

Toward a typographical linguistics

*Keynote address for the ATypI 1997 Conference,
University of Reading, 13 September 1997*

David Crystal

It must be twenty years since I was a member of a rewarding series of interdisciplinary seminars held under the auspices of Michael Twyman and the Typography Department in Earley Gate at the University of Reading, in which I found myself for the first time discussing problems to do with graphic language at the shadowy interface between my subject and yours. Much of the focus of those seminars was on the criteria for evaluating reading materials for children, and I very clearly recall the lack of a common vocabulary for describing the graphic features of those materials, and much confusion over what was most salient – or even how the notion of graphic salience itself could be best characterized. It seemed to me then that there was an interdisciplinary focus awaiting invention – a “typographical linguistics,” as I have since occasionally called it – but given the remarkable paucity of typographers who knew anything about linguistics, or the even greater paucity of linguists who knew anything about typography, there seemed little chance of such a focus emerging – and so it has since proved. I have done my best to incorporate a typographic dimension into my own general writing on linguistics – and there are some here today who have helped me do so. But linguistics texts on the whole ignore typography. Nor, I understand, is it routine for linguistics terms to figure prominently in the indexes of typographic works. This is a shame, because it seems to me that the explication of printed language needs the expertise of both typographers and linguists, in order to provide a complete description of its forms and structures and a satisfactory explanation of its functions and effects. I therefore very much welcome opportunities to engineer a close encounter between the two fields, as with the present conference – though (in my case, at least) it has to be a limited encounter, only of the first kind, doing little more than

identifying some of the data which any theory of typographical linguistics would have to explain.

In investigating linguistic data, linguists very much rely on the notion of linguistic “levels” – a notion which pervades the domain of linguistic inquiry, though it seems not to have much resonance within typography. The idea has a very simple origin. It arises out of the recognition that language is the most complex domain of human behavior and that it is impossible to capture everything to do with its structure within a single type of generalization. I can illustrate this complexity by making some crude quantitative estimates: while I am talking to you now, in English, I am using some 44 contrasting sounds (the vowels and consonants), combining them in some 300 possible ways to form syllables, adding over 100 patterns of melody and rhythm, putting my syllables to work in sequences to make words (over 50,000 available in common usage), and using these words to build up phrases, clauses, sentences, and stretches of discourse (over 3,000 distinguishable grammatical properties). Now, if you tried to describe what was going on, you would find that the kinds of statement you could make would vary greatly as you focused first on one aspect of this complexity, then another. For example, if you chose to study the way the sounds worked, you would find yourselves talking (at least in part) about the physical mechanisms of articulation and hearing, and using terminology which derived from the study of human anatomy and physiology. By contrast, if you chose to study the way the words worked, you would find yourselves talking (at least in part) about the nature of meaning, and about how words group together to express the structure of the objects, of phonetics – seems totally different from that of the second – semantics; and yet both comfortably coexist within the stream of spoken or written language. How many such “worlds” are there, comfortably coexisting in this way? That is what the notion of “level” was developed to elucidate.

It turns out that there is no indisputable answer to the number and nature of the linguistic levels it proves illuminating to recognize when studying language. When you are trying to model something as complex as language, there is no single way of viewing the

behavior: some models make distinctions which other models ignore. Most linguists find five levels useful. When studying spoken language, we can identify first the domain of pronunciation, which is conventionally analyzed into two levels: (1) *phonetics*, which I mentioned above, and (2) *phonology*, which I shall discuss later. When studying the domain of written language (in the broadest sense, to include handwriting, type, print, electronic, and any other visually mediated expression of language), some (by no means all) linguists draw an analogous distinction of levels between (1) *graphetics* and (2) *graphology*, which I shall also discuss below. Next, there is the domain of grammar, applicable to both spoken and written language, and often handled as just a single level, though for present purposes divided into two levels: (3) sentence-grammar (*syntax*) and (4) word-grammar (*morphology*). Finally, there is the domain of meaning – what the sounds and grammar are there to express – also applicable to both spoken and written language, and usually identified as a single level, (5) semantics. Other possibilities have been suggested, but for present purposes I shall restrict myself to these five.

Ultimately, all levels play their part in the structuring and expression of meaning. Sounds and letters are obviously there to differentiate words (phonology/graphology): *pin* is different from *bin*, and *bin* from *bun*, and *bun* from *bus*. In many instances, it only needs the replacement of one sound to change a meaning. Changes in grammar also facilitate contrasts of meaning – singular vs. plural, past tense vs. present tense, alterations of word order (*Man chases dog, Dog chases man*). And the vocabulary of the language, which is at the heart of semantics, is obviously all to do with meaning: we learn “the meaning of a new word” and are confused if we do not know what “that word means.” Explaining how we communicate meanings to each other, using the spoken or written medium, is what linguistics is all about. So in relation to the present conference, the interesting question, for me as a linguist, is: How do the various features of typography relate to the need to communicate meaning? More specifically: To what extent do the various features of typography convey linguistic meaning? To what extent do they impede the communication of that meaning?

To what extent are they ambiguous in the communication of that meaning? and so on.

For typography to convey linguistic meaning, we would need to be able to identify those typographic features which are the source of the way a particular word, phrase, sentence, or text is to be interpreted. In exactly the same way that the substitution of a particular sound changes the meaning of a word, and a different word changes the meaning of a phrase, and a different phrase changes the meaning of a sentence, and a different sentence changes the meaning of a text, I ask: Could the substitution of one typographic feature within a word cause that word to change its meaning? Could changing the typography of a word within a phrase cause that phrase to change its meaning? Could changing the typography of a phrase within a sentence cause that sentence to change its meaning? Could changing the typography of a sentence within a text cause that text to change its meaning?

My impression is that these questions are not often asked, because it is only after this point that typographic considerations normally apply. If I have understood your field correctly, you are chiefly concerned with the typographic properties of texts as wholes – whether the text is as small as a logo or a bus ticket or as large as a book – as suggested by such characterizations of the subject matter of typography as “the determination of the appearance of the printed page” (see, e.g., Warren E. Preece, in the *Encyclopedia Britannica* article on typography). Linguists too are interested in texts as wholes – indeed, there are several domains within my subject which are devoted to finding out what are the linguistic properties of whole texts, such as stylistics and the aptly named textlinguistics. And it is here, as a matter of fact, that our two subjects have hitherto come closest together. Stylistic studies of written texts quite routinely try to describe the graphic features which make them distinctive, and borrow terminology from typography in order to do so. And when linguists talk about the way in which texts, whether spoken or written, make use of particular linguistic features – such as the use of vocative O in a religious text, or the use of an impersonal style in a scientific text, or the use of the present tense in sports commentary – they do so using such notions as

appropriateness, predictability, distinctiveness, and aesthetics – notions which you make use of too. Your concern with “how a text looks” is exactly parallel to ours: we too are concerned with “how a text (including both the handwritten and the typographic) looks,” as well as with the associated notions, “how a text sounds” and “how a text is structured.” But because this overlap between our fields is well recognized, I shall not dwell on it here. I want to focus, rather, on the extent to which, *within* a text – within the printed page, if you will – it is possible for us to use typographic features to change the linguistic meaning of what we write.

Before illustrating this point, I must refer to a second reason why this topic is not commonly addressed, and this arises out of the familiar observation (expressed, for example, by Beatrice Warde) that typography ought to be an invisible carrier of content. Now I understand this view, for it is found within linguistics also, in relation to the notion that there needs to be a level of physical substance in order for linguistic communication to take place at all – an actual medium of transmission. This is typically sound waves in the air in the case of spoken language and visual marks on a surface in the case of written language. And at this level, we are not talking about meaning – rather, we are talking about the way in which phonic substance can “carry” a stream of meaningful speech and graphic substance can “carry” a stream of meaningful writing (or print). In both cases, the carrier should not be noticed. It is invisible or, better, transparent, allowing us to get at the message it carries.

However, the properties of phonic and graphic substance are not without interest. People make idiosyncratic use of this substance, giving it a personal shaping which is present in everything they say or write. In the case of speech, these permanently present features of speech act as a background identity feature, sometimes referred to as “voice quality” or “articulatory setting.” Normally, while a person is speaking, we are paying attention to what is being said, and voice quality is auditorily invisible. But we can have our attention drawn to it, as when someone speaks with an abnormal voice (perhaps because of a sore throat), or when we need to recognize someone by voice alone, or when someone needs to perform

a certain social function and adopts a characteristic form of vocal expression in order to do so (as in the harsh tones of the drill sergeant). Graphic substance, likewise, has a person- or group-identifying capability. Thus, as we all know, we can recognize someone by their handwriting and many groups or products by their characteristic typography (e.g., newspapers). It is not so much what they are writing about, or even the way they are writing about it, but the way that their written language looks, while writing about it, which is the thing.

Phoneticians are the professionals who study the way human beings employ phonic substance. They are a long-established group within the linguistic sciences. There is no analogous group of professional grapheticians. For anything comparable, we have to look outside my subject. The nearest thing we get to grapheticians are (for handwriting) the “graphologists” who specialize in the interpretation of personal identity, whether forensically or in tents at Blackpool – or, of course, for printed language, typographers. And this is where an interesting contrast emerges between our fields. Phoneticians, on the whole, have spent most of their time analyzing just those properties of phonic substance which are relevant for the communication of linguistic meaning, in the form of vowels, consonants, and such vocal effects as stress and intonation. They have spent relatively little time researching the background features of speech – voice quality and so on – precisely because these features are not involved in the communication of meaning. They are not the focus of interest. Typographers, by contrast, seem to spend most of their time on the permanently present background features of (printed) graphic expression and very little on the way typographic features can directly influence the linguistic meaning.

The result of this difference of focus is that there is a very real gap between our two subjects. Linguists are traditionally not interested in studying formal features operating at the level of phonic or graphic substance unless there is an immediate pay-off in terms of their contribution to linguistic meaning. They have traditionally been skeptical of accepting a vaguer pay-off in pragmatic or stylistic terms, as when one claims that a feature has a particular

“impact” on its listener or reader. It is not that linguists are uninterested in such things: on the contrary, a whole field of study has developed within pragmatics to analyze them. We are ready, along with the speech-act philosophers, to talk about the intended effect of an utterance and the effect it actually has on its audience, and we have developed terminology in order to do so (*illocutionary force*, for the former; *perlocutionary effect*, for the latter). We encounter stylisticians, sociolinguists, and others who are prepared to talk about these effects in terms of elegance, balance, vividness, contrast, and a host of other aesthetic and functional notions. But linguists recognize that this is a vaguer and less readily formalized world – that it is much more difficult to phrase a precise answer to the question “What is the *effect* of a particular choice of language?” than to answer “What is the *meaning* of a particular choice of language?” That is why we draw as clear a line as we can between pragmatics, which studies the effects of linguistic choices, and semantics, which studies their meanings.

If our two subjects are to come closer together, then, there seem to be only two ways of doing so. One is for linguists to become more interested in the properties of graphic substance, either in the form of individual typefaces or in the way the cumulative effects of using a particular typeface result in judgments about its effect. The other is for typographers to become more interested in the linguistic properties of printed language. I hope typographers will have occasion to speak to gatherings of linguists in relation to the former point; today, I have the chance to illustrate some of the features of graphic language which interest me. And they are all to do, inevitably, with meaning. Which typographic features cause a change in linguistic meaning? I am not here talking about spelling – the selection of particular sequences of letters to represent the sounds of words. Rather, given a word (or phrase or sentence) with a particular spelling, the question is: Are there variations in the way it can be typographically presented which will cause its meaning to change?

Obviously, switching from serif to sans-serif will not cause a semantic change. Equally obviously, a change from roman to italics or boldface can do so. In the first example below, the implications

of the sentence are altered; in the second, the meaning of a word is different.

I've lost my red slippers.
(I've lost a pair of slippers, which happen to be red)

I've lost my *red* slippers.
(i.e., not my blue ones)

I have been reading about America in the paper.
(i.e., the country)

I have been reading about *America* in the paper.
(i.e., the book by Alastair Cook)

However, the range of effects conveyed by switching between roman and bold and italic are quite well understood, so I will not spend time on them here. Less obvious are the effects conveyed by lower-level typographical contrasts, such as case.

Not all uses of case contrast in English are semantically relevant. When we use an uppercase I for the first-person pronoun, there is no semantic contrast, because there is no lowercase i used as an isolated word which means something different. To use an *i* would just look odd. Similarly, a sentence beginning with a lowercase letter would still be a sentence; the presence of the preceding punctuation mark would simply indicate that the person had forgotten to capitalize. Initial sentence capitalization is not contrastive: there are several languages which make no such distinction in their writing (e.g., Hebrew). Most proper names, likewise, convey no semantic contrast if their capitalization is altered: *london* is still the same *London*, *john smith* is still the same *John Smith*.

But with proper names we do find some possible contrasts – where a meaning difference is signaled by the choice of one case rather than the other. To travel in *western Australia*, with a small *w* is to visit different parts of the country than to travel in *Western Australia*, with a large *W*. There are many other such cases:

Moon/moon	Bible/bible	God/god
Sun/sun	Liberal/liberal	Depression/depression
the Bar/the bar	the Church/the church	the Stones/the stones

In speech, we sometimes grope for a way of expressing this kind of contrast when we say “Bible with a capital B,” or “the church with a small c,” and the like. It is unclear just how much of this kind of thing goes on, but the existence of word lists of capitalization conventions compiled by copy editors in manuals of house style suggests that it is not insignificant.

Another type of example is when we write words with initial capitals in order to give them extra significance, as in the *Good Old Days*. Sometimes there is a clear suggestion of a different tone of voice. John Mortimer introduced us to *Mrs Hilda Rumpole (known to me only as She Who Must Be Obeyed)*. There is a storable difference of meaning between saying that someone has made a very important point and that that person has made a *Very Important Point*. A.A. Milne uses this device a great deal in *Winnie-the-Pooh*:

“Would you mind coming with me, Piglet, in case they turn out to be Hostile Animals?”

“If I know anything about anything, that hole means Rabbit,” he said, “and Rabbit means Company,” he said, “and Company means Food and Listening-to-Me-Humming and such like.”

Eeyore doesn't normally talk in initial capitals, except when he is making a *Very Important Point*:

“That Accounts for a Good Deal,” said Eeyore gloomily. “It Explains Everything. No Wonder.”

“Thank you, Pooh,” answered Eeyore. “You're a real friend,” said he. “Not Like Some,” he said.

Are there effects which can be achieved regardless of case? Here we are looking for examples where it is possible to change the typographic form of an individual letter or letters, within a word, in order to reinforce, contradict, or extend the word's meaning in some way. A familiar example is the way illuminated biblical man-

uscripts would often highlight the initial letter of *Jesus*, or any saint's name, to show special status. This kind of thing does not seem to happen much today – though one does occasionally encounter characters in stories who can be identified by the typography in which their words are set (such as the speech balloons in the *Pogo* comic strips of a few years ago in the United States). Its value is that, once a typographic convention is established, contrasts could be made. For example, it would be possible to use some kind of black-letter font as a symbol of religious status and then print

James spoke to James

which would mean that St. James spoke to non-St. James. We get close to this kind of thing in the world of modern graphic design, where an individual letter can be visually transformed to make some extra point – such as an *i* becoming a candle, or two *o*'s a pair of eyes. But this is not changing the meaning of the word: an ad for some firm in which, for example, the slogan *We illuminate!* has both of the *i*'s drawn as if they were candles does not change the meaning of the word *illuminate*. It simply adds a symbol which relates to what is being advertised. This is using visual communication to reinforce linguistic communication, just as one might add other semiotic modalities – smell, taste, music, animation – to provide extra impact. It is not a change in the nature of the linguistic communication as such.

The possibilities suggested by the “James spoke to James” example are more often found applied to whole words and phrases. For example, the archaic connotations surrounding various kinds of Old English typefaces are sufficiently established to allow a contrast between *Tea Shop* and **Tea Shop** – and the latter of course is often further reinforced by adopting an archaic spelling, **Tea Shoppe**. Similarly a *Christmas Fair* is not the same as a **Christmas Fair**, and would be even less like a **Christmas Fayre**. Correspondingly, we would not expect to find a firm called **MODERN ALARMS**. I believe there has been a certain amount done in your field on this notion of typographic appropriateness; it is certainly one of the topics which comes closest to the interests of linguists.

Linguists are also concerned to move away from issues of performance toward those of competence. By this I mean that linguistics is always trying to explore what it is we know about language – our awareness of the rules which govern what we say or write, both the acceptable utterances and the unacceptable ones. I know that *The cat sat on the mat* is an acceptable English sentence, whereas *Mat cat sat on the the* is not. I also know that, if I were to encounter someone saying *I spoke to Mary three whiskies ago*, I would be able to interpret that sentence even though I may never have heard it before in my life. The same principle applies to our investigation of typographic contrastivity. If we begin to think generatively – about what is possible and what is not, and about how we use our linguistic knowledge to extract meaning from new linguistic experiences – then we must go beyond the *Tea Shop* kind of example, and ask what it is theoretically possible to do with typography, in order to convey linguistic meaning. I have only time for one example, and I choose it from the domain of poetry.

One of the noticeable features of a great deal of poetry is the way patterns of sound (or the associated letters) are built up throughout a poem. Poets like to bring together words which have similarities of sound. Sometimes these similarities occur at the ends of words at the ends of lines – in which case we call them rhyme. Sometimes they occur at the beginnings of words – in which case we talk of alliteration. And there are several other types of pattern. But in all cases, these similarities perform two functions. There is a phonetic function, usually an aesthetic one: the similarities “sound nice,” in some way. More important, there is a phonological function: to paraphrase poet and critic William Empson, the similarities of sound make us think of similarities of sense. Words which are linked by sound make us look for semantic links too. And the connotations from one part of a rhyming pair can carry through and affect the other. Here is an example from “Mr. Eliot’s Sunday Morning Service”:

The sable presbyters approach
The avenue of penitence;
The young are red and pustular
Clutching piaculative pence.

The second half of the poem is an ironic comment on the first. The alliterative effect of presbyter and penitence is reinforcing: the two meanings sit happily together. But then this nice semantic association is blown sky-high: the unpleasant pustular takes us back to presbyter – there are several phonological similarities between the words – and the materialistic pence appears inside the other-worldly penitence, thereby putting it in its place. These are real semantic effects, in which connotations from one word unexpectedly affect another.

Can typography be put to work to do this kind of thing? Would it be possible to write a poem in which, through typographic contrastivity alone, the connotations from one word (or phrase) would be transferred to another – or even the denotation of one word altered by another? I looked for examples in my collection of poetry and could not find any – so am reduced to writing some short poems myself, to illustrate what I am getting at. I make no literary claims for these occasional pieces, which are here only to make a point. Whether there is already a genre out there of what I will call “graphic poetry” I do not know. I welcome your contributions.

If the effects truly reflect our typographical competence, such poems need no explanation. Because people already know the rules governing the way typographic features work, they should be able to deduce its meaning. Is this so?

Hagidiscography

Peter and **P**aul.

Simon and Garfunkel.

Matthew and **M**ark.

Peters and Lee.

Luke and **J**ohn.

Captain and Tenille.

Cosmas and **D**amian.

George and Ringo.

Cyril and **M**ethodius.

Nina and Frederick.

Elvis?

Those to whom I have shown this poem seem to get the point straight away. I suppose what it is illustrating is a kind of typographic line-rhyme. Note that the meaning is different if one prints the poem in the same typeface throughout:

Peter and Paul.

Simon and Garfunkel.

Matthew and Mark.

Peters and Lee.

Luke and John.

Captain and Tenille.

Cosmas and Damian.

George and Ringo.

Cyril and Methodius.

Nina and Frederick.

Elvis?

Two things happen here. The structure of the poem becomes opaque; and the conclusion of the poem reduces in force, because it leaves it completely open which group Elvis is to be assigned to. In the first setting, there is no question: Elvis’s sainthood is being asserted, then queried.

Could one reduce the linguistic unit involved in this exercise so that the process became more like the alliteration and syllable rhyme of the Eliot example? There are several possibilities. Here is what this poem would look like if it were typographically rhyming on the last stressed syllable in the line.

Peter and **P**aul.

Simon and Gar**f**unkel.

Matthew and **M**ark.

Peters and *Lee*.

Luke and **J**ohn.

Captain and *Tenille*.

Cosmas and **D**amian.

George and *Ringo*.

Cyril and **M**ethodius.

Nina and *Frederick*.

Elvis?

And here is what it would look like if it were being alliterative:

Peter and **P**aul.

Simon and **G**arfunkel.

Matthew and **M**ark.

Peters and **L**ee.

Luke and **J**ohn.

Captain and **T**enille.

Cosmas and **D**amian.

George and **R**ingo.

Cyril and **M**ethodius.

Nina and **F**rederick.

Elvis?

Is the effect there? Yes, but the reader has to work harder. Certainly, neither, to my mind, is as successful as the first. And it prompts the thought that maybe typographic effects are most efficiently linguistically contrastive only when they are used with linguistic units which are already meaningful – such as a word, a phrase, or a sentence.

Or a part of a word? A prefix or a suffix? This next poem suggests that it is possible for typographic contrastivity to work at the morphological level. I hasten to add that I have no evidence whatsoever of nefarious goings-on in the Typography Department in Earley Gate in Reading, but if there had been, and I wanted to make a nasty comment in the local paper, might I get my point across in this way?

Just Kidding

Watergate

abnegate

instigate

Irangate

investigate

interrogate

Westlandgate

profligate

litigate

Contragate

fumigate

expurgate

Lychgate

Starting gate

Wrought iron gate

Aldgate

Lancaster Gate

Notting Hill Gate

Earley Gate

These examples raise two questions. First, what is the optimum size of linguistic units in which typographic effects are best perceived? And then, which kinds of typeface are best able to convey semantic contrasts? And which typefaces are so similar that, if they were used, no effect would be perceived at all (other than by professional typographers). These, it seems to me, are the sort of questions that a typographical linguistics would explore. Would the sainthood effect in the first poem be apparent if an arbitrary typeface were chosen, instead of one which already has some kind of religious association? The poem might then look like this.

Peter and Paul.

Simon and Garfunkel.

Matthew and Mark.

Peters and Lee.

Luke and John.

Captain and Tenille.

Cosmas and Damian.

George and Ringo.

Cyril and Methodius.

Nina and Frederick.

Elvis?

Such examples do not of course tell the whole story about the way typography can enter into poetic expression. For instance,

they do not say anything about the way spatial orientation can directly influence linguistic meaning. That it can is clear from this short epic by José Paolo Paes, translated by Edwin Morgan, called “The Suicide, or Descartes à Rebour’s”:

cogito

ergo

boom

To print this simply as *cogito ergo boom* would be to nullify the effect. There needs to be time to allow the build-up of a linguistic expectation which is then frustrated, when *sum* becomes *boom*. The spacing permits this to happen.

My examples also take us in the direction of concrete poetry, though there the kind of symbolic “meanings” typically conveyed are moving away from the more precisely specifiable linguistic effects which I am looking for above. A linguistic effect, by definition, works within a particular language. Insofar as a particular concrete poem is recognizable by people from different languages (e.g., a text shaped like an altar or a pair of wings), it has moved away from the strictly linguistic into a world which is more directly iconic (though a truly universal iconicity may be chimerical, given the many cultural factors which condition our abilities to perceive and interpret). There is doubtless a continuum here, along which one could locate texts which balance the linguistic and the iconic elements in different degrees. This is another direction for future inquiry.

These examples and questions are only some of the topics which would enter into the domain of a putative typographical linguistics. There are many others – as can be seen if we pursue the analogy with phonology a little further. Phonologists have spent a great deal of time analyzing what exactly it is that makes phonemes different: it is possible to show that /p/ is different from /b/ in just one salient respect – one unit uses vocal cord vibration, the other

does not. Similarly, *tick* is different from *kick* because of the location of the tongue. *Tick* is different from *sick* not because of the location of the tongue but because of the way the sound is released – sharply, in the case of *tick* (it is a plosive); gradually, in the case of *sick* (it is a fricative). A sophisticated descriptive apparatus has been devised to enable phonologists to talk about sound systems, and the thrust of much current research is to find the best way of developing this apparatus so it can be used with equal plausibility for all the languages in the world.

A typographical linguistics would follow up this perspective for printed language. What exactly is it that makes graphemes different? The descriptive apparatus already available seems much more rudimentary, by contrast with phonology – much more selective, concentrating on the most noticeable features (such as uppercase, lowercase, ascender, descender). There seems to be nothing like the International Phonetic Alphabet, which identifies all the sounds the vocal tract can make that are capable of playing a role in some language. But why should there not be an International Graphetic Alphabet, identifying all the marks the human hand can make that are capable of playing a contrastive role in some language – the array of straight lines of varying length and orientation, curves, dots, thicknesses, and so on, which when combined result in written letters, syllables, and logograms? This would enable us to ask some interesting questions, such as: Are some writing systems more economical than others, in terms of the number of graphic features they employ? Which graphic features carry the most functional load, in a language? Which writing systems permit most graphic ambiguity? The answers to such questions are a long way off, given that there is not even a standard terminology for describing the distinctive letters and diacritics that turn up in different alphabets, let alone non-alphabetic systems. The need to standardize nomenclature has received a considerable boost with the emergence of languages other than English on the Internet, but we are nowhere near a universal framework for describing typographic contrasts with structural relevance. And it is this which I see as a long-term goal of a typographical linguistics.