

Metafont, Metamathematics, and Metaphysics

Comments on Donald Knuth's Article "The Concept of a Meta-Font"

Douglas R. Hofstadter

It is argued that readers are likely to carry away from Donald Knuth's article "The Concept of a Meta-Font" a falsely optimistic view of the extent to which the design of typefaces and letterforms can be mechanized through an approach depending on describing letterforms by specifying the settings of a large number of parameters. Through a comparison to mathematical logic, it is argued that no such set of parameters can capture the essence of any semantic category. Some different ways of thinking about the problem of the "spirit" residing behind any letterform are suggested, connecting to current research issues in the field of artificial intelligence.

The "Mathematization of Categories" and Metamathematics

Donald Knuth has spent the past several years working on a system allowing him to control many aspects of the design of his forthcoming books—from the typesetting and layout down to the very shapes of the letters! Never has an author had anything remotely like this power to control the final appearance of his or her work. Knuth's TEX typesetting system has become well-known and available in many countries around the world. By contrast, his Metafont system for designing families of typefaces has not become as well known or as available.

In his article "The Concept of a Meta-font" [Knuth 82], Knuth sets forth for the first time the underlying philosophy of Metafont, as well as some of its products. Not only is the concept exciting and clearly well executed, but in my opinion the article is charmingly written as well. However, despite my overall enthusiasm for Knuth's idea and article, there are some points in it that I feel might be taken wrongly by many readers, and since they are points that touch close to my deepest interests in artificial intelligence and esthetic theory, I felt compelled to make some comments to clarify certain important issues raised by "The Concept of a Meta-font".

Although his article is primarily about letterforms, not philosophy, Knuth holds out in it a philosophically tantalizing prospect for us: that with the arrival of computers, we can now approach the vision of a unification of all typefaces. This can be broken down into two ideas: (1) that underneath all "A" 's there is just one grand, ultimate abstraction

that can be captured in a finitely parametrizable computational structure—a “software machine” with a finite number of “tunable knobs” (we could say “degrees of freedom” or “parameters,” if we wished to be more dignified); and (2) that every conceivable particular “A” is just a product of this machine with its knobs set at specific values.

Beyond the world of letterforms, Knuth’s vision extends to what I shall call the “mathematization of categories”: the idea that any abstraction or Platonic concept can be so captured—i.e., as a software machine with a finite number of knobs. (For more on this notion, see [Hofstadter 82b].) Knuth gives only a couple of examples—those of the “meta-waltz” and the “meta-shoe”—but by implication one can imagine a “meta-chair,” a “meta-person,” and so forth.

This is perhaps carrying Knuth’s vision further than he ever intended. Indeed, I suspect so; I doubt that Knuth believes in the feasibility of such a “mathematization of categories” opened up by computers. Yet any imaginative reader would be likely to draw hints of such a notion out of Knuth’s article, whether Knuth intended it that way or not. It is my purpose in this article to argue that such a vision is exceedingly unlikely to come about, and that such intriguingly flexible tools as meta-shoes, meta-fonts, modern electronic organs (with their “oom-pah-pah” and “cha-cha-cha” rhythms and their canned harmonic patterns), and other many-knobbed devices will only help us see more clearly why this is so. The essential reason for this I can state in a very short way: I feel that to fill out the full “space” defined by a category such as “chair” or “A” or “waltz” is an act of infinite creativity, and that no finite entity (inanimate mechanism or animate organism) will ever be capable of producing all possible “A” ’s and nothing but “A” ’s (the same could be said for chairs, waltzes, etc.).

I am not making the trivial claim that, because life is finite, nobody can make an infinite number of creations; I am making the nontrivial claim that nobody can possess the “secret recipe” from which all the (infinitely many) members of a category such as “A” can in theory be generated. In fact, my claim is that no such recipe exists. Another way of saying this is that even if you were granted an infinite lifetime in which to draw all the “A” ’s you could think up, thus realizing the full potential of any recipe you had, no matter how great it might be, you would still miss vast portions of the space of “A” ’s.

In metamathematical terms this amounts to positing that any conceptual (or “semantic”) category is a “productive” set, a precise notion whose characterization is a formal counterpart to the description in the previous paragraphs (namely, a set whose elements cannot be totally enumerated by any effective procedure without overstepping the bounds of that set, but which can be approximated more and more fully by a sequence of increas-

ingly complex effective procedures). The existence and properties of such sets first became known as a result of Gödel's Incompleteness Theorem of 1931 [Gödel 31]. It is certainly not my purpose here to explain this famous result, but a short synopsis might be of help. (Other useful references are: [Chaitin 75], [DeLong 70], [Hofstadter 79], [Nagel 58], [Rucker 82], [Smullyan 61], [Smullyan 78].)

An Intuitive Picture of Gödel's Theorem

Gödel was investigating the properties of purely formal deductive systems in the sphere of mathematics, and he discovered that such systems—even if their ostensible domain of discourse was limited to one topic—could be viewed as talking “in code” about themselves. Thus a deductive system could express, in its own formal language, statements about its own capabilities and weaknesses. In particular, System X could say of itself through the Gödelian code, “System X is not powerful enough to demonstrate the truth of Sentence S.” It sounds a little bit like a science-fiction robot called “Robot 15” droning in a telegraphic monotone, “Robot-15 unfortunately unable to complete Task T-12—very sorry.” Now what happens if Task T-12 happens, by some crazy coincidence, to be not the assembly of some strange cosmic device but merely the act of uttering the preceding telegraphic monotone? (I say “merely” but of course that is a bit ironic.) Then Robot-15 could get only partway through the sentence before choking: “Robot-15 unfortunately unable to comp--.”

Now in the case of a formal system, System X, talking about its powers, suppose that Sentence G, by an equally crazy coincidence, is the one that says, “System X is regrettably not powerful enough to demonstrate the truth of Sentence G.” In such a case, Sentence G is seen to be an assertion of its own unprovability within System X. In fact we do not have to rely on crazy coincidences, for Gödel showed that given any reasonable formal system, a G-type sentence for that system actually exists. (The only exaggeration in my English-language version of G is that in formal systems there is no way to say “regrettably.”) In formal deductive systems this foldback takes place of necessity by means of a Gödelian code, but in English no Gödelian code is needed and the peculiar quality of such a loop is immediately visible.

If you think carefully about Sentence G, you will discover some amazing things. Could Sentence G be provable in System X? If it were, then System X would contain a proof for Sentence G, which asserts that System X contains no proof for Sentence G. Only if System X is blatantly self-contradictory could this happen—and a formal reasoning system that is self-contradictory is no more useful than a submarine with screen doors. So, provided we are dealing with a consistent formal system (one with no self-con-

traditions), then Sentence G is not provable inside System X. And since this is precisely the claim of Sentence G itself, we conclude that Sentence G is true--true but unprovable inside System X.

One last way to understand this curious state of affairs is afforded the reader by this small puzzle. Choose the more accurate of the following pair of sentences:

(1) Sentence G is true *despite* being unprovable.

(2) Sentence G is true *because* it is unprovable.

You'll know you've really caught on to Gödel when both versions ring equally true to your ears, when you flip back and forth between them, savoring that exceedingly close approach to paradox that G affords. That's how twisted back on itself Sentence G is!

The main consequence of G's existence within each System X is that there are truths unattainable within System X, no matter how powerful and flexible System X is, as long as System X is not self-contradictory. Thus, if we look at truths as objects of desire, no formal system can have them all; in fact, given any formal system we can produce on demand a truth that it cannot have, and flaunt that truth in front of it with taunting cries of "Nah, nah!" The set of truths has this peculiar and infuriating quality of being uncapturable by any finite system, and worse, given any candidate system, we can use what we know about that system to come up with a specific Gödelian truth that eludes provability inside that system.

By adding that truth to the given system, we come up with an enlarged and slightly more powerful system--yet this system will be no less vulnerable to the Gödelian devilry than its predecessor was. Imagine a dike that springs a new leak each time the proverbial Dutch boy plugs up a hole with his finger. Even if he had an infinite number of fingers, that leaky dike would find a spot he hadn't covered. A system that contains at least one unprovable truth is said to be "incomplete," and a system that not only contains such truths but that cannot be rescued in any way from the fate of incompleteness is said to be "essentially incomplete." Another name for sets with this wonderfully perverse property is "productive" [Rogers 67].

My claim--that semantic categories are productive sets--is, to be sure, not a mathematically provable fact but a metaphor. This metaphor has been used by others before me--notably, the logicians Emil Post and John Myhill--and I have written of it myself before (see [Post 44], [Myhill 52], [Hofstadter 79], and [Hofstadter 82a]).

Completeness and Consistency

Note that it is important to have the potential to fill out the full (infinite) space, and equally important not to overstep it. However, merely having

infinite potential is not by any means equivalent to filling out the full space. After all, any existing Metafont “A”-schema--even one having just one degree of freedom!--will obviously give us infinitely many distinct “A” ’s as we sweep its knob (or knobs) from one end of the spectrum to the other. Thus to have an “A”-making machine with infinite variety of potential output is not in itself difficult; the trick is to achieve *completeness*: to fill the space.

And yet, isn’t it easy to fill the space? Can’t one easily make a program that will produce all possible “A” ’s? After all, any “A” can be represented as a pattern of pixels (dots that are either off or on) in an $m \times n$ matrix--hence a program that merely prints out all possible combinations of pixels in matrices of all sizes (starting with 1×1 and moving upwards to 2×1 , 1×2 , 3×1 , 2×2 , 1×3 , etc., as in Georg Cantor’s famous enumeration of the rational numbers) will certainly cover any given “A” eventually. This is quite true. So what’s the catch?

Well, unfortunately, it is hard--very hard--to write a screening program that will retain all the “A” ’s in the output of this pixel-pattern program, and at the same time will reject all “K” ’s, pictures of frogs, octopi, grandmothers, and precognitive photographs of traffic accidents in the twenty-fifth century (to mention just a few of the potential outputs of the generation program). The requirement that one must stay within the bounds of a conceptual category could be called *consistency*--a constraint complementary to that of completeness.

In summary, what might seem desirable from a knobbed category-machine is the joint attainment of two properties--namely, (1) completeness: that all true members of a category (such as the category of “A” ’s [Figure 1] or the category of human faces [Figure 2] should be potentially producible eventually as output; and (2) consistency: that no false members of the category (“impostors”) should ever be potentially producible. In short, that the set of outputs of the machine should coincide exactly with the set of members of the intuitive category.

The twin requirements of consistency and completeness are metaphorical equivalents of well-known notions by the same names in metamathematics, denoting desirable properties of formal systems (theorem-producing machines)--namely, (1) completeness: that all true statements of a theory (such as the theory of numbers or the theory of sets) should be potentially producible eventually as theorems; and (2) consistency: that no false statements of the theory should ever be potentially producible. In short, that the set of theorems of the formal system should coincide exactly with the set of truths of the informal theory.

The import of Gödel’s Incompleteness Theorem is that these two idealized goals are unreachable simultaneously for any “interesting” theory (where “interesting” really means “sufficiently complex”); nonetheless, one can



Figure 1. The category of "A"'s (drawn from [Letraset 81]).



Figure 2. The category of human faces (drawn from [Strich 81]).

approach the set of truths by stages, using increasingly powerful formal systems to make increasingly accurate approximations. The goal of total and pure truth is, however, as unreachable by formal methods as is the speed of light by any material object. I suggest that a parallel statement holds for any “interesting” category (where again, “interesting” means something like “sufficiently complex,” although it is a little harder to pin down): namely, one can do no better than approach the set of its members by stages, using increasingly powerful knobbed machines to make increasingly accurate approximations.

Intuition at first suggests that there is a crucial difference between the (metamathematical) result about the nonformalizability of truth and the (metaphorical) claim about the nonmechanizability of semantic categories; this difference would be that the set of all truths in a mathematical domain such as set theory or number theory is objective and eternal, whereas the set of all “A”’s is subjective and ephemeral. However, on closer examination, this distinction begins to blur quite a bit. The very fact of Gödel’s proven nonformalizability of mathematical truth casts serious doubt on the objective nature of such truth. Just as one can find all sorts of borderline examples of “A”-ness, examples that make one sense the hopelessness of trying to draw the concept’s exact boundaries, so one can find all sorts of borderline mathematical statements that are formally undecidable in standard systems and which, even to a keen mathematical intuition, hover between truth and falsity. And it is a well-known fact that different mathematicians hold different opinions about the truth or falsity of various famous formally undecidable propositions (the axiom of choice in set theory is a classic example). Thus, somewhat counterintuitively, it turns out that mathematical truth has no fixed and eternal boundaries, either. And this suggests that perhaps my metaphor is not so much off the mark.

A Misleading Claim for Metafont

Whatever the validity and usefulness of this metaphor, I shall now try to show some evidence for the viewpoint that leads to it, using Metafont as a prime example of a “knobbed category machine.” In his article, Knuth comes perilously close, in one throwaway sentence, to suggesting that he sees Metafont as providing us with a mathematization of categories. I doubt he suspected that anyone would focus in on that sentence as if it were the key sentence of the article—but as he did write it, it’s fair game! That sentence ran:

The ability to manipulate lots of parameters may be interesting and fun, but does anybody really need a 6 1/7-point font that is one fourth of the way between Baskerville and Helvetica?

This rhetorical question is fraught with unspoken implications. It suggests that Metafont as it now stands (or in some soon-available or slightly modified version) is ready to carry out, on demand, for any user, such an interpolation between two given typefaces. There is something very tricky about this proposition that I suspect most readers will not notice: it is the idea that jointly parametrizing two typefaces is no harder, no different in principle, from just parametrizing one typeface in isolation.

Indeed, to many readers, it would appear that Knuth has actually carried out such a joint parametrization. After all, in printing Psalm 23 [Figure 3] didn't he move from an old-fashioned, compact, serified face with

Figure 3. Donald Knuth's virtuoso Metafont rendition of Psalm 23, in which the font for each character is determined by the settings of 28 knobs, all of which change slowly but steadily as the psalm progresses.

The LORD is my shepherd;
I shall not want.
He maketh me to lie down
in green pastures:
he leadeth me
beside the still waters.
He restoreth my soul:
he leadeth me
in the paths of righteousness
for his name's sake.
Yea, though I walk through the valley
of the shadow of death,
I will fear no evil:
for thou art with me;
thy rod and thy staff
they comfort me.
Thou preparest a table before me
in the presence of mine enemies:
thou anointest my head with oil,
my cup runneth over.
Surely goodness and mercy
shall follow me
all the days of my life:
and I will dwell
in the house of the LORD
for ever.

relatively tall ascenders and descenders and small x-height all the way to the other end of the spectrum: a modern-looking, extended, sans-serif face with relatively short ascenders and descenders and large x-height? Yes, of course--but the critical omitted point here is that these two ends of the spectrum were not pre-existing, prespecified targets; they just happened to emerge as the extreme products of a knobbed machine designed so that one more or less intermediate setting of its knobs would yield a particular target typeface (Monotype Modern Extended 8A).

In other words, this particular set of knobs was inspired solely and directly by an attempt to parametrize one typeface (Monotype Modern). The two extremes shown in the psalm are both variations on that single theme; the same can be said of every intermediate stage as well. There is only one underlying theme (Monotype Modern) here, and a cluster of several hundred variants of it, each one of which is represented by a single character. The psalm does not represent the marriage of two unrelated families, but simply exhibits many members of one large family.

Joint Parametrization of Two Typefaces:

A Far Cry from Parametrizing One Typeface

You can envision all the variants of Monotype Modern produced by twiddling the knobs on this particular machine as constituting an "electron cloud" surrounding a single "nucleus" [Figure 4a]. Now by contrast, joint parametrization of two pre-existent, known typefaces (say, Baskerville and Helvetica, as Knuth suggests [Figure 5] would be like a cloud of electrons swarming around two nuclei, like a chemical bond [Figure 4b].

In order to jointly parametrize two typefaces in Metafont, you would need to find, for each pair of corresponding letters (say Baskerville "a" and Helvetica "a") a set of discrete geometric features (line segments, serifs, extremal points, points of curvature shift, etc.) that they share and that totally characterize them. Each such feature must be equated with one or more parameters (knobs), so that the two letterforms are seen as produced by specific settings of their shared set of knobs. Moreover, all intermediate settings must also yield valid instances of the letter "a". That is the very essence of the notion of a knobbed machine, and it is also the gist of the quote, of course: that we should now (or soon) be able to interpolate between any familiar typefaces merely by knob-twiddling.

Now I will admit that I think it is perhaps feasible--though much more difficult than parametrizing a single typeface--to jointly parametrize two typefaces that are not radically different. It is not trivial, to cite just one sample difficulty, to move between Baskerville's round dot over the "i" to Helvetica's square dot--but it is certainly not inconceivable. Conversely, it is not inconceivable to move between the elegant swash tail of the Basker-




Figure 4. (a) An electron cloud surrounding a single nucleus; (b) A cloud of electrons around two nuclei, like a chemical bond.

abcdefghijklmnopqrstuvwxyz
 ABCDEFGHIJKLMNOPQRSTUVWXYZ
 abcdefghijklmnopqrstuvwxyz ß
 ABCDEFGHIJKLMNOPQRSTUVWXYZ

Figure 5. Baskerville (above) and Helvetica.

i Q
 i Q

Figure 6. Blow-ups of Baskerville and Helvetica “i” and “Q”.

ville “Q” and the stubby straight tail of the Helvetica “Q”—but it is certainly not trivial [Figure 6].

Moving from letter to letter and comparing them will reveal that each of these two typefaces has features that the other totally lacks. (You should disregard lowercase “g”, since the “g”’s of our two typefaces are as different from each other as Baskerville “B” is from Helvetica “H”; in both cases the two letterforms being compared derive from entirely different underlying “Platonic essences.” (It is Metafont’s purpose to mediate between different stylistic renditions of a single Platonic essence, not between distinct Platonic essences.) Presumably, in a case where one typeface possesses some distinct feature that the other totally lacks, there is a way to fiddle with the knobs that will make the feature nonexistent in one but present in the other. For instance, a knob setting of zero might make some feature totally vanish. Sometimes it will be harder to make features disappear—it might require several knobs to have coordinated settings. Nonetheless, des-

pite all the complex ways that Baskerville and Helvetica differ, I repeat, it is conceivable that somebody with great patience and ingenuity could jointly parametrize Helvetica and Baskerville. But the real question is this: Would such a joint parametrization easily emerge out of two separate, independently carried-out parametrizations of these typefaces? Hardly!

The Baskerville knobs do not contain even a hint of the Helvetica qualities--or the reverse. How can I convince you of this? Well, just imagine how great the genius of John Baskerville would have had to be for his design to have implicitly defined another typeface--and a typeface only discovered (or invented) two centuries later! To see this more concretely, imagine that someone who had never seen Helvetica naively created a Metafont rendition of Baskerville (that is, a metafont centered on Baskerville in the same sense as Knuth's sample metafont is centered on Monotype Modern). Now imagine that someone else who does know Helvetica comes along, twiddles the knobs of this Baskerville metafont, and actually produces a perfect Helvetica! It would be nearly as strange as having a marvelous music-composing program based exclusively on the style of G. F. Handel (who composed in England in a baroque, elegant 18th-century style) that was later discovered, totally unexpectedly, to produce many pieces indistinguishable in style from the music of Ernest Bloch (who composed in Switzerland in a sparse, crisp 20th-century style) when various melodic, harmonic, and rhythmic parameters were twiddled. To me, this is simply inconceivable.

Interpolating between an Arbitrary Pair of Typefaces

The worst is yet to come, however. Presumably Knuth did not wish us to take his rhetorical question in such a limited way as to imply that the numbers $6\frac{1}{7}$ and $1\frac{1}{4}$ were important. Pretty obviously, they were just examples of arbitrary parameter settings. Presumably, if Metafont could easily give you a $6\frac{1}{7}$ -point font that is $1/4$ of the way between Baskerville and Helvetica, it could as easily give you an $11\frac{2}{3}$ -point font that is $5/17$ of the way between Baskerville and Helvetica--and so on. And why need it be restricted to Baskerville and Helvetica? Surely those numbers weren't the only "soft" parts of the rhetorical question! Common sense tells us that Helvetica and Baskerville were also merely arbitrary choices of typeface. Thus the hidden implication is that, as easily as one can twiddle a dial to change point size, so one can twiddle another dial (or set of dials) and arrive at any desired typeface, be it Helvetica, Baskerville, or whatever. Knuth might just as easily have put it this way: "The ability to manipulate lots of parameters may be interesting and fun, but does anybody really need an X-point font that is Y percent of the way between typeface T1 and typeface T2?" For instance, we might have set the four knobs to the following settings:

X: 36
Y: 50%
T1: Magnificat
T2: Stop

Each of these two typefaces [Figure 7a, b] is ingenious, idiosyncratic, and visually intriguing. I challenge any reader to even imagine a blend halfway between them, let alone draw it! And to emphasize the flexibility implied by the question, how about trying to imagine a typeface that is (say) one third of the way between Cirkulus and Block Up [Figure 7c, d]? Or one that is somewhere between Explosion and Shatter [Figure 7e, f]?

Figure 7. (a) Magnificat, (b) Stop, (c) Cirkulus, (d) Block Up, (e) Explosion, (f) Shatter, and (g) Helvetica Medium Italic.



“A Posteriori” Knobs and the Frame Problem of AI

Shatter, incidentally, provides an excellent example of the trouble with viewing everything as coming from parameter settings. If you look carefully, you will see that Shatter is indeed a “variation on a theme,” the theme being Helvetica Medium Italic [Figure 7g]. But does that imply that any meticulous parametrization of Helvetica would automatically yield Shatter as one of its knob-settings? Of course not. That is absurd. No one in their right mind would anticipate such a variation while parametrizing Helvetica, just as no one in their right mind when delivering their Nobel Lecture would say, “Thank you for awarding me my first Nobel Prize.” When someone wins a Nobel Prize, they do not immediately begin counting how many they have won. Of course, if they win two, then a knob will spontaneously appear in most people’s minds, and friends will very likely make jokes about the next few Nobel Prizes. Before the second prize, however, the “just-one” quality would have been an unperceived fact.

This is closely related to a famous problem in cognitive science (the study of formal models of mental processes, especially computer models) called the “frame problem” [Dennett 81], which can be epitomized this way: How do I know, when telling you I’ll meet you at 7 at the train station, that it makes no sense to tack on the proviso, “as long as no volcano erupts along the way, burying me and my car on the way to the station,” but that it does make reasonable sense to tack on the proviso, “as long as no traffic jam holds me up”? And, of course, there are many intermediate cases between these two. The frame problem is about the question, “What variables (knobs) is it within the bounds of normalcy to perceive?” Clearly, no one can conceivably anticipate all the factors that might somehow be relevant to a given situation; one simply blindly hopes that the species’ evolution and the individual’s life experiences have added up to a suitably rich combination to make for satisfactory behavior most of the time. There are too many contingencies, however, to try to anticipate them all, even given the most powerful computer. One reason for the extreme difficulty in trying to make machines able to learn is that we find it very hard to articulate a set of rules defining when it makes sense and when it makes no sense to perceive a knob.

This brings us back to Shatter, seen as a variation on Helvetica. Obviously, once you’ve seen such a variation, you can add a knob (or a few) to your Metafont “Helvetica machine,” enabling Shatter to come out. (Indeed, you could add similar “Shatterizing” knobs to your “Baskerville machine,” for that matter!) But this would all be a posteriori: after the fact. The most telling proof of the artificiality of such a scheme is, of course, that no matter how many variations have been made on (say) Helvetica, people can still

come up with many new and unanticipated varieties, such as: Helvetica Rounded, Helvetica Rounded Deco, Helvetican Flair, and so on [Figure 8a, b, c] [Graphic 81].

Incidentally, it is important that I make it clear that although I find it easier to make my points with somewhat extreme or exotic versions of letters, these points hold just as strongly for more conservative letters. One simply has to look at a finer grain size, and all the same kinds of issues reappear.

No matter how many new knobs--or even new families of knobs--you add to your Helvetica machine, you will have left out some possibilities. People will forever be able to invent novel variations on Helvetica that haven't been foreseen by a finite parametrization, just as musicians will forever be able to devise novel ways of playing "Begin the Beguine" that the electronic-organ builders haven't yet built into their elaborate repertoire of canned rhythms, harmonies, and so forth. To be sure, the organ builders can always build in extra possibilities after they have been revealed, but by then a creative musician will have long since moved on to other styles. One can imagine Helvetica modified in many novel ways in-

Figure 8 (top to bottom). (a) Helvetica Rounded, (b) Helvetica Rounded Deco, and (c) Helvetican Flair [Graphic 81].

ABCDEFGHIJKLMNOPQRSTUVWXYZ&abcdefghijklmnopqrstuvwxyz\$¢1

ABCDEFGHIJKLMNOPQRSTUVWXYZ&abcdefghijklmnopqrstuvwxyz e

A A A ABCDEEF FG GH I J K L
LM M MN N OP P Q R R S T T U U V
W X Y Z & a a b c d e f f g g h h i j j k k l m



Figure 9 (top to bottom). (a) Sunrise, (b) Buster, (c) Stack, and (d) Double.

spired by various extant typefaces such as Sunrise, Buster, Stack, Double, and so on [Figure 9a, b, c, d]. I leave it to readers to try to imagine such variants.

A Total Unification of All Typefaces?

The worst is still yet to come! Knuth's throwaway sentence unspokenly implies that we should be able to interpolate any fraction of the way between any two arbitrary typefaces. For this to be possible, any pair of typefaces would have to share the exact same set of knobs (otherwise, how could you set each knob to an intermediate setting?). And since all *pairs* of typefaces have the same set of knobs, transitivity implies that *all* typefaces would have to share a single, grand, universal, all-inclusive, ultimate set of knobs. (The argument is parallel to the following one: If, given any two people, they have the same number of legs, then all people have the same number of legs.)

Thus we realize that Knuth's sentence casually implies the existence of a "universal 'A'-machine"--a single Metafont program with a finite set of parameters, such that any combination of settings of them will yield a valid "A", and conversely, such that any valid "A" will be yielded by some combination of settings of them. Now how can you possibly incorporate all of the previously shown typefaces into one universal schema?

Or look again at the 56 capital "A" 's of Figure 1. Can you find in them a set of specific, quantifiable features? (For a comparable collection for each letter of the alphabet, see [Kuwayama 73].) Imagine trying to pinpoint a few dozen discrete features of the Magnificat "A" (A7) and

simultaneously finding their “counterparts” in the Univers “A” (D3). Suppose you have found enough to characterize both completely. Now remember that every intermediate setting also must yield an “A”. This means we will have every shade of “cross” between the two typefaces.

This intuitive sense of a “cross” between two typefaces is common and natural, and occurs often to typeface lovers when they encounter an unfamiliar typeface. They may characterize the new face as a cross between two familiar typefaces (“Vivaldi is a cross between Magnificat and Palatino Italic Swash”) or they may see it as an exaggerated rendition of a familiar typeface (“Magnificat is Vivaldi squared”) [Figure 10]. What truth is there to such a statement? All one can really say is that each Magnificat letter looks “sort of like” its Vivaldi counterpart, only about “twice as fancy” or “twice as curly” or something vague along these lines. But how could a single “curliness” knob account for the mysteriously beautiful meanderings, organic and capricious, in each Magnificat letter?

Can you imagine twisting one knob and watching thin, slithery tentacles begin to grow out of the Palatino Italic “A”, snaking outwards eventually to form the Vivaldi “A”, then continuing to twist and undulate into ever more sinuous forms, yielding the Magnificat “A” in the end? And who says that that is the ultimate destination? If Magnificat is Vivaldi squared, then what is Magnificat squared?

Specialists in computer animation have had to deal with the problem of interpolation of different forms. For example, in a television series about evolution, there was a sequence showing the outline of one animal slowly transforming into another one. But one cannot simply tell the computer, “Interpolate between this shape and that one!” To each point in one there must be explicitly specified a corresponding point in the

Figure 10 (top to bottom). (a) Palatino Italic Swash caps, (b) Vivaldi caps, and (c) Magnificat caps.



other. Then one lets the computer draw some intermediate positions on one's screen, to see if the choice works. A lot of careful "tuning" of the correspondences between figures must be done before the interpolation looks good. There is no recipe that works in general.

The Essence of "A"-ness is Not Geometrical

Despite all the difficulties described above, some people, even after scrutinizing the wide diversity of realizations of the abstract "A"-concept, still maintain that they all do share a common geometric quality. They sometimes verbalize it by saying that all "A"'s have "the same shape" or are "produced from one template." Some mathematicians are inclined to search for a topological or group-theoretical invariant. A typical suggestion might be: "All instances of 'A' are open at the bottom and closed at the top." Well, in Figure 1, sample A8 (Stop) seems to violate both of those criteria. And many others of the sample letters violate at least one of them. In several examples, such concepts as "open" or "closed" or "top" or "bottom" apply only with difficulty. For instance, is G7 (Sinaloa) open at the bottom? Is F6 (Calypso) closed at the top? What about A4 (Astra)?

The problem with the Metafont "knobs" approach to the "A" category is that each knob stands for the presence or absence (or size or angle, etc.) of some specifically *geometric* feature of a letter: the width of its serifs, the height of its crossbar, the lowest point on its left arm, the highest point along some extravagant curlicue, the amount of broadening of a pen, the average slope of the ascenders, and so forth and so on. But in many "A" 's such notions are not even applicable. There may be no crossbar, or there may be two or three or more. There may be no curlicue, or there may be a few curlicues.

Chauvinism versus Open-Mindedness: Fixed Questionnaires versus Fluid Roles

A Metafont joint parametrization of two "A" 's presumes that they share the same features, or what might be called "loci of variability." It is a bold (and, I maintain, absurd) assumption that one could get any "A" by filling out an eternal and fixed questionnaire: "How wide is its crossbar? What angle do the two arms make with the vertical? How wide are its serifs?" (and so forth). There may be no identifiable part that plays the crossbar role, or the left-arm role; or some role may be split among two or more parts. You can easily find examples of these phenomena among the 56 "A" 's in Figure 1. Some other examples of what I call role splitting, role sharing, role transferral, role redundancy, and role elimination are shown in Figure 11. These terms describe the ways that conceptual roles




Figure 11.

are apportioned among various geometric entities, which are readily recognized by their connectedness and gentle curvatures.

When I was 12, my family was about to leave for Geneva, Switzerland, for a year, so I tried to anticipate what my school would be like. The furthest my imagination could stretch was to envision a school that looked exactly like my one-story Californian stucco junior high school, only with classes in French (twiddling the “language” knob), and with the schoolbus that would pick me up each morning perhaps pink instead of yellow (twiddling the “schoolbus color” knob). I was utterly incapable of anticipating the vast difference that there actually turned out to be between the Geneva school and my California school.

Likewise, there are many “exobiologists” who have tried to anticipate the features of extraterrestrial life, if it is ever detected. Many of them have made assumptions that to others appear strikingly naive. Such assumptions have been dubbed “chauvinisms” by Carl Sagan [Sagan 73]. There is, for instance, “liquid chauvinism,” which refers to the phase of the medium in which the chemistry of life is presumed to take place. There is “temperature chauvinism,” which assumes that life is restricted to a temperature range not too different from that here on the planet earth. In fact, there is planetary chauvinism—the idea that all life must

exist on the surface of a planet orbiting a certain type of star. There is carbon chauvinism, assuming that carbon must form the keystone of the chemistry of any sort of life. There is speed chauvinism, assuming that there is only one "reasonable" rate for life to proceed at. And so it goes.

If a Londoner arrived in New York, we might find it naive (or perhaps pathetic) if he or she asked "Where is your Big Ben? Where are your Houses of Parliament? Where does your Queen live? When is your teatime?" The idea that the biggest city in the land need not be the capital, need not have a famous bell tower in it, and so on, seem totally obvious after the fact, but to the naive tourist it can come as a surprise.

The point here is that when it comes to fluid semantic categories such as "A", it is equally naive to presume that it makes sense to refer to "the crossbar" or "the top" or to any constant feature. It is quite like expecting to find "the same spot" in any two pieces of music by the same composer. The problem, I have found, is that most people continue to insist that any two instances of "A" have "the same shape," even when confronted with such pictures as Figure 1.

The analogy between Britain and the United States is a useful one to continue for a moment. The role that London plays in England is certainly multifaceted, but two of its main roles are "chief commercial city" and "capital." These two roles are played by different cities in the U.S. On the other hand, the role that the American President plays in the U.S. is split into pieces in Britain, part being carried by the Queen (or King), and part by the Prime Minister. Then there is a subsidiary role played by the President's wife--the "First Lady." Her counterpart in Britain is also split, and moreover, these days "wife" has to be replaced by "husband," whether one is thinking that the "President of England" is the Queen or the Prime Minister. (See [Hofstadter 81] for an extended discussion of such analogy problems and their relation to machine intelligence.)

To think one can anticipate the complete structure of one country or language purely on the basis of being intimately familiar with another one is presumptuous and, in the end, preposterous. Even if you have seen dozens, you have not exhausted the potential richness and novelty in such domains. In fact, the more instances you have seen, the more circumspect you are about making unwarranted presumptions about unseen instances, although certainly your ability to anticipate the unanticipated (or unanticipable) improves! The same holds for instances of any letter of the alphabet or other semantic category.

The “A” Spirit

Clearly there is much more going on in typefaces than meets the eye--literally. The shape of a letterform is a surface manifestation of deep mental abstractions. It is determined by conceptual considerations and balances that no finite set of merely geometric knobs could capture. Underneath or behind each instance of “A” there lurks a concept, a Platonic entity, a *spirit*.

This Platonic entity is not an elegant shape such as the Univers “A”, not a template with a finite number of knobs, not a topological or group-theoretical invariant in some mathematical heaven, but a mental abstraction--a different sort of beast. Each instance of the “A” spirit reveals something new about the spirit without ever exhausting it. The mathematization of such a spirit would be a machine with a specific set of knobs on it, defining all its “loci of variability” for once and for all. I have tried to show that to expect this is simply not reasonable. In fact, I made the following claim, above: “No matter how many new knobs--or even new families of knobs--you add to your . . . machine, you will have left out some possibilities. People will forever be able to invent novel variations . . . that haven’t been foreseen by a finite parametrization. . . .”

Of what, then, is such an abstract “spirit” composed? Or is it simply a mystically elusive, noncapturable essence that defies the computational--indeed, the scientific--approach totally? Not at all, in my opinion. I simply think that a key idea is missing in what I have described so far. And what is this key idea? I shall first describe the key misconception. It is to try to capture the essence of each separate concept in a separate “knobbed machine”--that is, to isolate the various Platonic spirits. The key insight is that those spirits overlap and mingle in a subtle way.

Happy Roles, Unhappy Roles, and Quirk-Notes

The way I see it, the Platonic essence lurking behind any concrete letterform is composed of conceptual “roles” rather than geometric parts. (A related though not identical notion called “functional attributes” was discussed by Barry Blesser and co-workers nearly ten years ago in *Visible Language* [Blesser 73].) A role, in my sense of the term, does not have a fixed set of parameters defining the extent of its variability, but it has instead a set of tests or criteria to be applied to candidates that might be instances of it. For a candidate to be accepted as an instance of the role, not all the tests have to be passed; not all the criteria have to be present. Instead, the candidate receives a score computed from the tests and criteria, and there is a threshold point above which the role is “happy,” and below which it is “unhappy.” Then below that, there is a cut-off point below which the role is totally dissatisfied, and rejects the candidate outright.

An example of such a role is that of “crossbar.” Note that I am not

saying “crossbar in capital ‘A’, but merely “crossbar.” Roles are modular: they jump across letter boundaries. The same role can exist in many different letters. This is, of course, reminiscent of the fact that in Metafont a serif (or generally, any geometric feature shared by several letters) can be covered by a single set of parameters for all letters, so that all the letters of the typeface will alter consistently as a single knob is turned. The difference is that my notion of “role” doesn’t have the generative power that a set of specific knobs does. From the fact that a given role is “happy” with a specific geometric filler, one cannot deduce exactly how that filler looks. There is, of course, more to a role’s “feelings” about its filler than simply happiness or unhappiness; there are a number of expectations about how the role should be filled, and the fulfillment (or lack thereof) can be described in “quirk-notes.” Thus, quirk-notes can describe the unusual slant of a crossbar [Figure 1, E1 (Arnold Böcklin)], the fact that it is filled by two strokes rather than one [Figure 1, E3 (Aircraft)], or the fact that it fails to meet (or has an unusual way of meeting) its vertical mate [Figure 1, A2 (Eckmann Schrift), F5 (Le Golf), and many others].

These quirk-notes are characterizations of stylistic traits of a perceived letterform. They do not contain enough information, however, to allow a full reconstruction of that letterform, whereas a Metafont program does contain enough information for that. However, they do contain enough information to guide the creation of many specific letterforms that have the given stylistic traits. All of them would be, in some sense, “in the same style.”

Modularity of Roles

The important thing is that this modularity of roles allows them to be exported to other letters, so that a quirk-note attached to a particular role in “A” could have relevance to “E”, “L”, or “T”. Thus stylistic consistency among different letters is a by-product of the modularity of roles, just as the notion of letter-spanning parameters in Metafont gives rise to internal consistency of any typeface it might generate.

Furthermore, there are connections among roles so that, for instance, the way in which the “crossbar” role is filled in one letter could influence the way that the “post” or “bowl” or “tail” role is filled in other letters. This is to avoid the problem of overly simplistic mappings of one letter onto another, analogous to the Londoner asking an American where the American Houses of Parliament are. Just as one must interpret “Houses of Parliament” liberally rather than literally when “translating” from England to the U.S., so one may have to convert “crossbar” into some other role when looking for something analogous in the structure of another letter than “A”, such as “N”. In certain typefaces the diagonal stroke in “N” could well be the counterpart of the crossbar in “A”. But it is important to emphasize that no fixed (i.e., typeface-

independent) mapping of roles in “A” onto roles in “N” will work; only the specific letterforms themselves (via their quirk-notes) can determine what roles (if any) should be mapped onto each other. Such cross-letter mappings must be mediated by a considerable degree of understanding of what functions are fulfilled by all the roles in the two particular letters concerned. (This fluid mapping of roles is discussed in more detail in [Hofstadter 82b].)

Typographical Niches and Rival Categories

So far I have sketched very quickly a theory of “Platonic essences” or “letter spirits” involving modular roles—roles shared among several letters. This sharing of roles is one aspect of the overlapping and mingling that I spoke of above. There is a second aspect, which is suggested by the phrase “typographical niche.” The notion is analogous to that of “ecological niche.” When, in the course of perception of a letterform, a group of roles have been activated and have decided that they are present (whether happily or unhappily), their joint presence constitutes evidence that one of a set of possible letters is present. (Remember that since a role is not the property of any specific letter, its presence does not signal that any specific letter is in view.)

For instance, the presence of a “post” role and a “bowl” role in certain relative positions would suggest very strongly that there is a “b” present. Sometimes there may be evidence for more than one letter. The eye-mind combination is not happy with any such unstable state for long, and strains to make a decision. It is as if there is a very steep and slippery ridge between valleys, and a ball dropped from above is very unlikely to come to settle on top of the ridge. It will tumble to one side or the other. The valleys are the typographical niches.

Now the overlapping of letters comes about because each letter is aware of its typographical rivals, its next-door neighbors, just over the various ridges that surround its space. The letter “h”, for instance, is acutely sensitive to the fact that it has a close rival in “k”, and vice versa [Figure 12]. The letter “T” is very touchy about having its crossbar penetrated by the post below, since even the slightest penetration is enough to destroy its “T”-ness and to slip it over into “T”’s arch-rival niche, “t”. It’s a low ridge, and for that reason, “T” guards it extra-carefully.

The Intermingling of Platonic Essences

This image is, I hope, sufficiently strong to convey the second sense of overlapping and intermingling of Platonic essences. “No letter is an island,” one might say. There has to be much mutual knowledge spread about among all the letters. Letters mutually define each others’ essences, and this is why an isolated structure supposedly representing a single letter in all its glory is doomed to failure.

A letterform-designing computer program based on the above-sketched notions of typographical roles and niches would look very different from one that tried to be a full "mathematization of categories." It would involve an integration of perception with generation, and moreover an ability to generalize from a few letterforms (possibly as few as one) to an entire typeface in the style of the first few. It would not do so infallibly; but of course it is not reasonable to expect "infallible" performance, since stylistic consistency is not an objectively specifiable quality.

In other words, a computer program to design typefaces (or anything else with an esthetic or subjective dimension) is not an impossibility; but one should realize that, no less than a human, any such program will necessarily have a "personal" taste--and it will almost certainly not be the same as its designers' taste. In fact, to the contrary, the program's taste will quite likely be full of unanticipated surprises to its programmers (as well as to everyone else), since that taste will emerge as an implicit and remote consequence of the interaction of a myriad features and factors in the architecture of the program. Taste itself is not directly programmable. Thus, although any esthetically programmed computer will be "merely doing what it was programmed to do," its behavior will nonetheless often appear idiosyncratic and even inscrutable to its programmers, reflecting the fact--well known to programmers--that often one has no clear idea (and sometimes no idea at all) just what it is that one has programmed the machine to do!

Figure 12. Versions of "h" and "k" as rivals for the same typographical niche.



a	b	c	d	e	f...
a	b	c	d	e	f...
a	b	c	d	e	f...
a	b	c	d	e	f...
a	b	c	d	e	f...
a	b	c	d	e	f...
:	:	:	:	:	:

Figure 13. The vertical and horizontal problems.

The “Vertical” and “Horizontal” Problems: Two Equally Important Facets of One Problem

I have made a broad kind of claim: that true understanding of letterforms depends on more than understanding something about each Platonic letter in isolation; it depends equally much on taking into account the ways that letters and their pieces are interrelated, on the ways that letters depend on each other to define a total style. In other words, any approach to the impossible dream of the “secret recipe” for “A-ness” requires a simultaneous solution to two problems, which I call the “vertical” and the “horizontal” problems [Figure 13]. The former is the question, “What do all the items in any column have in common?” The latter is the question, “What do all the items in any row have in common?”

Actually, there is no reason to stop with two dimensions; the problem seems to exist at higher degrees of abstraction. We could lay out our table of comparative typefaces more carefully; in particular, we could make it consist of many layers stacked on top of each other, as in a cake. On each layer would be aligned many typefaces made by a single designer. This idea is illustrated in Figure 14, showing a few faces designed by Hermann Zapf (Optima, Palatino, Melior, Zapf Book, Zapf International, Zapf Chancery) (see [Zapf 60]). Along with the Zapf layer, one can imagine a Frutiger layer, a Lubalin layer, a Goudy layer, and so on. One could try to arrange the typefaces in such a way that “corresponding” typefaces by various designers are aligned.

Now in this three-dimensional cake, the two earlier one-dimensional questions still apply, but there is also a new two-dimensional question: “What do all the items in a given layer have in common?” The third dimension can be explored as one moves from one layer to another, asking what all the typefaces in a given “shaft” have in common.




Figure 14. A few faces designed by Hermann Zapf (top to bottom). (a) Optima, (b) Palatino, (c) Melior, (d) Zapf Book, (e) Zapf International, and (f) Zapf Chancery.

Moreover, a fourth dimension can be added if you imagine many such “layercakes,” one for each distinguishable period of typographical design. Thus our fourth dimension, like Einstein’s, corresponds to time. Now one can ask about each layercake, “What do all the items herein have in common?” This is a three-dimensional question. Presumably, one could carry this exercise even further.

If we go back to the “simplest” of these questions, the original “vertical” question of Figure 13, a naive answer to it could be stated in one word: “Letter.” And likewise, a naive answer to the “horizontal” question of Figure 13 is also statable in one word: “Spirit.” In fact, the word “spirit” is applicable, in various senses of the term, to all the higher-dimensional questions, such as “What do all the typefaces produced in the Art Deco era have in common?” There is such a thing, ephemeral though it may be, as “Art Deco spirit,” just as there is undeniably such a thing as “French spirit” in music or “impressionistic spirit” in art (see [Loeb 75]).

Stylistic moods permeate whole periods and cultures and indirectly determine the kinds of creations that people in them come up with. They exert gentle but definite “downward” pressures. As a consequence, not only are the alphabets of a given period and area distinctive, but one can even recog-

nize “the same spirit” in such things as teapots, coffee cups, furniture, automobiles, architecture, and so on [Bush 75]. And one can also be inspired by a given typeface to carry its ephemeral spirit over into another alphabet, such as Greek, Hebrew, Cyrillic, or Japanese. In fact, this has been done in many instances [Figure 15]. The problem I am most concerned with in my research is whether (or rather, how) susceptibility to such a “spirit” can be implanted in a computer program.

Letter and Spirit

These words “letter” and “spirit,” of course, recall the contrast between the “letter of the law” and the “spirit of the law,” and the way in which our legal system is constructed so that judges and juries will base their decisions on precedents. This means that any case must be “mapped,” in a remarkably fluid way, by members of a jury, onto previous cases. It is up to the opposing lawyers, then, to be advocates of particular mappings; to try to channel the jury members’ perceptions so that one mapping dominates over another. It is quite interesting that jury decisions are supposed to be unanimous, so that in a metaphorical sense, a “phase transition” or “crystallization” of opinion must take place. The decision must be solidly locked in, so that it reflects not simply a majority or even a consensus, but a totality, a unanimity (which, etymologically, means “one-souledness”). (For discussions of such “phase transitions,” see [Hofstadter 82c] and [Hofstadter 83], and for descriptions of computer models of perception in which a form of collective decision making is carried out, see [Reddy 76] and [Winston 75].)

In law, extant rules, statutes, and so on are never enough to cover all possible cases (reminding us once again of the fact that no fixed and rigid set of “A”-defining rules can anticipate all “A” ’s). The legal system depends on the notion that people, whose experience covers much more than the specific case and rules at hand, will bring to bear their full range of experience not only with many categories but also with the whole process of categorization and mapping. This allows them to transcend the specific, rigid, limited *rules*, and to operate according to more fluid, imprecise, yet more powerful *principles*. Or, to revert to the other vocabulary, this ability is what allows people to transcend the letter of the law and to apply its spirit. It is this tension between rules and principles, between the letter and the spirit, that is so admirably epitomized for us by the work of Knuth and others exploring the relationship between artistic design and mechanizability. We are entering a very exciting and important phase of our attempts to realize the full potential of computers, and Knuth’s article points to many of the significant issues that must be thought through very carefully.

Figure 15. The “spirit” of some Roman typefaces carried over into Cyrillic and Greek typefaces [Compugraphic 82] and into Hebrew and Japanese typefaces [Biggs 77]. The related Kana and Latin letters were designed by Yasaburo Kuwayama for the Nissan Company.

TASTE IN PRINTING DETERMINES THE FORM **English Times**
 Typography is to take. The selection of a congruous typeface **Bold**
 the quality and the suitability for its purpose of the paper

ШРИФТОТЕКА КОМПЬЮГРАФИК СОДЕРЖИТ **Cyrillic**
 дысяза гарнитуров шрифта включающихкак традици **Times Bold**
 ио нные так современные рисунок шрифта, кото рые

TASTE IN PRINTING DETERMINES THE FORM **Oracle Italic**
Typography is to take. The selection of a congruous typeface, the
quality and suitability for its purpose of the paper being us

Ἡ καλαισθησιὰ καὶ ἡ ἀπόδοσις τῇ ἐκτύπωσιν προσ **Greek**
 διο ρίξει μορφὴν πού θάρει τό τυπωμένο κείμενο τῇ
 ἐπιλογῇ τοῦ ἀνάλογου ὀρθογράφου, τῇ ποιότητι καὶ

אות עברית זאת, הראשונה שעוצבה **This type face, the first**
בארבעה משקלות, הותאמה במיוחד **in Hebrew to be available**

לשימוש עם סדרת יוניברס הלטינית. **in four weights, is also**

זאת הפעם הראשונה שאות עברית **the first designed specially**

הותאמה לאותיות הקטנות באלף בית לטיני **to align with the lower case**

כך שיתאימו במיוחד, זו בצד זו, לטקסטים **of a Latin type face, for use**

ידו לשוניים ארוכים, בהם השימוש באותיות **together in bilingual printing**

הגדולות (caps) בלבד יכגע מאוד בקריאות. **of extended texts.**

アイウエオカキク	ABCDEFGHI
ケコサシスセソタ	JKLMNOPQR
チツテトナニヌネ	STUVWXYZ&
ノハヒフヘホマミ	1234567890: ,
ムメモヤユヨラリ	abcdefghijkl -
ルルロロウウガグ	mnopqrstuvw
ゲザジズダヅバビ	xyz
パピプペーアイウエ	

To conclude, then, I wish to state that the mathematization of categories is an elegant goal, a wonderful beckoning mirage before us, and the computer is the obvious medium to exploit to try to realize this goal. Donald Knuth, whether he has been pulled by a distant mirage or by an attainable middle-range goal, has contributed immensely, in his work on Metafont, to our ability to deal with letterforms flexibly, and has cast the whole problem of letters and fonts in a much clearer perspective than ever before. Readers, however, should not pull a false message out of his article: they should not confuse the chimera of the mathematization of categories with the quest after a more modest but still fascinating goal. In my opinion, one of the best things Metafont could do is to inspire readers to chase after what Knuth has rightly termed “the intelligence” of a letter, making use of the explicit medium of the computer to yield new insights into the elusive “spirits” that flit about so tantalizingly, hidden just behind those lovely shapes we call “letters.”

References

- [Biggs 77] Biggs, John R. *Letterforms and Lettering*, Blandford Press, 1977.
- [Blessner 73] Blessner, B., R. Shillman, C. Cox, T. Kuklinski, J. Ventura, and M. Eden. “Character Recognition Based on Phenomenological Attributes,” *Visible Language*, Vol. VII, Summer 1973.
- [Bush 75] Bush, Donald J. *The Streamlined Decade*, George Braziller, 1975.
- [Chaitin 75] Chaitin, Gregory. “Randomness and Mathematical Proof,” *Scientific American*, Vol. 232, May 1975, pp. 47-52.
- [Compugraphic 82] *A Portfolio of Text and Display Type*. Compugraphic Corporation, 1982.
- [DeLong 70] DeLong, Howard. *A Profile of Mathematical Logic*. Addison-Wesley, 1970.
- [Dennett 81] Dennett, Daniel. *Cognitive Wheels: A Philosophers’ Introduction to the Frame Problem of Artificial Intelligence* (forthcoming).
- [Gödel 31] Gödel, Kurt. *On Formally Undecidable Propositions of Principia Mathematica and Related Systems*, translated by B. Meltzer. Basic Books, 1962.
- [Graphic 81] *Format Catalogue No. 7*. Graphic Products Corporation, 1981.
- [Haab 54] Haab, A., A. Stocker, and W. Haettenschweiler. *Lettera 1* and *Lettera 2*. Arthur Niggli, Publishers, 1954 and 1961.
- [Hofstadter 79] Hofstadter, Douglas R. *Gödel, Escher, Bach: an Eternal Golden Braid*. Basic Books, 1979.
- [Hofstadter 81] Hofstadter, Douglas R. “Metamagical Themas: Roles and Analogies in Human and Machine Thought,” *Scientific American*, Vol. 245, September 1981, pp. 18-30.
- [Hofstadter 82a] Hofstadter, Douglas R. “Metamagical Themas: On Spheexishness, ‘Jootsing,’ and Creativity,” *Scientific American*, Vol. 247, September 1982, pp. 18-34.

- [Hofstadter 82b] Hofstadter, Douglas R. "Metamagical Themas: On Variations on a Theme, Slippability, and Creativity," *Scientific American*, Vol. 247, October 1982, pp. 20-29.
- [Hofstadter 82c] Hofstadter, Douglas R. "Who Shoves Whom Around Inside the Careenium?" Indiana University Computer Science Department Technical Report No. 130, July 1982. (Also to appear in *Synthese*.)
- [Hofstadter 83] Hofstadter, Douglas R. "Artificial Intelligence: Subcognition as Computation," in F. Machlup and Una Mansfield (eds.), *The Study of Information*. To be published in 1983 by John Wiley and Sons.
- [Kim 81] Kim, Scott E. *Inversions: A Catalogue of Calligraphic Cartwheels*. Byte Books, 1981.
- [Knuth 82] Knuth, Donald E. "The Concept of a Meta-Font," *Visible Language*, Vol. XVI, Winter 1982, pp. 3-27.
- [Kuwayama 73] Kuwayama, Yasaburo. *Trademarks and Symbols, Volume I: Alphabetical Designs*. Van Nostrand Reinhold, 1973.
- [Larcher 76] Larcher, Jean. *Fantastic Alphabets*. Dover Press, 1976.
- [Letraset 81] *Graphic Art Materials Reference Manual*. Letraset, 1981.
- [Loeb 75] Loeb, Marcia. *New Art Deco Alphabets*. Dover Press, 1975.
- [Myhill 52] Myhill, John. "Some Philosophical Implications of Mathematical Logic: Three Classes of Ideas," *Review of Metaphysics*, Vol. VI, December 1952, pp. 165-198.
- [Nagel 58] Nagel, Ernest, and James R. Newman. *Gödel's Proof*. New York University Press, 1958.
- [Post 44] Post, Emil. "Recursively Enumerable Sets of Positive Integers and Their Decision Problems," *Bulletin of the American Mathematical Society*, Vol. 50, 1944, pp. 284-316; reprinted in Martin Davis (ed.), *The Undecidable*. Raven Press, 1965.
- [Reddy 76] Reddy, D. Raj, et al. "Working Papers in Speech Recognition, IV: The Hearsay II System," Carnegie-Mellon University Computer Science Department Technical Report, 1976.
- [Rogers 67] Rogers, Hartley. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, 1967.
- [Rucker 82] Rucker, Rudy. *Infinity and the Mind*. Birkhäuser, 1982.
- [Sagan 73] Sagan, Carl (ed.), *Communication with Extraterrestrial Intelligence*. MIT Press, 1973.
- [Smullyan 61] Smullyan, Raymond. *Theory of Formal Systems*. Princeton University Press, 1961.
- [Smullyan 78] Smullyan, Raymond. *What Is the Name of This Book?* Prentice-Hall, 1978.
- [Strich 81] Strich, Christian. *Fellini's Faces: 418 Photographs from the Archives of Federico Fellini*. Holt, Rinehart, and Winston, 1981.
- [Winston 75] Winston, Patrick H. (ed.) *The Psychology of Computer Vision*. McGraw-Hill, 1975.
- [Zapf 60] Zapf, Hermann. *About Alphabets*. MIT Press, 1960.