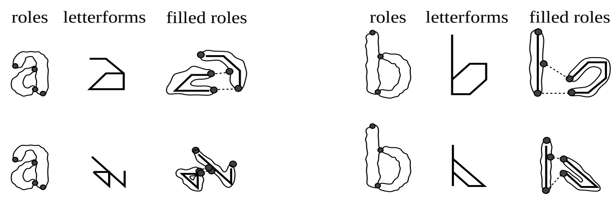


Figure 2: Parsing letterforms into high-level parts occurs under the top-down influence of roles. The top two examples show straightforward parsings (or, *role-fillings*) requiring little plasticity. The bottom two are more complicated, with role-plasticity performing a critical function during recognition.

Roles are "plastic" — their plasticity resulting from the fact that a role can be filled by a whole host of differently shaped parts. Roles are defined in terms of *norms*. Sometimes norms associated with roles must be *violated* in order to accommodate a given letterform's parts. The plasticity of a role is context-sensitive and varies according to conceptual pressures brought to bear by the situation.

## Conceptual-level letter recognition

Human letter recognition has many distinct flavors. Letter recognition used while reading sentences in a uniform book-face, for example, is far different than recognition of a style-rich letter in a display face for advertising. The sort of recognition that we strive to model is more closely-related to display-face recognition. Specifically, our research concentrates on recognizing one letter at a time with no word-level context. Particular letters range all the way from very normal to stylistically-loaded. The human results that we use for comparative purposes were collected during just this sort of recognition task — that is, the recognition of single letters with no word-level context. Our model is driven by the need for a flexible and powerful letter-recognizer to be included in an analogy-based model of typeface creation called Letter Spirit.

Our model is based on the tenets of high-level perception, in which concepts dynamically provide top-down influence on the formation of perceptual structures during the process of categorization. The resulting perceptual structures, which emerge from the stochastic activities of a large number of tiny processing agents, are well-suited for further analogical processing (Mitchell, 1993). These perceptual structures include important information about the stylistic attributes of a letterform (as opposed to a mere categorization, like "is an 'a'"). This "conceptual parsing" is critical to the design of stylistically uniform alphabets — the ultimate aim of Letter Spirit.

Why focus on cognitive plausibility of recognition models instead of engineering efficiency? Because human recognition of letterforms (especially highly-stylized ones) is still far superior to that of machines. We believe that a human-like approach will significantly enhance the capacity of computers to correctly recognize a wide variety of non-standard letterforms.

## Letter Spirit

The Letter Spirit project is an attempt to model central aspects of human high-level perception and creativity on a computer, focusing on the creative act of artistic letter-design.[2] The aim is to model the process of rendering the 26 lowercase letters of the roman alphabet in many different, artistically coherent styles. Two important and orthogonal aspects of letterforms are basic to the project: the *categorical sameness* possessed by instances of a single letter in various styles (*e.g.,* the letter 'a' in Baskerville, Palatino, and Helvetica) and the *stylistic sameness* possessed by instances of various letters in a single style (*e.g.,* the letters 'a', 'b', and 'c' in Baskerville). Figure 3 shows the relationship of these two ideas. Initial work on the Letter Spirit project has been focused on the "letter" aspect. The model described in this paper is able to successfully recognize hundreds of letters from all 26 categories. We aim to show that the model does this in a similar fashion to the way people do.

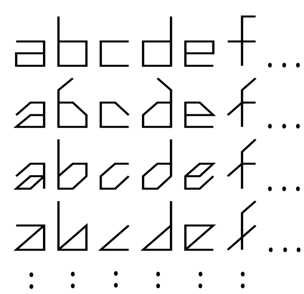

Figure 3: Items in any column have *letter* in common. Items in any row have *spirit* in common.

To avoid the need for modeling low-level vision and to focus attention on the deeper aspects of letter recognition, we developed an idealized micro-domain. Letterforms are restricted to short line segments on a fixed grid of 21 points arranged in a $3 \times 7$ array. Legal line segments, connect a point to any of its nearest neighbors. There are 56 possible segments, as shown in Figure 4. This restriction allows much of low-level vision to be bypassed and forces concentration on higher-level cognitive processing, particularly the abstract and context-dependent character of concepts.

[2]For information about the on-going Letter Spirit project see (Hofstadter & McGraw, 1993) and (McGraw & Hofstadter, 1993), available on the World Wide Web through URL `http://www.cogsci.indiana.edu`.

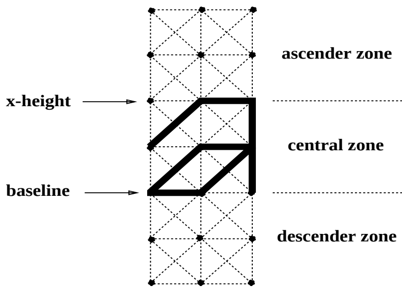


Figure 4: The Letter Spirit grid, with line segments instantiating one of many possible 'a's turned on.

## Emergent letter perception

The fully-implemented gridletter recognizer in Letter Spirit is called the role model. All perceptual and creative processes in the role model are emergent, in the sense that they result from the actions of a large number of independent *codelets* — computational micro-agents that create, examine, and modify structures representing parts, roles, letters, stylistic traits, and so forth. Codelets perform these activities in (simulated) parallel. Newly created codelets wait to be run in a structure called the *Coderack*, which can be thought of as a stochastic waiting room. In contrast to a standard operating-systems queue, where processes wait before being deterministically given their slice of CPU time, the Coderack features stochastic selection of actions. To each codelet is attached an *urgency* value — a number that determines its probability of being chosen next. Urgency values are based on how well a codelet's possible effect coheres with structures already built.

Actions of every sort — gluing, labeling, scanning, matching, adjusting, regrouping, destroying, and so on — are carried out by codelets. The effect of each codelet considered by itself is very slight; however, as many codelets run, their independent effects build upon one another into a coherent collective behavior.

Over a long period of time, *processes* are interleaved in a manner reminiscent of time-sharing. (A process consists of many codelets, which *ex post facto* can be seen to have been acting in concert.) One notable difference between this and conventional time-sharing is that the biased nondeterministic selection of codelets amounts to having different processes run at different speeds. The speeds themselves are regulated over time, by varying the urgencies of the codelets involved, in an effort to favor more-promising directions over less-promising ones. Since codelets have very small effects, it is never critical that any particular codelet get selected. What does matter is that certain broad-stroked courses of action as a whole run faster than others. Probabilistic selection based on urgencies allows this to happen.

Many of the ideas behind this model (and its Copycat predecessor (Mitchell, 1993)) were originally inspired by the Hearsay II speech-understanding system (Erman et al., 1980). Hearsay II introduced the idea of simultaneous bottom-up and top-down influences interacting in the process of perception. A complete comparison of the role model and Hearsay II, including a discussion of many important distinctions and differences can be found in (McGraw, 1995).

### Processing in the role model

The role model operates roughly as follows, although this outline may give the impression that processing is more serial and well-ordered than it really is. In reality, processing occurs in a more parallel manner, with various aspects described below often proceeding concurrently. See (McGraw, 1995) for a thorough account of the role model.

- Line segments are probabilistically bonded together (by local perceptual codelets) with different amounts of "glue". (For example, more glue tends to be deposited at straight junctions than at angles.) The gluing codelets agents execute in a completely bottom-up fashion.
- When enough glue has been deposited, the glued shape is metaphorically "shaken". This amounts to probabilistically breaking the glued shape into chunks of segments at weak joints, resulting in a set of parts (usually made up of between two and four line segments).
- Each part is scanned by multiple codelets that probabilistically attach *syntactic labels* to the part. Labels reflect very simple properties of parts like curviness, length, width, location, and so on, and in no way involve the set of categories (either at the level of roles or wholes) into which these stimuli will eventually be channeled. Each syntactic label has a real number associated with it, standing for the strength with which the label applies to its part.
- When a part has accumulated enough syntactic labels (once again, a probabilistically determined event), it is allowed to send activation to one or more *roles*.
- The presence of a particular label on a given part serves as a cue that tends to lightly activate one or more roles with which the label is associated (*i.e.*, roles of which the given label is at least somewhat diagnostic). For instance, the labels "left-side", "straight", "skinny", and "tall" would tend to activate the "left-post" role found in 'b', 'h', 'k', and sometimes 'l'. It is important to understand that even a label such as "straight" is probabilistic, in the sense that a not-totally-straight part *might* get that label, with probability diminishing with its non-straightness. The real number associated with each label reflects this "goodness of fit".
- A part that does not strongly activate any roles will be slated for destruction, with its constituent line segments subjected to the part-forming process all over again.
- The various light activations coming from a given part's labels sum up to a *total* activation-level for each role that the part matches sufficiently. If activated highly enough, roles associate themselves with labeled parts, with the best-matching roles getting the most activation from a part. Each particular role may be filled by one part at any given time, although inter-part competition for the role's attention is ongoing and sometimes fierce.
- As roles and parts attempt to "mate", a given part may need to be slightly altered in order to be a good mate for a given role. A quantum or two may need to be stolen from one or more neighboring parts to make the part in question more attractive to a possible match. Likewise, small pieces that seem to make a part ugly in the eyes of

possible role-mates may need to be given away or simply detached. The resulting structures composed of groups of the initial line segments are now results of the combined influence of bottom-up and top-down processing. As such, these parts are no longer totally syntactic entities, and we call them *(semantically) adjusted parts.*

- Roles compete for parts throughout the letter, adjusting the parts as they go. When this adjustment-and-association phase is over, there is a fairly strong match-up between roles and semantically-adjusted parts.
- Each instantiated role has a few tags attached, stating *how well* the given part instantiates the role. This information focuses on how the part deviates from various norms associated with the role.
- Each realized role begins to alert one or more *wholes* (*i.e.*, full role-sets, such as those shown in Figure 1) for which it provides evidence, in the sense of fitting a particular conceptualization of that letter. Role/whole coupling is analogous to part/role coupling, only it occurs at a higher (more semantic) level. Particular letter-conceptualizations are activated according to how strongly their component roles are realized in the actual grid letter.
- The activation level of each hypothesized whole is adjusted according to whether the whole's r-roles (*i.e.*, inter-role relations) approve of the structure discovered so far. When r-roles for a particular whole are checked, activation is taken away from any whole whose roles are not appropriately related or filled. This is a critical inhibitory aspect of the categorization process.
- Different wholes thus become activated to different extents. Each sufficiently-activated whole attempts to match itself up with the shape on the grid.
- If there is a clear leading contender among the wholes, it is deemed the winner. If there is a close race between several, the letterform is deemed ambiguous and therefore unacceptable. In a borderline case between clear-winner and close-race, a probabilistic decision is made that chooses between the two courses of action.
- In the end, the winning whole has been parsed into constituent, non-overlapping parts. These parts fill specific roles in the whole to a greater or lesser extent. The tags attached to roles according to how well they are filled by their associated parts are also available when processing ends. For example, a 't' whose spine is too tall or is bent over at the top, or whose crossbar is too short, too high, or tilted, will have tags stating such things attached to its filled conceptualization. This information can be used in further processing, including checking the style of a letterform and designing related letterforms by analogy.

The entire labeling and role-association process happens in parallel for each of the initial parts created after shaking. In general, processing proceeds from low-level syntactic processing to high-level semantic (or conceptual-level) processing, but is not completely linear in nature. Conflicting perceptual structures compete against each other in the context of current perceptual trends. Weak structures tend to be destroyed and strong ones to be strengthened. A perceptual parsing in terms of role-filling parts emerges as a result of this perceptual "competition".

## Emergence, subsymbols, and symbols

It is difficult to pigeonhole the role model as belonging to any particular well-known cognitive-science school, as it is neither fully symbolic nor fully connectionist. Instead, it takes some of the central features of both paradigms and mixes them together, and thus could be said to fall somewhere between symbolic AI and connectionism.

In common with connectionism, the role model has many important subsymbolic characteristics of the sort that Smolensky advocates (Smolensky, 1988). In the subsymbolic paradigm, cognitive representations are built of *subsymbols* that in turn give rise to symbol-like structures. In systems of the subsymbolic vein, symbols are *statistically emergent* entities that are represented by patterns of activation over large numbers of subsymbols. The role model's fine-grained parallelism, local actions, competition for limited computational resources, spreading activation, and emergent concepts are all faithful to the subsymbolic enterprise. Also closely related is the interaction of top-down and bottom-up processing in the model.

The role model's representations also have something in common with more traditional symbolic methods. The emergent representations that the system develops are able to be quickly and easily *referenced* and *explicitly manipulated* since they are cut from symbolic cloth (*i.e.*, they are made up of Scheme structures, albeit with attached and dynamically varying activation values). Furthermore, these symbolic structures are built up in a workspace similar in some respects to a short-term memory — something not often found in connectionist models. The notion of *reference* is a critical one in models of higher-level cognitive activities such as analogy-making (Indurkhya, 1992). By their very nature, symbols — even of the active, emergent variety that we model — provide a natural avenue for such reference.

## Performance of the role model

An intuitive way to illustrate the flexibility inherent in and emergent from the role model's architecture is to consider a series of runs on groups of letters. Doing this gives some idea of the sorts of letterforms the model is capable of recognizing and the sorts of letterforms it fails on. Figure 5, below, shows the extent to which the role model can handle cases where stylistic aspects of a letterform begin to overcome its category.

The particular group of letters we discuss here is taken from page 424 of (Hofstadter & FARG, 1995). The dataset consists of 88 lowercase gridfont 'a's. We ran the role model at least ten times on each of the 'a's in the original illustration, in order to discover which ones are easily recognized by the model and which are not.

In Figure 5 we have arranged the 'a's of the original chart into a bull's-eye pattern, whose center is made up of 'a's that were correctly recognized 10 times out of 10. Concentric rings surrounding the bull's-eye the remaining 'a's into "recognition bins", the first consisting of 'a's recognized from 5 to 9 times over 10 runs, the second being those recognized from 2 to 4 times. Finally, outside of the outermost closed curve are those 'a's that the role

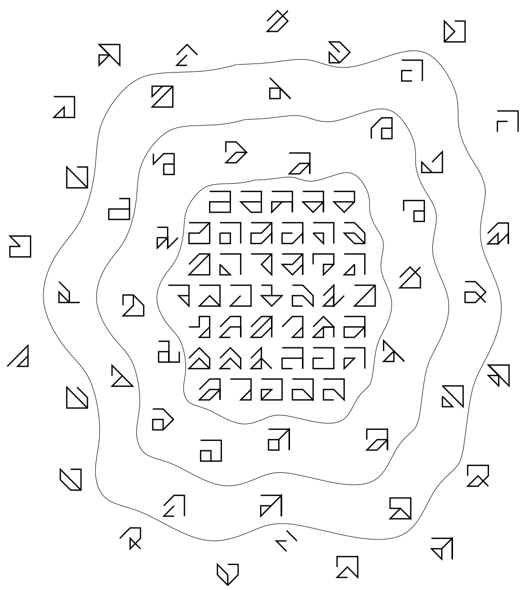

Figure 5: How the role model fares on 88 'a's.

model never categorized correctly. The more poorly a letter is categorized by the role model, the farther from the center of the picture it is placed. It is very interesting to note what the role model *can* do in terms of flexibly stretching its roles, and what the role model *can't* do. A detailed analysis of the model's behavior on this and other datasets may be found in (McGraw, 1995).

Overviews such as this provide a general feel for categorization in the role model. As can be seen from the letterforms in the center, the role model exhibits a remarkable flexibility in its categorization, making its notion of 'a'-ness very rich. Generally speaking, letterforms that are never properly recognized (those shown at the outskirts of the picture) are indeed letterforms at the "fringes of the category **a**", although there are a few exceptions that, by disappointing us, keep life lively.

## Comparison with human data

A more objective measure of the role model's performance may be had by comparing its recognition data with human gridfont recognition data over the same dataset. The series of psychology experiments from which our comparative data is drawn can only briefly be described here. However, they are fully explained in (McGraw et al., 1994) and (McGraw, 1995). The idea is very basic: 35 subjects were presented with a series of grid-bound letterforms one at a time on a macintosh monitor and asked to identify the letters as quickly and accurately as possible. The letterforms were presented as darkened line segments on a lighter grid such that the grid provided some degree of noise. The computer tracked both reaction time and accuracy data. Error-making, which often provides insight into the behavior of complex cognitive systems, was carefully tracked. Predictions based on the role hypothesis were confirmed during analysis reported in (McGraw et al., 1994). In addition to this simple experiment, another study (alluded to on the first page) showed that given letterforms as stimuli, people prefer parts that correspond to roles over parts that are supposedly "better" according to Palmer's rules. The study also included as a control non-letter stimuli made of flipped and inverted letterforms in which the "role-based" parts were not prefered over the Palmer parts.

The entire dataset reported here, called PSYCH, is made up of 544 tokens coming from all 26 letter categories. The full dataset can be split into two subsets: NORMALS (388 relatively strong letters) and FONTS (156 letters ranging from somewhat stylized to completely eccentric). Division of the dataset was initially completed by a human letterform expert. This division was confirmed empirically through *post facto* analysis.

Table 1 shows accuracy values of humans and the role model on the PSYCH dataset. The values were computed by averaging the correct-response percentages of each letter category.

| Dataset | Humans | Role model |
|---|---|---|
| PSYCH | 80.1 | 76.6 |
| NORMALS | 85.0 | 93.8 |
| FONTS | 65.4 | 51.4 |

Table 1: Accuracy percentages.

Large 544×26-entry confusion matrices can be built for both the human data and the role model. These matrices differ considerably from others like them in the psychology literature since they include data about a *large variety* of lowercase alphabetic styles instead of just one. Even though the dataset is made up of arguably idiosyncratic gridletters, our stimuli are more realistic than many past datasets, in the sense that they capture more of the natural variability found among letterforms. This leads to a more thorough treatment of errors than has been evident in past work.

| Dataset | r-value | points |
|---|---|---|
| PSYCH | .8872 | 10764 |
| NORMALS | .9511 | 6708 |
| FONTS | .7274 | 4056 |

Table 2: Token-level correlation of role model and human error matrices. (All correlations are significant $p > 0.0001$.)

Table 2 shows the correlation values of the *token-level* confusion matrix of the role model against the human confusion matrix over our large and varied dataset.[3] Since error-making tends to highlight the type of processing that a cognitive system is doing, it is important to take errors into account during correlation. Favorable comparisons of these correlation values with those of simple connectionist models and a brute-force symbolic model can be found in (McGraw, 1995). Unfortunately, space constraints do not allow us to introduce those results or comparisons here. In-depth analysis of

[3] Category-level correlations are routinely higher (*e.g.*, the category-level correlation over PSYCH is 0.9821), but offer less resolution for inter-model comparison of the sort in (McGraw, 1995).

particular trends in error-making are also analyzed in (McGraw, 1995) and corroborate the general claims regarding the role-model's strength.

## Conclusions

Our fully-implemented role model provides one possible implementation of the role hypothesis. It has proven to be a very strong model of letter recognition, as it clearly explains much of human letter-recognition behavior (especially with regard to error-making). Among the critical portions of the recognition process in the role model are these: building preliminary parts in a bottom-up fashion from low-level data, adjusting parts under top-down influence from roles, noting aspects of style — norm violations — that result from filling roles with particular parts, evaluating prospective filled role-sets as matches of different letter categories, tracking the strengths of activation of competing categories, and searching for new perceptual parsings if initial attempts at recognition result in only weak categorization.

Unlike many recognition programs, the role model does not simply perform very well on one style, only to crash and burn on others. It was designed specifically to handle a huge variety of styles. Tests using large datasets show that we have captured at least some of the perceptual fluidity exhibited by people. The role model is able to recognize both standard and stylistically eccentric letterforms.

A distinct advantage that the role model holds over simpler models lies in the nature of its output. The role model returns not only the usual category label, but also a parsing of a letterform in terms of its constituent parts. Part-level parsings correspond to role-level conceptualizations of letters. Further processing — for instance, the extraction of style information, or the analogical design of other letterforms — is possible only with the kind of structural information that the role model provides. The ability to parse a letterform into natural parts corresponding to roles is of critical importance to Letter Spirit's capability to design stylistically-consistent alphabets.

## References


Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review*, 94(2), 115–147.

Blesser, B., Shillman, R., Cox, C., Kuklinski, T., Ventura, J., and Eden, M. (1973). Character recognition based on phenomenological attributes. *Visible Language*, 7(3), 209–223.

Erman, L., Hayes-Roth, F., Lesser, V., Raj Reddy, D. (1980). The Hearsay II speech-understanding system: Integrating knowledge to resolve uncertainty. *Computing Surveys*, 12(2), 213–253.

Gaillat, G. and Berthod, M. (1979). Panorama des techniques d'extraction de traits caractéristiques en lecture optique des caractères. *Revue Technique Thomson-CSF*, 11(4), 943–959.

Hock, H., Tromley, C., and Polmann, L. (1988). Perceptual units in the acquisition of visual categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 75–84.

Hofstadter, D. and FARG (1995). *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought.* New York: Basic Books. "FARG" is an acronym for the "Fluid Analogies Research Group".

Hofstadter, D. and McGraw, G. (1995). Letter Spirit: Esthetic perception and creative play in the rich microcosm of the alphabet. In Hofstadter, D. and FARG, *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought* (pp. 407–466). New York: Basic Books.

Indurkhya, B. (1992). *Metaphor and Cognition: An Interactionist Approach.* Norwell, MA: Kluwer.

McGraw, G. (1995). Letter Spirit (part one): Emergent High-Level Perception of Letters Using Fluid Concepts. Doctoral dissertation. Bloomington, IN: Indiana University, CRCC. Available on the web http://www.cogsci.indiana.edu/farg/mcgrawg/thesis.html.

McGraw, G. and Hofstadter, D. (1993). Perception and creation of diverse alphabetic style. *Artificial Intelligence and Simulation of Behaviour Quarterly*, 85, 42–49.

McGraw, G., Rehling, J., and Goldstone, R. (1994). Letter perception: Toward a conceptual approach. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 613–618). Hillsdale, NJ: Lawrence Erlbaum Associates.

Mitchell, M. (1993). *Analogy-making as Perception.* Cambridge, MA: MIT Press/Bradford Books.

Mori, S., Yamamoto, K., and Yasuda, M. (1984). Research on machine recognition of handprinted characters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(4), 386–405.

Naus, M. and Shillman, R. (1976). Why a Y is not a V: A new look at the distinctive features of letters. *Journal of Experimental Psychology: Human Perception and Performance*, 2(3), 394–400.

Palmer, S. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, 9, 441–474.

Palmer, S. (1978). Structural aspects of visual similarity. *Memory & Cognition*, 6(2), 91–97.

Rich, E. (1983). *Artificial Intelligence.* New York: McGraw-Hill.

Sanocki, T. (1986). Visual Knowledge Underlying Letter Perception: A Font-Specific Approach. Doctoral dissertation. Madison, WI: University of Wisconsin, Department of Psychology.

Sanocki, T. (1991). Effects of early common features on form perception. *Perception & Psychophysics*, 50(5), 490–497.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioural and Brain Sciences*, 11, 1–74.

Treisman, A. and Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(12), 97–136.