Phonology as Cognition

Mark Hale
Concordia University, Montréal
hale1@alcor.concordia.ca

Charles Reiss
Concordia University, Montréal
reiss@alcor.concordia.ca

1 Form and substance in phonology

This paper attempts to ground phonology within psychology. That is, we are interested in phonology as a branch of the study of mental representation, the psychology of mind. In order to develop this 'phonology of mind' we need to understand the relationship between form and substance in linguistic representation. A coherent account of this distinction has yet to be proposed for either phonology or syntax. We attempt to contribute to this necessary inquiry in the domain of phonology by first defining 'form' and 'substance', then critiquing some recent work which implicitly or explicitly touches on the relationship between the two. We will argue that current trends in phonology fail to offer a coherent conception of form and substance and are also inconsistent with basic principles of science. Since we are not proposing a complete alternative model of phonology, we invite the reader to reflect on how our proposals could be implemented or on how our assumptions (which we believe are widely shared in principle, if not in practice) should be modified.

It has proven quite useful for linguists to conceive of a grammar as a relationship between i) a set of symbols—entities such as features and variables, constituents like syllables, feet, NPs, *etc.*; and ii) a set of computations—operations whose operands are drawn from the set of symbols, such as concatenation, deletion, *etc*. The set of symbols and relations together describe the formal properties of the system. Relevant questions in discussing formal properties include 'Is the system rule- and/or constraint-based?'; 'Do operations apply serially or in parallel?'; and 'Are there limits on the number of operands referred to in the course of a given phonological computation?'

The issue of substance essentially arises only with respect to the set of symbols and the extent to which their behavior in phonological computation is driven by what they symbolize. For the sake of simplicity we restrict ourselves in this discussion to the set of phonological primitives known as distinctive features and to the representations which can be defined as combinations of distinctive features.

We will concentrate in this paper on this notion of substance in phonological representation. In brief the question we are interested in is the following:

(1) Do the phonetic correlates (*i.e.*, the substance) of a particular distinctive feature or feature bundle have any non-arbitrary bearing on how that feature or feature bundle is treated by the computational system?

It is trivial to show that languages differ in that their computational systems treat specific features or feature bundles differently — for example, Standard German has coda obstruent devoicing and English does not. From this we can conclude that languages *can* treat the same symbols differently. A more challenging problem arises when we find an apparent example of cross-linguistically universal, seemingly non-arbitrary treatment of a feature or feature bundle. In such cases we must ask ourselves the following:

(2) Is the observed pattern a reflection of substantive constraints on the computational system (i.e., the grammar), or is the pattern due to other causes?

Other *a priori* plausible causes include, as we shall show in what follows, the process of language change, the nature of the language acquisition device, sampling errors, *etc*. From the standpoint of grammatical theory, factors such as sampling errors are obviously uninteresting. However, language change and the nature of the learning path are also, strictly speaking, not part of grammatical

theory. The modular approach to linguistics, and to science in general, requires that we both model the interactions between related domains, and also sharply delineate one domain from another. Occam's razor demands that, in doing so, we avoid redundancy and the postulation of unnecessary factors.

Even before proceeding to our argument that generalizations that bear on patterns of phonetic substance are not relevant to phonological theory as we define it, we can see that there is potentially much to gain from this modular approach in that it posits that universal phonology should be studied not just across languages, but also across modalities. What is shared by the phonologies of signed and spoken languages? We believe that phonology consists of a set of formal properties (*e.g.*, organization into syllables and feet, feature spreading processes) that are modality independent and thus not based on phonetic substance. The goal of phonological theory should be to discover these formal properties.[1] Failure to appreciate this goal has resulted in rampant 'substance abuse' in the phonological community.

We discuss various aspects of substance abuse in sections 2-5. In section 6, we offer a modest contribution to a substance-free phonology. In section 7, we return to substance with a discussion of the putative phenomenon of phonetic enhancement in grammars. Section 8 ties together the preceding sections with arguments against functionalist 'explanation' in linguistics. We argue that *dysfunctionalist* reasoning fares as well as its better known rival. Section 9 provides a concluding plea for a modular approach to the study of sound patterns in human languages.

2 THREE EXAMPLES OF SUBSTANCE ABUSE IN GRAMMATICAL THEORY

2.1 *Positional Faithfulness in Beckman (1997)*

Beckman (1997) proposes the constraints in (3a-b) as members of the universal constraint set:

(3) a. IDENT-σ_1(hi)
A segment in the root-initial syllable in the output and its correspondent in the input must have identical values for the feature [high].

b. IDENT(hi)
Correspondent segments in output and input have identical values for the feature [high].

As Beckman explains, this set of constraints allows faithfulness to a feature, like [high], to be maintained in some contexts, but not others, since the context sensitive constraint (3a) can be ranked above a markedness constraint that is violated by, say, the presence of high vowels, *HIGH, which in turn is ranked above the general constraint in (3b). In other words the ranking in (4) will allow surface high vowels only in root-initial syllables.

[1]This paper expands on Hale & Reiss (2000). We are grateful to audiences at *Montreal-Ottawa-Toronto Phonology Workshop 1998* at the University of Ottawa and at the Berkeley Phonology Laboratory, as well as to Noel Burton-Roberts, Morris Halle, Bill Idsardi, Madelyn Kissock, Afton Lewis, Jean-Philippe Marcotte and Ida Toivonen for discussion and challenging criticism that improved the paper. The authors' names appear in alphabetical order.

(4) IDENT-σ_1(hi) >> *HIGH >> IDENT(hi)

This is assumed to be a welcome result:

> The high ranking of positional faithfulness constraints, relative to both the more general IDENT constraints and markedness constraints, yields the result that features and/or contrasts in *just those positions which are psycholinguistically or perceptually salient* are less susceptible to neutralisation than in other locations which are not protected. [Beckman 1997:8. emphasis in original.]

Beckman (p.5) cites more than ten psycholinguistic studies to support her claim that word-initial material is more salient than medial or final material.[2] We believe that the correct conclusion to be drawn from this psycholinguistic evidence is the *exact opposite* of that which Beckman draws.[3] Encoding the findings of psycholinguistic experimentation in the grammar is a mistake because it is possible to achieve the same empirical coverage without positing new mechanisms like positional faithfulness.[4] Consider the following alternative account.

We know that children acquire spoken language primarily on the basis of acoustic input from speakers in their environment, with UG providing constraints on the hypothesis space.[5] We also know that phonological contrasts are best distinguished and recalled when occurring in certain positions. Imagine a child exposed to a language $\mathcal{L}_1$ which allows high vowels in all syllables — initial, medial and final. Imagine further that $\mathcal{L}_1$ has initial stress and that stress is realized as relatively increased duration and intensity. Given this scenario, it is easy to see that a child constructing $\mathcal{L}_2$ on the basis of ouput from $\mathcal{L}_1$ could consistently fail to acquire a contrast between mid and high vowels in relatively short, quiet syllables (those that are non-initial and thus unstressed), but succeed in acquiring this distinction in initial syllables, which are stressed and thus longer and louder. This type of relationship between $\mathcal{L}_1$ and $\mathcal{L}_2$ is known as 'sound change' (in particular, as a 'conditioned merger'). On the other hand, it is highly implausible that an acquirer would consistently fail to correctly analyze the mid/high contrast in longer, louder (stressed) syllables, yet successfully analyze the contrast in relatively short, quiet syllables. Note that this implausibility is manifest independently of our view of the nature of UG.

We see therefore that the existence of positional faithfulness phenomena can be understood as merely reflecting the nature of the learning situation and not a reflection of any grammatical principle:[6]

(5) If the acoustic cues of a given contrast in the target language are correctly analyzed by the acquirer in a context where they are relatively weak, they will also be analyzed correctly in a context where they are relatively strong.

[2]It is unclear whether this generalization would hold, say, in a language with non-initial stress. It is also unclear whether Beckman's extension of psycholinguistic findings concerning word-initial syllable to *root*-initial syllables is justified. However, we will assume, for purposes of this discussion, that Beckman has stated the relevant generalizations correctly.

[3]We wish to stress that we are not singling Beckman out for any reason except for the fact that her paper appeared recently in a widely read journal and is well-written and clear in its arguments and assumptions.

[4]For other arguments against context-sensitive faithfulness see Reiss (1996:315).

[5]It may be a useful idealization to assume that UG does not just constrain the learning path, but completely determines it. We suspect that such a position will prove most fruitful in sketching an explicit theory of acquisition, but justification for this goes beyond the scope of this paper.

[6]This idea is discussed more thoroughly in Hale (forthcoming).

Note that (5) is essentially definitional, since the strength, or acoustic salience, of a contrast is just a measure of how easy it is to perceive. What is most important to understand is that the theory proposed here is not meant to *replace* a synchronic account of the data. So, the best synchronic analysis must somehow be able to generate vowel neutralization in noninitial syllables. (5) is meant to guide us in choosing a theory of grammar in which to couch that synchronic account, but (5) is not part of the grammar. Whatever theory of phonology one adopts, it must be able to synchronically generate the type of patterns that Beckman describes, but the predictions generated by the correct theory, *qua* phonological theory, need not replicate the predictions derivable from (5).

By adopting the view of sound change proposed here, we see that many supposedly phonological tendencies, or markedness patterns, are actually emergent properties, that is epiphenomenal. 'Positional faithfulness' is due, not to the nature of *phonology*, but to the 'sifting effect' of acquisition on the incidental, arbitrary nature of the phonetic substance. Since effects such as those observed by Beckman already have a coherent extragrammatical account within acquisition theory (and it is necessary, in any event, to have an acquisition theory), building positional faithfulness into a theory of universal phonology is a misuse, or abuse, of phonetic substance in theory construction.

2.2 */r/-insertion in McCarthy (1993)*

McCarthy's (1993) discussion of intervocalic *r*-insertion in Massachusetts English is fairly well-known, so an example should be sufficient for illustration. In this dialect, an underlying sequence like *Wanda arrived* is realized with a 'linking' [r]: *Wanda[r] arrived*. As McCarthy himself notes (and as discussed by LaCharité & Paradis 1993 and Halle & Idsardi 1997) "r is demonstrably not the default consonant in English" (189). That is, it is not the maximally unmarked consonant that an OT account predicts would emerge in such a situation. In order to account for the insertion of [r] McCarthy proposes a special *rule* of *r*-insertion: "a phonologically arbitrary stipulation, one that is outside the system of Optimality" (190). There are several problems with this proposal, many of which are insightfully discussed by Halle & Idsardi. However, we propose that one of their criticisms requires elaboration. Halle & Idsardi rightly point out that "reliance on an arbitrary stipulation that is outside the system of Optimality is equivalent to giving up on the whole enterprise (337)," but these authors do not discuss what we consider to be the most important point: grammars do contain arbitrary processes. McCarthy's grammar has an arbitrary component (containing rules like *r*-insertion) and a non-arbitrary component (containing the substantive OT constraints). Such a theory is empirically non-distinct from the theory we propose below, which posits that *all* grammatical computations are arbitrary with respect to phonetic substance. This is because the set of phenomena predicted to exist by our theory (with only arbitrary processes) is identical to the set of phenomena predicted to exist by McCarthy's theory (with both non-arbitrary and arbitrary processes). Since McCarthy must adopt a model which allows arbitrary phenomena (like *r*-insertion), the addition to the theory of a special subcomponent to account for alleged 'non-arbitrary' phenomena violates Occam's Razor.

The diachronic source of *r*-insertion is transparent—the relevant dialects also exhibit *r*-deletion in codas, so insertion reflects rule-inversion triggered by hypercorrection. Again, the diachronic facts do not make a synchronic account unnecessary, but they show us that basically idiosyncratic

historical events affect specific grammars — and, in part, how they may do so.

2.3 *Structural constraints on non-structures*

Perhaps one of the most problematic cases of substance abuse we have come across is McCarthy's (1994) appeal to parameterized constraints to account for opacity effects in Hebrew spirantization by invoking the notion of constraint schema. McCarthy makes some reasonable simplifying assumptions in this first attempt:

> I will assume that every constraint is a prohibition or negative target defined over no more than two segments, α and β. That is, the canonical constraint is something like $*\{\alpha,\beta\}$, with appropriate conditions imposed on α and β. These conditions are as follows:
>
> (i) a specification of the featural properties of α and β as individual segments.
>
> (ii) a specification of the linear order relation between α and β ($\alpha < \beta$, $\beta < \alpha$, or both in the case of mirror-image rules ...
>
> (iii) a specification of the adjacency relation between α and β (e.g., strict adjacency, vowel-to-vowel adjacency ...
>
> The decomposition of the conditions imposed by a phonological constraint will be crucial in accounting for the range of opacity phenomena. Even more important, though, is this: each condition—the featural composition of α, the featural composition of β, linear order and adjacency—must also name the level (underlying, surface, or either) at which it applies. Correspondence Theory allows us to make sense of conditions applying at one level or the other. As a bookkeeping device, I will state the constraints in the form of a table ...

We reproduce here the schema-based constraint that McCarthy proposes to account for Tiberian Hebrew Post-vocalic Spirantization.

(6) Constraint for Opacity in Hebrew Spirantization (McCarthy 1994)

	Condition	Level
α	V	Indifferent
β	[-son, -cont]	Surface
Linear Order	$\alpha > \beta$	Indifferent
Adjacency	Strict	Indifferent

As McCarthy says "In correspondence terms, the meaning of this constraint is this: the constraint is violated if a surface stop β or its underlying correspondent is immediately preceded by a vowel."

As pointed out in Reiss (1997), this powerful constraint type has several problems. First, it compromises the OT notion of a universal, innate constraint set by allowing apparently language-specific parameterized constraints. This may not be a serious problem, since it represents an attempt to define the form of possible constraints. In other words, McCarthy could be interpreted as

presenting a theory in which the intensional description of the set of constraints is universal, but languages vary in which constraints they actually incorporate (based on evidence presented to the learner).[7]

Most relevant to our present purposes, however, is the fact that such constraints undermine implicit and explicit appeal to phonetic grounding of well-formedness constraints in McCarthy's work. For example, McCarthy and Prince (1995:88) refer to a constraint *VgV as the "phonologization of Boyle's Law". It is incoherent to argue that a constraint is motivated by the facts of phonetics, when the structures which violate this constraint need not be surface structure strings. In fact, they need not exist as strings at *any level of representation*.

3 NEO-SAUSSUREANISM

The conclusion we wish to draw from the above examples and many others like them is that the best way to gain an understanding of the computational system of phonology is to assume that the phonetic substance (say, the spectral properties of sound waves) that leads to the construction of phonological entities (say, feature matrices) *never* reflects how the phonological entities are treated by the computational system. The computational system treats features and the like as arbitrary symbols. What this means is that many of the so-called *phonological universals* (often discussed under the rubric of markedness) are in fact epiphenomena deriving from the interaction of extragrammatical factors like acoustic salience and the nature of language change. It is not surprising that even among its proponents, markedness 'universals' are usually stated as 'tendencies'. If our goal as generative linguists is to define the set of *computationally possible* human grammars, 'universal tendencies' are irrelevant to that enterprise.

We therefore propose extending the Saussurean notion of the arbitrary nature of linguistic signs to the treatment of phonological representations by the phonological computational system. Phonology is not and should not be grounded in phonetics since the facts which phonetic grounding is meant to explain can be derived without reference to *phonology*. Duplication of the principles of acoustics and acquisition inside the grammar constitutes a violation of Occam's razor and thus must be avoided. Only in this way will we be able to correctly characterize the universal aspects of phonological computation.

John Ohala (*e.g.*, 1990) has done the most to demonstrate that many so-called markedness tendencies can be explained on phonetic grounds and thus should not be explained by principles of grammar. Examples discussed by Ohala include patterns of assimilation and the contents of phonemic inventories. For an extensive bibliography on this topic see Ohala (1998). We differ from Ohala in our use of the term 'phonology' (which for him covers *all* aspects of the sound systems of human language) but wholeheartedly endorse his approach.

[7] McCarthy does not explicitly make this argument, but it seems to us to be a better theory than the standard OT claim that all constraints are literally present in all grammars. Of course, adopting our suggested interpretation will force OT practitioners to revise their views on acquisition and, especially, *emergence of the unmarked*. This view of OT would also make it much closer to a theory of learned rules.

3.1 *Substance in* SPE

It is obvious that our proposal runs contrary to most of the discussion in Chapter 9 of *SPE* (Chomsky and Halle 1968). This chapter starts out with an 'admission' that the theory developed in the earlier chapters is seriously flawed:

> The problem is that our approach to features, to rules and to evaluation has been overly formal. Suppose, for example, that we were systematically to interchange features or to replace [αF] by [-αF] (where α is +, and F is a feature) throughout our description of English structure. There is nothing in our account of linguistic theory to indicate that the result would be the description of a system that violates certain principles governing human languages. To the extent that this is true, we have failed to formulate the principles of linguistic theory, of universal grammar, in a satisfactory manner. In particular, we have not made use of the fact that the features have intrinsic content. [400].

Later in the chapter Chomsky and Halle themselves acknowledge that, with the above-quoted assertion, they are on the wrong track:

> It does not seem likely that an elaboration of the theory along the lines just reviewed will allow us to dispense with phonological processes that change features fairly freely. The second stage of the Velar Softening Rule of English (40) and of the Second Velar Palatalization of Slavic strongly suggests that the phonological component requires wide latitude in the freedom to change features, along the lines of the rules discussed in the body of this book [428].

In other words, Chomsky & Halle ultimately recognize that the truly important parts of the phonology, in the sense of the ones that are unnatural, are those which cannot be derived from functional considerations of naturalness. This conclusion is echoed elsewhere: "Where properties of language can be explained on such 'functional' grounds, they provide no revealing insight into the nature of mind. Precisely because the explanations proposed here are 'formal explanations,' precisely because the proposed principles are not essential or even natural properties of any imaginable language, they provide a revealing mirror of the mind (if correct)" (Chomsky 1971:44).

We propose that switching the feature coefficients as described in the first quotation might lead to the description of systems that are *diachronically* impossible human languages (ones that could never arise because of the nature of language change), but not to ones that are *computationally* impossible. The goal of phonologically theory, as a branch of cognitive science, is to categorize what is a computationally possible phonology, given the computational nature of the phonological component of UG.[8]

3.2 *Computation* vs. *Transduction: A Place for Substance*

It is important to note that the preceding discussion is not meant to imply that the mapping of sound to features is arbitrary. It is only the treatment of phonological representations within the

[8] This argument, as well as other ideas in this paper, was anticipated by Hellberg (1980) See also Burton-Roberts, this volume, section 5.

computation that is arbitrary. Articulatory and acoustic substance *are* related to the representations we construct, but not within the grammar. The nature of this relationship is part of the theory of *transduction*—the mapping between the physical and the symbolic (Pylyshyn 1984). As Bregman (1990:3) points out "In using the word representations, we are implying the existence of a two-part system: one part forms the representations and another uses them to do such things as calculate . . . " Bregman is concerned with the auditory system which does not have an output module – in discussing language, we also need to model output transducers that map from surface (featural) representations to articulatory gesture. For our purposes, Bregman's distinction corresponds to speech perception (construction of featural representations, ultimately from auditory signals) and grammar, which performs symbolic computation. We know from the existence of visual and auditory illusions that the transduction process is not simple. The perceptual system does not just form a direct record of physical stimuli. As Bregman points out, we know that representations are being constructed, because only then could they be constructed incorrectly, leading to illusions.

Pylyshyn (1984) provides the following discussion: "This, then is the importance of a transducer. By mapping certain classes of physical states of the environment into computationally relevant states of a device [*e.g.* a human], the transducer performs a rather special conversion: converting computationally arbitrary physical events into computational events. A description of a transducer function shows how certain nonsymbolic physical events are mapped into certain symbolic systems" (152). Pylyshyn points out that the "*computationally relevant* states are a tiny subset of [a system's] physically discriminable states" and that the "former are typically a complex function of the latter" (150). In (7) we paraphrase Pylyshyn's criteria for a psychological transducer (153-4), that is a transducer from physical signals to representations.

(7) Criteria for a psychological transducer

- The function carried out by a transducer is itself *nonsymbolic*; it is part of the functional architecture of the system.
- A transducer is stimulus bound, operating independent of the cognitive system.
- The behavior of a transducer is described as a function from physical events to symbols:
  a. The domain of the function (the input) is couched in the language of physics.
  b. The range of the function (the output) must be computationally available, discrete atomic symbols (for example, feature matrices).
  c. The transformation from input to output must be follow from the laws of physics.

This is where issues of substance arise—the physical aspects of the acoustic signal serve as the input into the transducer function. From that point on, in the manipulations of the constructed symbolic representations, substance is irrelevant to computation. Only the *formal* properties of such representations are relevant to the computational system.

It is worth contrasting Pylyshyn's well-articulated modular approach to that of Prince & Smolensky (1993) who directly reject the kind of extreme formalist position we advocate here.

> We urge a reassessment of this essentially formalist position. If phonology is separated from the principles of well-formedness (the 'laws') that drive it, the resulting loss of constraint and theoretical depth will mark a major defeat for the enterprise [Prince & Smolensky 1993: 198, see also p.3].

This view of the goals of phonology stems from a failure to observe the critical transducer vs. grammar distinction, that is, from extensive 'substance' abuse. It is also at odds with the well-established goals of cognitive science in general:

> [I]f we confine ourselves to the scientific and intellectual goals of understanding psychological phenomena [as opposed to predicting observed behavior–mh & cr] one could certainly make a good case for the claim that there is a need to direct our attention away from superficial "data fitting" models toward deeper structural theories" [Pylyshyn 1973:48].

As our discussion of markedness below will indicate, we do not believe that there is any reason to believe that 'principles of well-formedness' exist, aside from those that constrain the set of possible representations. That is, we find that the evidence for markedness based constraints to be unconvincing.

The 'principles of well-formedness' that Prince & Smolensky refer to and adopt as the basis of OT constraints are merely derived from the heuristic devices that constitute the intuitions of an experienced linguist. For example, we may intuitively believe that a sequence like [akra] will more likely be syllabified as [a.kra] rather than as [ak.ra] in a random sample of grammars, although both syllabifications are found, for example, in the Ancient Greek dialects. Lacking information to the contrary, it may be useful to assume that the more common syllabification is present in a new, unfamiliar language. This will allow the formulation of hypotheses that may then be tested, and the guess will turn out to be correct more often than not, if our intuitions have any basis. However, it is a mistake to assume that our intuitions reflect the nature of the system we are studying in any direct manner. The intuition that heavy things fall faster than light things is very useful when someone drops something from a window, but the intution needs to be transcended to understand the workings of gravity. Heuristics are used by the analyst to make useful guesses about data, and guesses can be wrong. This is why OT constraints need to be violable, unlike all other scientific laws.

The pervasiveness of such 'data-over-principles' approaches to phonology can be appreciated by the following quote from an influential pre-OT paper: "The goal of phonology is the construction of a theory in which cross-linguistically common and well-established processes emerge from very simple combinations of the descriptive parameters of the model" (McCarthy 1988). By concentrating on what is 'common', rather than what is possible, phonology will provide (or rather has provided) plentiful material for descriptive work at some level of sophistication, but it is clear that no science should be concerned with making it particularly simple to express that which happens often. The goal of any science is to define a coherent domain of inquiry and to establish a common vocabulary for *all* events in that domain. This involves reducing the common *and* the rare events (*e.g.*, planetary motion and the Big Bang) to special cases of an abstract set of primitive notions. All of this suggests that while a change of course for phonological theory was definitely needed in the early nineties, Optimality Theory has been a change in exactly the wrong direction.

3.3 *Acoustophilia: a warning*

Sapir (1925:37) points out that "it is a great fallacy to think of the articulation of a speech sound as [merely] a motor habit". A corresponding error is committed in many of the studies (e.g., Flem-

ming 1995) that argue for the increased use of acoustic information to model human phonological computation. This work tends to establish units of analysis in terms of measurements taken over the acoustic signal itself. We believe that this technique shows the negative effects of 'acoustophilia' – the mental state arising from the deep and abiding satisfaction which comes from having *something* concrete to measure, in this case the acoustic signal. There is, we believe, a fairly serious difficulty with such an approach: we know with a great deal of confidence that human perception does not show the kind of direct dependency on the signal which the methodology of the acoustophiliacs requires.[9] This attitude towards the study of language echoes the overly positivist brand of empiricism adopted by the behaviorists, an attitude that was discredited already in the nineteen fifties.

An example may make this clearer. Flemming 1995 argues from an examination of F_2 interactions in an experimental setting that it is necessary to have the grammar generate a statistical pattern which forms a reasonable match to his experimental results. A parallel from the field of the cognition of vision would examine the properties of an image as measured with, e.g., a photometer, and require of us that our 'grammar of vision' generate a representation like that measured on the page. So, in Figure 1 below, it would require – since the triangle we see is of precisely the same color and brightness as the background (as can be verified by the use of a photometer) – that we construct a human visual system that does *not* see the triangle projecting from the page. This is of course the wrong result – the human visual system, given the input in Figure 1, constructs a 'percept' which is very different from the patterns we might infer from photometric readings (see Hoffman 1998).