Letter Spirit: An Architecture for Creativity in a Microdomain

Gary McGraw and Douglas Hofstadter

Istituto per la Ricerca Scientifica e Tecnologica
Loc. Pantè di Povo, I38100 Trento, Italia

Center for Research on Concepts and Cognition, Indiana University
510 North Fess Street, Bloomington, Indiana 47408, USA

gem@cogsci.indiana.edu dughof@cogsci.indiana.edu

Abstract. The Letter Spirit project explores the creative act of artistic letter-design. The aim is to model how the 26 lowercase letters of the roman alphabet can be rendered in many different but internally coherent styles. Viewed from a distance, the behavior of the program can be seen to result from the interaction of four emergent agents working together to form a coherent style and to design a complete alphabet: the Imaginer (which plays with the concepts behind letterforms), the Drafter (which converts ideas for letterforms into graphical realizations), the Examiner (which combines bottom-up and top-down processing to perceive and categorize letterforms), and the Adjudicator (which perceives and dynamically builds a representation of the evolving style). Creating a gridfont is an iterative process of guesswork and evaluation carried out by the four agents. This process is the "central feedback loop of creativity". Implementation of Letter Spirit is just beginning. This paper outlines our goals and plans for the project.

1 The Motivation of 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. Implementation of Letter Spirit is just beginning. This paper outlines our goals and plans for the project. The aim is to model the process of rendering the 26 lowercase letters of the roman alphabet in many different, internally 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 1 illustrates the relationship of these two ideas. Starting with one or more seed letters representing the beginnings of a style, the program will attempt to create the rest of the alphabet in such a way that all 26 letters share the same style, or *spirit.*

Letters in the domain are formed exclusively from straight segments on a 3 × 7 grid (see Figure 2) in order to make decisions smaller in number and

more discrete. 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.

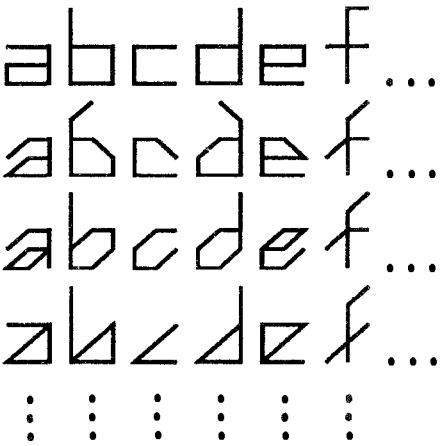

Fig. 1. Items in any column have *letter* in common. Items in any row have *spirit* in common. (Also see Fig. 3.)

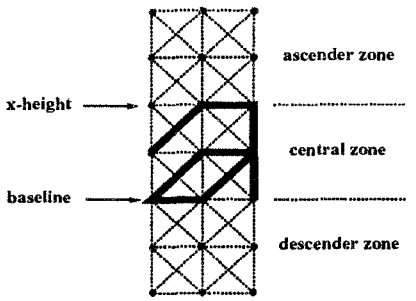

Fig. 2. The Letter Spirit grid, with one of the many possible sets of quanta instantiating an 'a' turned on.

While at first glance, the Letter Spirit domain might be shrugged off as a "toy domain", this would grossly underestimate its subtlety. In spite of, or rather *because* of, the reduction to the grid, the Letter Spirit challenge is, in terms of cognitive-science issues, extremely rich. The cognitive issues are magnified, not reduced, by the act of simplifying the domain. All that has been thrown out is the need for expertise. One need not be a professional typeface designer or lifelong student of letterforms to appreciate the consistency of a well-designed gridfont.

Even a novice can design a passable gridfont, though creating a sophisticated one is very difficult.

2 Letters as Concepts

In order to better distinguish the *concept* of a letter from various geometric shapes that may instantiate it, we introduce some terminology. We distinguish three conceptual levels, running from abstract to nearly concrete as they move toward the actual geometric letterform. The term *letter-concept* refers to the most abstract idea for drawing a letter without reference to style. This level is comprised of a set of *letter-conceptualizations*. A typical letter-conceptualization would be the notion that a 'b' consists of two "roles" — a *post* on the left side attached in two places to an open *bowl* on the right side, sitting on the baseline. A rival conceptualization for the same letter also consists of two roles — a *post* on the left side attached in one place to a closed *loop* on the right side, sitting on the baseline. These conceptualizations, possibly augmented by others, constitute the *letter-concept* of 'b'. Once a specific letter-conceptualization has been chosen, notions of style give rise to a more specific and detailed letter-conceptualization that partially specifies *how each role should be realized* (of course this conceptualization still could be realized in infinitely many ways). This is called a *letter-plan*. A letter-plan is present in a designer's mind before any marks are put on paper. The actual shape drawn on paper is a *letterform*. Letter Spirit is concerned with all these levels: play with letter-conceptualizations, creation of letter-plans, and the design of letterforms based on letter-plans.

A vivid example of the shape/concept distinction involves lowercase 'x'. For most US-educated adults, the only conceptualization for 'x' consists of a forward slash and a backward slash of the same size that cross somewhere near the middle. English children, by contrast, are taught to draw a lowercase cursive 'x' as a pair of small crescents facing away from each other but "kissing" in the middle. If we look at a printed 'x' in this way, we are suddenly struck by this new conceptualization. The shape on our retina is the same, but what is constructed in our mind's eye is very different.

The conceptual pieces into which a letter is broken in the mind's eye are its *roles*. For example, the two crossing slashes in an imagined 'x' are roles. So also are their four tips, and the crossing-point in the middle. Each role has a different degree of *importance* to the letter — the degree to which its presence or absence matters. Of course, different shapes instantiate a given role more strongly or more weakly than others. In other words, roles are also concepts with somewhat nebulous boundaries, just as *wholes* (complete letters) are. The difference is, membership in a role is easier to characterize than in a whole, so that reducing wholes to collections of interacting roles is a step forward in simplification.

The internal structure of a category is represented as a collection of interacting roles. Category membership at the whole-letter level is partially determined by category membership at the lower level of roles. In addition, *stylistic* appropriateness of a shape is judged in terms of *how roles are filled* — in other words, *how norms are violated*. Any such violation is a stylistic hallmark that must be propagated (via analogy) to other letters.

3 A Sketch of the Architecture

Letter Spirit is motivated by the belief that creativity is an automatic outcome of the existence of flexible and context-sensitive concepts. Its architecture is based on the principles of emergent computation, wherein complex high-level behavior emerges as a statistical consequence of the bottom-up cooperation of many small computational actions influenced by many dynamically changing top-down conceptual pressures. Micro-agents known as "codelets" build and destroy perceptual structures in a nondeterministic parallel manner, guided throughout by letter-concepts[1].

The Letter Spirit program will contain four dynamic memories, each concerned with different levels of concreteness and abstraction of shapes (and concepts pertaining to shapes). These memories are:

- the **Scratchpad**, which is *a virtual piece of paper* on which all the letters of a font are drawn and modified; as such it is more a type of external memory than an aspect of mental activity;
- the **Visual Focus**, which is *the site where perception of a given letterform occurs*; in it, perceptual structures are built up and converge to stable categorical and stylistic interpretations;
- the **Thematic Focus**, which is the program's *dynamically changing set of ideas about the stylistic essence of the gridfont under way*; in it are recorded stylistic observations of all sorts concerning letters already designed, and if and when some of these observations are perceived as falling into patterns, those patterns can be taken as determinant of the style, meaning they can be elevated to the status of explicit *themes* — ideas that play an active role in guiding further design decisions, in the sense of serving as "pressures" on the construction of further letters;
- the **Conceptual Memory**, which is the program's *locus of permanent knowledge and understanding of its domain*, and which, for each concept, has three facets: (1) a set of *category-membership criteria*, which specify the recognition requirements for instances of the concept in terms of more primitive concepts; (2) a set of *explicit norms*, which encode aspects of the concept's "core"; and (3) an *associative halo*, consisting of links having time-varying lengths connecting the concept with related concepts, thus giving a sense of where the concept is located in "conceptual space" by saying what it most resembles.

Viewed from a distance, the behavior of the program can be thought of as resulting from the interaction of just four large-scale emergent agents working together to form a coherent style and to design a complete alphabet. These four conceptually separable types of large-scale activities emerge from the many small actions of codelets:

1. the high-level conceptual activity of *devising a new letter-plan* (*i.e.*, either an idea for an as-yet undesigned letter or a possibility for improving an already-designed letter);

[1] See [Mitchell, 1993] for more about codelets, pressures, and emergent computation.

2. the intermediary activity of *translating a new letter-plan into a concrete shape* on the Scratchpad;
3. the relatively concrete perceptual activity of *examining a newly-drawn shape and categorizing it* (*i.e.*, deciding which letter of the alphabet it is, and how unambiguously so);
4. the more abstract perceptual activity of *recognizing the stylistic attributes of a newly-drawn letter, and judging them* (*i.e.*, finding "exportable" ways of describing how a given letterform violates norms, and deciding how well the letter's attributes fit with those of other letters in the developing gridfont).

It is convenient to speak as if these emergent activities were carried out by four explicit and cleanly separable modules, together comprising the totality of the program. (These agents could be likened to the agents referred to in [Minsky, 1985].) We call these hypothetical agents the *Imaginer*, the *Drafter*, the *Examiner*, and the *Adjudicator*. It must be borne in mind that these modules are in some sense convenient descriptive fictions, in that each is simply an emergent by-product of the actions of many codelets, and their activities are so intertwined that they cannot be disentangled in a clean way.

The interaction of these four agents, whereby ideas are suggested, critiqued, revised, possibly abandoned and regenerated, and so on, jibes with our intuitive sense of what human creativity really is. It seems to us fair to say that this kind of emergent, unpredictable processing constitutes a program's *making its own decisions*. This continuous process of suggestion and revision makes up the "central feedback loop of creativity". The full realization of the Letter Spirit program will, we believe, shed significant light on the mechanisms of human creativity.

References

[Hofstadter, 1985] Hofstadter, D. (1985). *Metamagical Themas: Questing for the Essence of Mind and Pattern*. Basic Books, New York. See especially Chapters 10, 12, 13, 23, 24, and 26.

[Hofstadter, 1987] Hofstadter, D. (1987). Introduction to the letter spirit project and to the idea of "gridfonts". Technical Report 17, Center for Research on Concepts and Cognition, 510 N. Fess, Bloomington, IN, 47408.

[Hofstadter and McGraw, 1993] Hofstadter, D. and McGraw, G. (1993). Letter spirit: An emergent model of the perception and creation of alphabetic style. Technical Report 68, Center for Research on Concepts and Cognition, 510 N. Fess, Bloomington, IN 47405.

[Minsky, 1985] Minsky, M. (1985). *The Society of Mind*. Simon and Schuster, New York.

[Mitchell, 1990] Mitchell, M. (1990). *Copycat: A Computer Model of High-level Perception and Conceptual Slippage in Analogy Making*. PhD thesis, University of Michigan, Ann Arbor, Michigan.

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

Standard Square

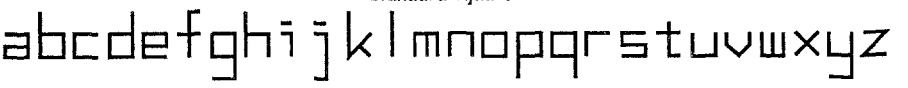

Double Backslash

Hint Four

Zigzag

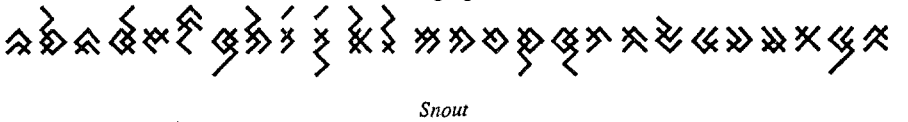

Snout

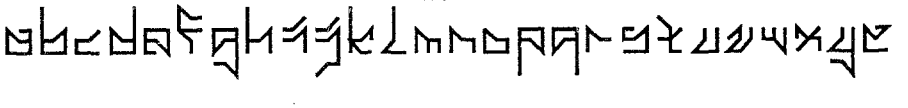

Bowtie

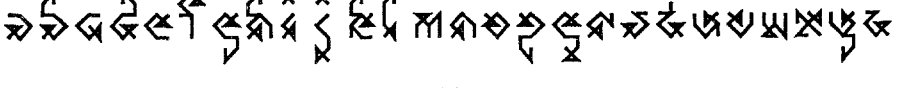

Weird Arrow

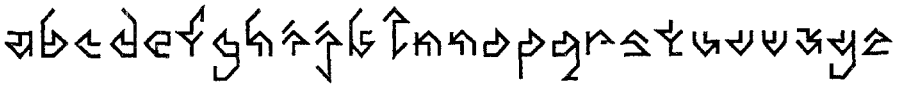

Sabretooth

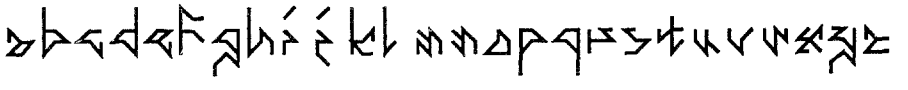

Sluice

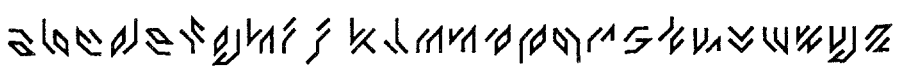

Flournoy Ranch

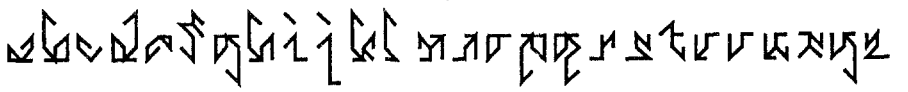

Fig. 3. This figure, like Fig. 1, can be thought of as illustrating both the vertical problem (categorical sameness) and the horizontal problem (stylistic sameness) for which Letter Spirit is named. All of these gridfonts were designed by people.