Psychological Review
2004, Vol. 111, No. 1, 261–273

0033-295X/04/$12.00 DOI: 10.1037/0033-295X.111.1.261

Holographic Goodness Is Not That Bad: Reply to Olivers, Chater, and Watson (2004)

Peter A. van der Helm and Emanuel L. J. Leeuwenberg
University of Nijmegen

The holographic approach (HA) to goodness (P. A. van der Helm & E. L. J. Leeuwenberg, 1991, 1996, 1999) is an ideal-observer theory at, in D. Marr's (1982) terms, the computational and algorithmic levels of description. It provides an explanation of the detectability of visual regularities such as mirror symmetry, repetition, and Glass patterns. C. L. N. Olivers, N. Chater, and D. G. Watson (2004) gave a picture of HA as if it were a flawed theory. However, they gave a flawed picture containing factual errors and misconceptions. Most of their alleged counterevidence actually supports HA's unified account of perfect regularities, perturbed regularities, and nested regularities. Recent evidence indicates that HA may even lead to deeper implementational insight into the processing of spatial frequencies.

In van der Helm and Leeuwenberg (1991, 1996, 1999; henceforth vdHL91, vdHL96, and vdHL99, respectively), we developed the holographic approach (HA) to goodness. Goodness refers to the detectability of single regularities and combinations of nested regularities, whether or not perturbed by noise. Olivers, Chater, and Watson (2004; henceforth OCW) argued that HA is inadequate on both theoretical and empirical grounds. In this reply to OCW, however, we not only show that HA is much more adequate than they suggested but also rebut their underlying complaint that HA is an ideal-observer theory that does not take implementational peculiarities of the human system into account.

We first sketch the basics of HA's goodness model to give a framework for addressing the latter issue. We then give an overview of HA's theoretical and modeling ingredients to set the stage for our evaluation of OCW's theoretical and empirical arguments.

The Holographic Goodness Model

HA's goodness model, as proposed in vdHL96, is very simple insofar as it concerns perfect regularities. It is based on vdHL91's new formalization of regularity, according to which mirror symmetries have a point structure, repetitions a block structure, and Glass patterns a dipole structure (see Figure 1). The model uses these internal structures to quantify the goodness of a regularity by the weight of evidence (W) for the regularity in the stimulus (cf. MacKay, 1969). That is, it is quantified by $W = E/n$, where n is the number of elements in the stimulus (the dots in Figure 1) and E the number of nonredundant identity relationships (the arcs in Figure 1) between the substructures of the regularity.

The model thus predicts correctly the well-known phenomenon that symmetries and Glass patterns are about equally good and better than repetitions. More specifically, it quantifies the goodness of symmetries by a W-load of $W = .5$; Glass patterns get nearly the same W-load, but repetitions get a lower W-load that, moreover, depends heavily on the number of stimulus elements. Furthermore, by way of the holographic bootstrap model, vdHL99 showed that these strength differences can be translated faithfully into differences in detection speed. According to this translation, detection advances linearly for repetition but exponentially for symmetry and Glass patterns.

As argued in vdHL96 and vdHL99, HA's goodness models have considerable explanatory power. This assessment, however, is precisely what OCW challenged. Our reply begins by discussing the background against which their critique is to be seen.

Regularity Detection: Goal, Method, and Means

We agree with OCW (p. 242) that goodness is a property that may have many stimulus- and task-dependent appearances in the empirical practice. For the rest, however, OCW's scientific stance seems to be quite different from ours. This difference can be made more explicit by way of Marr's (1982) distinction among three levels at which information processing systems may be described: first, the computational level, at which a system's goal is specified; second, the algorithmic level, at which a system's method is specified; and third, the implementational level, at which a system's means are specified. As we discuss next, OCW seem to give nearly exclusive rights to the implementational level, whereas we concur with Marr's idea that perception research should aim at compatible (i.e., complementary and, necessarily, partly overlapping) descriptions at all three levels.

OCW (p. 259) argued that the variety of empirical goodness appearances is best accounted for by a variety of implementational processing factors related to, for instance, attention, spatial proximity, scale, optic flow, orientation, perspective, continuity, and common fate. Consider, for instance, Palmer's (1983) transformational approach at the computational level and Wagemans, Van Gool, Swinnen, and Van Horebeek's (1993) original bootstrap model at the algorithmic level. In the transformational approach, symmetry, repetition, and Glass patterns get a block structure, and in the original bootstrap model, all three regularities are assumed

Peter A. van der Helm and Emanuel L. J. Leeuwenberg, Nijmegen Institute for Cognition and Information, University of Nijmegen, Nijmegen, The Netherlands.

Correspondence concerning this article should be addressed to Peter A. van der Helm, Nijmegen Institute for Cognition and Information, University of Nijmegen, Montessorilaan 3, 6525 HR Nijmegen, The Netherlands. E-mail: peterh@nici.kun.nl

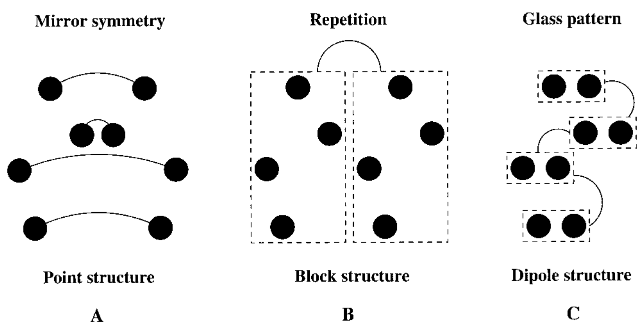

Figure 1. The holographic structure of visual regularities. A: Mirror symmetry is constituted by identities (the arcs) between single dots. B: Repetition is constituted by identities between repeats (the rectangles). C: The Moiré structure in Glass patterns is constituted by identities between dipoles (not between single dots, as Olivers et al., 2004, suggested).

to have a point structure. Hence, neither approach as such can account for even the most basic goodness differences between these regularities. To this end, both approaches have to resort to implementational factors. By the standards OCW applied to HA, this would mean that both approaches as such can be dismissed.

In this context, the most prominent implementational factor is the one called *local attention* or *proximity:* Stimulus elements that are closer to each other can be matched more easily. This factor, by the way, should not be confused with the Gestalt law of proximity, which is not about matching but about grouping. At first glance, "matching-by-proximity" seems to explain that symmetry and Glass patterns are better than repetition. Symmetry detection can already start to be successful by matching nearby stimulus elements around the axis of symmetry, whereas repetition detection can only start to be successful by matching elements that are one repeat apart. Similarly, in Glass patterns, corresponding elements are generally closer to each other than in repetition.

There are, however, indications that go against matching-by-proximity. First, it is true that a separation space between stimulus halves renders symmetry detection more difficult; however, it renders repetition detection less difficult (Corballis & Roldan, 1974). Second, imagine two identical dot patterns that overlap partially, that is, such that they form a configuration that is between a juxtaposition as in repetition and a nearly complete overlap as in Glass patterns. Matching-by-proximity would imply that the identity of such partially overlapping dot patterns can be detected more easily than the identity of two juxtaposed dot patterns as in repetition. However, such a stimulus is generally perceived as a random pattern.

One might argue that in the latter case the matching of corresponding elements is overruled by "grouping-by-proximity" of noncorresponding elements. First, however, it should then also be overruled in repetition, but it is not. Second, Maloney, Mitchison, and Barlow (1987) did not find such overruling: Glass patterns remained detectable even when every dipole dot had 6–10 noise dots closer by than its dipole mate.

The foregoing illustrates that most proximity effects are in fact phenomena that still need an explanation. That is, proximity may (co)determine the arena of processing, but it does not determine what happens inside that arena. The transformational approach and the original bootstrap model enter that arena, but as noted, they have to resort to proximity for the most basic goodness effects. HA, conversely, starts from a differentiation between the internal structures of regularities, not only at the computational level by way of the $W = E/n$ model but also at the algorithmic level by way of the holographic bootstrap model. As we reinforce in this article, the goodness effects that OCW attributed to proximity and other implementational factors are the same effects that in a unified way are implied by HA's structural differentiation. Hence, regarding many goodness effects, HA simply does not need to resort to implementational factors.

One might still object, as OCW (p. 245) did, that HA is an ideal-observer theory that does not take implementational peculiarities of the human system into account. However, the purpose of an ideal-observer theory is not to deny influences of such peculiarities but, rather, to bring unification so that one does not have to resort to postulating a new peculiarity each time a new perceptual phenomenon is discovered. In fact, it creates an ideal–nonideal contrast that gives better sight on what the implementational peculiarities actually are, thereby leading to a better understanding of the functionality of brain structures and brain processes. Later in this article, we substantiate this by concrete examples of how HA relates to functional and neural mechanisms in vision.

From the foregoing, we conclude that OCW's metatheoretical arguments may indicate that our goodness approach, like any approach thus far, is incomplete but not that it is inadequate. We can therefore turn to OCW's concrete arguments against our goodness model. We first pinpoint OCW's misconceptions about HA, in the context of an overview of HA's theoretical basis.

The Holographic Approach (HA)

HA's theoretical basis contains a mathematical part and a psychological part. The mathematical part comprises vdHL91's new formalization of the notion of *regularity* in terms of 1-D symbol strings. The psychological part comprises the generalization from regularity in 1-D symbol strings to regularity in 2-D visual stimuli. OCW's misconceptions concern both parts. We first summarize vdHL91's formalization in a stepwise fashion that may show that OCW's allusions to circularity are misplaced.

Holographic Regularity and Transparent Hierarchy

The first step in vdHL91's formalization is the introduction of *identity chains* to characterize arbitrary configurations of identity relationships among symbols in a symbol string. For example, the configuration in the string *abcpabc* can be characterized by, among others, the one-identity chain {(1 2 3) = (5 6 7)}, which simply represents that the substring consisting of the first, second, and third symbols is identical to the substring consisting of the fifth, sixth, and seventh symbols. The same configuration can also be characterized by the three-identity chain {(1) = (5), (2) = (6), (3) = (7)}. This is an ordered set of identities, which has 2 two-identity *subchains* (consisting of successive identities in the three-identity chain), namely, {(1) = (5), (2) = (6)} and {(2) = (6), (3) = (7)}. Furthermore, vdHL91 introduced *identity structures,* that is, sets of identity chains (with the same number of identities) that are alike in a sense that can be pinpointed mathematically. The core of this alikeness may be illustrated as follows: By way of the substitutions $y = ab$ and $z =$ pq, not only the strings *yzy* and *abpqab* can be said to be alike but also the identities (1) = (3) in *yzy* and (1 2) = (5 6) in *abpqab.*

Second, vdHL91 noted that among the infinite number of identity structures, some *n*-identity structures have the so-called *holographic* property that for each fixed m ($1 \leq m < n$), all *m*-identity substructures are the same. In other words, for all *n*-identity chains in such a holographic *n*-identity structure, all *m*-identity subchains belong to the same *m*-identity structure. Furthermore, vdHL91 discovered that such a holographic *n*-identity structure can be uniquely expanded, one identity at a time, into larger identity structures that are also holographic, leading to an infinite set of coherent *n*-identity structures ($n = 1, 2, \ldots, \infty$). Such a set is called a *holographic regularity.* As vdHL91 discovered, there are only 20 holographic regularities, each of which can be described by 4 slightly different coding rules yielding, in total, 80 holographic coding rules.

For instance, the holographic iteration rule (I-rule; see Table 1) captures configurations that comprise an indefinite number of identical juxtaposed substrings that each consist of an indefinite number of symbols. Furthermore, the holographic symmetry rule (S-rule; see Table 1) captures configurations that comprise an indefinite number of nested pairs of identical substrings that each consist of an indefinite number of symbols. In line with common sense, vdHL91 gave the label *repetition* to the holographic regularity that is described by the I-rule and the label *bilateral symmetry* to the one that is described by the S-rule. If a bilateral symmetry is formed by pairs of substrings, each consisting of only one symbol, then it is called a *mirror symmetry;* otherwise it is called a *broken symmetry* (cf. Weyl, 1952; see also the Glass Patterns Can Destroy or Boost Symmetry section below).

The aforementioned holographic expansion steps can now be illustrated as follows: For the I-rule, such a *growth* step is, for instance, the expansion of $4 * (y)$ into $5 * (y)$, involving the inclusion of one extra substring *y.* Because the substring *y* may consist of an arbitrary number of symbols, vdHL96 concluded that repetition has a *block* structure. Furthermore, for the S-rule, such a growth step is, for instance, the expansion of $S[(y)(z)]$ into $S[(x)(y)(z)]$, where the substrings *x, y,* and *z* may consist of an arbitrary number of symbols. If *x, y,* and *z* consist of only one symbol each, then both before and after the expansion, it concerns a mirror symmetry. Because a symbol is the smallest string part, vdHL96 concluded that mirror symmetry has a *point* structure.

As the third and final step, vdHL91 discovered that only 9 of the 80 holographic coding rules allow for so-called *transparent hierarchy:* This means that a hierarchical code of a symbol string always corresponds unambiguously to only one hierarchical organization in the symbol string. They noted further that everything that can be said and done with these 9 transparent holographic coding rules can also be said and done with only 4 of them, namely, the I-rule, the S-rule, and the two alternation rules (A-rules; see Table 1). Therefore, vdHL91 ended with the proposal to use only these *ISA-rules* in a visual coding model.

Table 1
Examples of Transparent Holographic Coding Rules and Encoding

Rule	String	Code
T-rule	$k_1k_1\ k_2k_2 \ldots k_nk_n$	$T[(k_1)(k_2)\ldots(k_n)]$
I-rule	$k\ k\ k \ldots k\ k$	$m * (k)$
S-rule	$k_1k_2 \ldots k_s\ p\ k_s \ldots k_2k_1$	$S[(k_1)(k_2)\ldots(k_s), (p)]$
A-rules	$k\ x_1\ k\ x_2 \ldots k\ x_n$	$\langle(k)\rangle/\langle(x_1)(x_2)\ldots(x_n)\rangle$
	$x_1\ k\ x_2\ k \ldots x_n\ k$	$\langle(x_1)(x_2)\ldots(x_n)\rangle/\langle(k)\rangle$
Encoding	*a b a b a b r s t p t r s*	$3 * (ab)\ S[(rs)(t), (p)]$
	p b c c b p b c c b	$2 * (p\ S[(b)(c)])$
	a b a b a b b a b a b a	$S[3 * ((a)(b))]$
	a b p a b q r a b s	$\langle(ab)\rangle/\langle(p)(qr)(s)\rangle$
	b a c a c a t a	$\langle(b)\ 2 * ((c))\ (t)\rangle/\langle(a)\rangle$

Note. The codes are possible codes, that is, not necessarily simplest codes. Furthermore, a code remains possible when in string and code, identical symbols are substituted by identical strings or codes. See the text for details on the T-, I-, S-, and A-rules.

Pinpointing Olivers et al.'s (2004) Misconceptions

OCW's theoretical arguments were aimed mainly at undermining our claim that HA implies that mirror symmetry gets a point structure and repetition a block structure. This implication of HA relies on the usage of only the ISA-rules, and among other things, OCW argued that by our own standards other coding rules qualify as well. As we evaluate next, however, their arguments suffer from four classical confoundings.

First, curiously, OCW (Footnote 3) advised us to ignore the nonambiguity demand in our definition of transparent hierarchy (which would give room to other coding rules that imply a point structure for repetition). After all, they argued, allowing hierarchy ambiguity would agree with the fact that we also allow multiple interpretations of one and the same stimulus. However, stimuli are inherently ambiguous, whereas codes should be unambiguous. That is, on the one hand, any stimulus has many interpretations that are possible in principle. In fact, a main objective of perception research is to find out how the human system solves this ambiguity, that is, how it succeeds in arriving at, generally, only one interpretation. On the other hand, if one investigates this stimulus ambiguity by means of codes that each stand for one specific interpretation, then the codes clearly should be unambiguous; otherwise, the coding system would be unreliable, just as ambiguous computer codes cause computer crashes. Hence, here, OCW's attempt to undermine our claim fails because they confounded stimulus ambiguity and code ambiguity.

Second, OCW (p. 246) argued that vdHL96 too quickly dismissed the transparent holographic *T-rule* (see Table 1). In fact, however, vdHL91 (p. 197) already assessed the T-rule as one of the redundant transparent holographic coding rules. For instance, the symbol strings *aabbccdd* and *abcdabcd* can be encoded into $T[(a)(b)(c)(d)]$ and $T[(abcd)]$, respectively, but by way of the I-rule, they can also be encoded into $2 * (a)\ 2 * (b)\ 2 * (c)\ 2 * (d)$ and $2 * (abcd)$, respectively, yielding the same complexities, hierarchical organizations, and *W*-loads. Hence, there simply is no reason to include the T-rule. OCW nevertheless tried to show that a point structure for repetition is implied not only by the T-rule but also by the so-called *P-rule,* which enables the encoding of a symbol string like *aabbccdd* into $P[(a)(b)(c)(d)]$. They tried to show this by means of scan paths in 2-D stimuli, which, however, have nothing to do with coding rules but with the way in which stimuli may be represented by symbol strings: Only once such a symbol string is given, coding rules come into play (see next subsection). Given this disentanglement, the P-rule can be nothing but identical to the T-rule, which, as before, remains redundant. Hence, this time OCW's attempt to undermine our claim fails because they confounded two coding steps, namely, the representation of a 2-D stimulus by a 1-D symbol string, on the one hand, and the encoding of such a symbol string by means of coding rules, on the other hand.

Third, OCW argued that the notion of holographic growth (see previous subsection) has no relevance for 2-D stimuli. Symmetrical or repetitive 2-D stimuli, they argued, can be expanded, at will, both pointwise and blockwise. First, however, the holographically pointwise growth of mirror symmetry is analogous to the gradual increase in body size of a (fairly symmetrical) person who grows bigger. Hence, OCW (see their Figure 4C) may add many separate point pairs to a mirror symmetry and call it a blockwise expansion, but in our view that is just a perversion of words. Second, the holographically blockwise growth of repetition is analogous to the growth of, for instance, a queue of (fairly identical) penguins. Clearly, the latter analogy applies to an increase in the number of penguins and not, as suggested by OCW (see their Figure 4F), to an increase in body size of individual penguins. This illustrates that holographic growth is about a stepwise expansion of a stimulus property like symmetry or repetition (adding one holographic identity per step) and not about an arbitrary expansion of a stimulus. Hence, this time OCW's attempt to undermine our claim fails because they confounded stimuli and stimulus properties.

Fourth, OCW (see their Figure 4G) argued that holographic growth implies that to expand a twofold mirror symmetry, one needs sets of four extra points (which, for some odd reason, they again called a blockwise expansion). However, they apparently forgot that holographic growth is about single regularities. That is, a twofold mirror symmetry is not one holographic regularity, but it is a hierarchical combination of two holographic regularities, namely, of 2 onefold mirror symmetries. OCW's (p. 248) argument in this respect looks like a game on the word *n-fold.* They seemed to have the idea that the (blockwise) structuring of an *n*-fold repetition (i.e., *n* juxtaposed repeats) should, merely because of the adjective *n-fold*, pertain to an *n*-fold mirror symmetry as well. This idea may be evoked by the traditional transformational approach, but it is not in line with HA. For $n \geq 2$, an *n*-fold repetition is in HA an instantiation of one holographic regularity, but as just indicated, an *n*-fold mirror symmetry is not. Hence, this time OCW's attempt to undermine our claim fails because they confounded different connotations of the term *n-fold.*

The foregoing shows that OCW failed to undermine our claim that HA implies that mirror symmetry gets a point structure and repetition a block structure, at least, insofar as this claim applies to symmetry and repetition conceived as stimulus properties. For a concrete symmetrical or repetitive 2-D stimulus, however, the foregoing still leaves open whether according to HA its perceptually preferred description is indeed the one that captures the symmetry or repetition property. After all, as OCW remarked correctly, any 2-D stimulus can be described in various ways, capturing various properties. In the next subsections, we address this issue, and we pinpoint further misconceptions by OCW.

Structural Information Theory

Structural information theory (SIT) was initiated by Leeuwenberg (1969, 1971). It forms a conceptual framework for research on the questions of which stimulus organizations the perceptual system can produce and which one it does produce for a given stimulus (see van der Helm, 2000, for an extensive discussion of this conceptual framework). Relevant to HA is that SIT conceives the perceptual organization process as a process that explores all candidate organizations to select the simplest one. This idea underlies SIT's coding model, which HA uses as an experimenter's tool to analyze visual stimuli, as follows.

First, a stimulus is represented by symbol strings that each form a so-called *spatially contiguous* reconstruction recipe for an interpretation of the stimulus. Below, we go into more detail on this so-called *semantic mapping*; but, note that a stimulus generally has many interpretations, each of which generally can be represented by many symbol strings. In principle, and contrary to OCW's (p. 246) allegation, all these possibilities are to be analyzed (fortunately, in practice, it generally suffices to analyze only a few).

Second, every symbol string is encoded by means of the transparent holographic ISA-rules. For every possible code, the complexity (*I*-load) is determined by counting the number of so-called structural information parameters (sips) in the code. In this article, we do not go into detail on this complexity measurement that we introduced in vdHL91 and that was discussed extensively in van der Helm, van Lier, and Leeuwenberg (1992): But, note that nearly all *I*-loads in OCW's article are incorrect. For a symbol string, the number of possible codes is combinatorially explosive. Yet, by exploiting the unique transparent and holographic character of the ISA-rules, the encoding algorithm PISA (Parameter load and ISA-rules) is able to select a guaranteed simplest code in a realistic computing time (van der Helm, 1988; vdHL91).

Finally, the stimulus interpretation with the overall simplest code is predicted to be the preferred interpretation. Like any code, this simplest code imposes a specific hierarchical organization on the stimulus, by way of a segmentation into objects and object parts. In the empirical practice, this predicted organization can be compared, for instance, with other organizations of the same stimulus or with predicted organizations of other stimuli. Currently relevant is that HA's *W*-load is a property of such overall-simplest codes.

So far so good, but to both opponents and proponents of SIT, the semantic mapping is a point of continuous concern regarding the application of SIT's coding model in the empirical practice. That is, for the time being, it is up to the experimenter to perform a psychologically plausible semantic mapping for a stimulus set at hand. This may not always be easy but is actually analogous to the fact that a computer language leaves it up to the computer programmer to come up with a good algorithm for a problem at hand. Because OCW's critique is strongly related to this issue, we elaborate on it next.

Semantic Mapping: A Brief Manual

As indicated, the semantic mapping is to yield symbol strings that each form a spatially contiguous reconstruction recipe for a stimulus interpretation. Here, we discuss 2-D interpretations only, and we go from theoretical specifications to concrete applications.

Theory. It is expedient to keep in mind that the semantic mapping does not have to reflect the way in which distal stimuli happened to have been constructed. For instance, it would be an error to think that after a change in a stimulus, the new stimulus automatically inherits the visual properties of the original stimulus. This is the type of error that invalidates Hulleman and Boselie's (1999) and Pothos and Ward's (2000) application of HA. It is also the type of error that pervades OCW's article (see the next sections for concrete examples). The point is that SIT's coding model, just as the human visual system, takes every stimulus as it is, that is, without using knowledge about the origin of stimuli.

The reconstruction demand implies that if the symbols are given concrete values, then, just as a computer code, a symbol string can be read as a sequence of instructions to render a stimulus on the basis of one of its interpretations. The question then is what are the elements that a symbol may refer to. Without further specification of the stimulus, the only sensible but not always pragmatic advice would be to represent its interpretations at the finest scale allowed by the perceptual resolution. Then, the symbols stand for parameters that specify, for example, retinal pixels and a path along these pixels.

Fortunately, it is good empirical practice to use homogeneous stimulus sets, that is, stimuli of a specific type (e.g., dot stimuli) that are random except for variations at the level of the topic of interest (e.g., regularities in the dot positions). Pragmatic advice regarding the semantic mapping can therefore be to represent interpretations at the level just below the level at which the experimentally relevant variation occurs. This advice means that the topic of interest is assumed to be tractable without having to bother about lower level details that by construction are hardly distinctive. This is analogous to writing computer algorithms in terms of higher level instructions, without having to bother about how these instructions are implemented at lower levels.

The second demand, spatial contiguity, simply implies that adjacent symbols in the symbol string correspond to adjacent elements in the interpretation. It is a direct offshoot of the hierarchical-transparency demand for coding rules in that it ensures that a hierarchical code of a symbol string corresponds to a hierarchical organization not only in the symbol string but also in the interpretation that is represented by this symbol string.

The reconstruction demand and the spatial-contiguity demand, together, imply that the symbol strings preserve characteristics of the 2-D spatial layout. The subsequent selection process then determines which of these characteristics are predicted to be perceptually relevant. Hence, both demands are part of our goodness explanation, but contrary to what OCW feared, they are not decisive. They are merely meant to ensure that interpretations get faithful symbolic representations before entering the selection process.

Applications. For dot stimuli, for instance, the semantic mapping may yield symbols that refer to a dot or to an angle or distance in a path from dot to dot (see Figure 2); unless, for example, the dots have varying shapes or colors, in which case one may have to start at a lower level to get a proper account of the stimulus variation. Similarly, for line drawings, the symbols may refer to line lengths and angles in a path through and between the line segments. In both cases, the symbols can be seen as referring to plotter instructions to reconstruct a stimulus on the basis of a 2-D interpretation of this stimulus. The spatial-contiguity demand then implies that a once-plotted line may not be crossed (because this would reflect a 3-D interpretation). In this case, the spatial-contiguity demand boils down to van Tuijl and Leeuwenberg's (1980) *object principle*, and despite OCW's denial, it is precisely what vdHL96 (p. 443) formulated in mathematical terms.

The foregoing may be illustrated by returning to the question of whether according to HA the perceptually preferred description of a symmetrical or repetitive 2-D stimulus is indeed the one that captures its point-structured symmetry property or its block-structured repetition property. OCW (see their Figure 3) challenged this by means of stimuli comparable with ⊏⊐ and ⊏⊏. They assigned symbols to the line segments (in our example, six per stimulus) and argued that the line segments can be scanned in any arbitrary order. That is, they argued that the semantic mapping may yield symbol strings (in our example, of six symbols) in which the symbols can have any arbitrary order, to the extent that the simplest codes of these symbol strings indicate that symmetrical and repetitive 2-D stimuli may both have a point structure but may also both have a block structure.

Indeed, the line segments can be scanned in any arbitrary order, but to satisfy the reconstruction demand, one must also represent the scan paths unambiguously. This means, contrary to OCW's (Footnote 2) suggestions, that angles of equal size but opposite sign cannot be represented by the same symbol. Then, as one may verify, the resulting symbol strings lead to simplest codes that impose a point structure on their symmetrical stimulus and a block structure on their repetitive stimulus. This indicates that, contrary

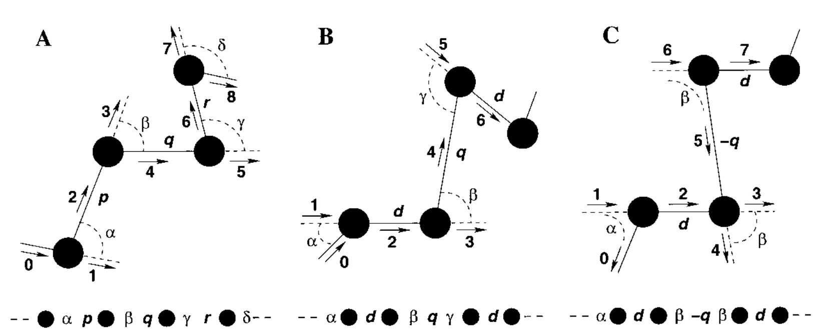

Figure 2. Examples of semantic mappings for dot stimuli. The string below each stimulus specifies the parameters of a path (the solid line) that scans the stimulus in the order given by the numerals; the arrows indicate the orientation of the scanner. A: Identical dots at random positions; between two successive dots, two path parameters are needed. B: Identical dipoles (two dots at a distance d) at random positions and in random orientations; between two successive dipoles, three path parameters are needed. C: Identical and parallel oriented dipoles at random positions; again, between two successive dipoles, three path parameters are needed, but now, two of these three parameters are identical. After the positive turn β between Points 3 and 4, the negative distance $-q$ implies, like for a car going backward, again a positive turn β between Points 5 and 6.

to OCW's fear, the holographic structuring of 1-D symbol strings does pertain to 2-D stimuli.

Goodness: From Simplicity to Weight of Evidence

Lastly, *figural goodness* is an intuitive Gestalt notion that has no formal definition. In the 1950s and 1960s, goodness got connotations such as learnability and rememberability of arbitrary stimuli. These connotations are thought to be related closely to complexity, which, at least nowadays, has a pretty clear and stable definition (see van der Helm, 2000).

More recently, goodness got the connotation of detectability or, to be more specific, the detectability of global regularities and the effect thereon of noise and local regularities. A growing body of empirical evidence showed that detectability is poorly correlated to the complexity of simplest codes. Therefore, in HA, we turned to another property of simplest codes, namely W, which in our view is pretty successful in quantifying detectability.

OCW did not seem to appreciate the foregoing distinction. For instance, OCW (see their Figure 11) compared a stimulus exhibiting a global mirror symmetry with a stimulus consisting of a number of juxtaposed local mirror symmetries. They concluded intuitively that the former has a higher goodness than the latter, and they argued that this counts against HA. However, did they think of yesteryear goodness connotations (e.g., learnability or rememberability) or of the present-day goodness connotation (i.e., detectability)? If OCW meant that a global regularity is more easily detected than a local regularity, then they were perfectly in line with HA's goodness measure W. That is, W predicts that symmetry degrades gracefully (see the next section), which, for instance, explains that in a visual search task a symmetrical target becomes less detectable as the number of asymmetrical distractors increases (as found by Olivers & van der Helm, 1998).

Thus far, this article shows that HA in fact integrates many aspects OCW seemed to cherish—even though they apparently prefer to aim at another account. Be that as it may, their theoretical misconceptions about HA persist in their empirical application of HA to the goodness issues of noise resistance, number effects, nested regularities, and symmetry effects. This is discussed in the remainder of this article.

Graceful Degradation for Symmetry but Not for Repetition

Usually, the noise resistance of a regularity is taken to be reflected by its detectability under varying noise levels. Here, however, we follow OCW, who defined the noise resistance of a regularity by the lowest noise level that renders it undetectable (*noise tolerance* would probably be a better term for this). They then argued that HA yields contradictory predictions on noise resistance because HA uncouples goodness from simplicity. As we show below, however, their argument is false.

In the previous section, we assessed that HA first measures complexity by way of the I-load to select the simplest code of a stimulus, after which HA measures goodness by way of the W-load of this simplest code. This "ingenious" move, as OCW (p. 245) called it, means that HA does not uncouple goodness from simplicity—at least, not completely. Indeed, I and W are two different properties of simplest codes, serving different purposes. OCW nevertheless pitted the I-load of a perfect regularity against its W-load as independent predictors of its noise resistance, which would give contradictory predictions because perfect symmetry and perfect repetition have equal I-loads but different W-loads.

OCW's move, however, is nonsensical—at least, in HA. Whether, and if so what, regularity is still detectable when noise has been added to a perfect regularity is determined in HA by the simplest code of the resulting stimulus, that is, not by the simplest code of the perfect regularity. Hence, in HA, the I-load and the W-load of a perfect regularity simply cannot be predictors of noise resistance, let alone contradictory predictors.

It seems that OCW did not realize that a stimulus, after a change, has to be analyzed anew (see also the previous section). They sketched a signal plus noise picture, in which they assumed that the simplest code of a perturbed regularity can, generally, still capture the original regularity. If so, the I-load and the W-load of the original regularity might indeed be predictors of noise resistance. However, OCW's signal plus noise picture is too simplistic. It may have correct implications for some cases of perturbed symmetry but generally not for perturbed repetition. This may be qualified further, as follows.

In HA's ideal-observer world, the point structure of symmetry implies that it remains detectable no matter the amount of truly random noise (see vdHL96); that is, symmetry degrades gracefully (as found by Barlow & Reeves, 1979). In the empirical practice, however, one also has to reckon with the following two nonideal restrictions. First, truly random noise does not exist. The simplest code of a perturbed symmetry may therefore be a code that captures spurious structures instead of the original symmetry. Second, human observers have limited resources. A poorly structured stimulus may therefore trigger the system to say "random" before an exhaustive analysis could be completed. It may also trigger an attention shift from high-spatial-frequency levels where the symmetry was to be found to low-spatial-frequency levels that put forward the envelope of the stimulus.

Repetition is far less noise resistant than symmetry because, apart from the aforementioned nonideal restrictions, its block structure also comes into play. It is true that depending on the location of the noise, SIT's A-rules may capture residual repetitiveness in a perturbed repetition. (The A-rules capture configurations of identical parts alternated with random parts; see Table 1.) Contrary to OCW's (p. 249) allegation, however, vdHL96 (p. 449) argued that in general repetition is very sensitive to noise. A fine example is the slightly perturbed twofold repetition in the symbol string *abcdefabcedf.* OCW (p. 249) suggested that its simplest code is <(*abc*)>/<<(*de*)(*ed*)>/<(*f*)>>, but this code is syntactically incomprehensible (it cannot be decoded into *abcdefabcedf*). The actual simplest code is *abc*S[(*d*)(*e*),(*fabc*)]*f,* which no longer describes a form of repetition.

Our analysis is corroborated by OCW's own Experiment 1 in which degraded fourfold repetitions were judged as being worse than perfect fourfold repetitions (see also Tyler & Chang, 1977). OCW's argument that their data count against HA is based on the symbol strings *abcdabcdabcdabcd, abcdabcxabcdabcy,* and *abcdaxyzabcdapqr*. The first string exhibits a perfect fourfold repetition and has a $4 * (abcd)$ as its simplest code with $I = 5$ sip. The other two strings exhibit perturbed fourfold repetitions. According to OCW, their simplest codes are <(*abc*)>/<<(*d*)>/<(*x*)(*y*)>> with $I = 10$ sip and <(*a*)>/<<(*bcd*)>/<(*xyz*)(*pqr*)>> with $I = 16$ sip, respectively (although syntactically wrong, both codes are comprehensible). OCW's point was that these codes describe perturbed repetitions with a higher W-load than perfect repetitions. However, the actual simplest codes are S[S[((*abc*)),((*d*))],(*x*)]*y* with $I = 7$ sip and S[S[((*a*)),((*bcd*))],(*xyz*)]*pqr* with $I = 12$ sip, respectively, which no longer describe forms of repetition. Hence, insofar as OCW claimed that perceptually repetitions are easily destroyed, they were in fact perfectly in line with HA. Furthermore, the fair correlation of OCW's data with I instead of W is not surprising. OCW's experimental task did not concern detectability but asked for subjective judgments of "regularity," "goodness," "complexity," and "pleasantness," which suits the yesteryear connotations of goodness (see the previous section).

Number Effect in Repetition but Not in Symmetry

As indicated earlier, HA's computational goodness model $W = E/n$ predicts quantitatively that varying the number n of stimulus elements affects repetition but not symmetry. Furthermore, HA's algorithmic holographic-bootstrap model predicts qualitatively the same number effect: Detection advances linearly in repetition but exponentially in symmetry (which agrees with Baylis & Driver's, 1994, serial–parallel distinction, although their type of repetition differs from ours). OCW's Experiments 2 and 3 yielded data that, as they argued, count heavily against HA's predictions on number effects. In this section, however, we show that their data in fact agree well with HA.

Goodness Effects, Although Numbers Stay Constant?

In Experiment 2, OCW did not test HA's prediction for varying n, other things being equal. Instead, they tested a complementary prediction; that is, they tested whether goodness remains constant under constant n, other things not being equal. They constructed twofold repetitions and mirror symmetries by placing 32 dots in a grid that could be wide (16×4), square (8×8), or narrow (4×16). The experimental task then was to discriminate such regular patterns from random patterns. OCW's data showed a main effect of grid width (better discrimination for narrower grids) that at first glance seems to disagree with the fact that HA predicts no effect for constant n.

First, however, note that OCW's (see their Figure 8) data for symmetry showed, in line with HA, hardly any effect of grid width. Second, the effect of grid width for repetition is due to OCW's discrimination task (i.e., perfect vs. random), which allowed subjects to rely on local repetitions. For instance, consider the wide (6×1) pattern *abc abc* and the narrow (2×3) pattern

$$\begin{matrix} a & a \\ b & b \\ c & c \end{matrix}$$

in which the symbols represent something figural. In both cases, the simplest code is $2 * (abc)$, but given the spatial-contiguity demand, local repetitions like $2 * (a)$ are present in the latter case only. Such local repetitions have to be detected as part of the search for the globally simplest code. For narrower grids, local repetitions as such involve less elements and are therefore, in line with HA, better detectable. They would not suffice for a perfect versus perturbed discrimination, but they suffice for the perfect versus random discrimination in OCW's experiment. Hence, OCW were correct when they proposed that for their data the crucial explanatory factor is the number of elements in between elements that are to be matched. Apparently, however, they did not realize that they in fact exposed a specific empirical appearance of the HA-predicted number effect.

Number Effect in Perturbed Symmetry?

In Experiment 3, OCW did vary the number n of pattern elements, albeit only in symmetry. They constructed starlike stimuli in which the number of contour segments could be 6, 12, 24, or 48 (our Figure 3 gives a gist). They further distinguished two conditions: discrimination between perfect symmetries and random patterns (the SvR condition) and discrimination between perfect and perturbed symmetries (the SvP condition). In both cases, they reported number effects that, as they argued, HA predicts to be absent.

However, for the SvR condition, OCW failed not only to report that they replicated work by Baylis and Driver (1994) and Tapiovaara (1990) but also to make the distinction between $n < \sim 20$ and $n > \sim 20$ that vdHL96 (p. 445) made when discussing the data of Baylis and Driver and Tapiovaara. For $n > \sim 20$, the latter studies as well as OCW's study showed hardly any number effect. For $n < \sim 20$, Baylis and Driver found that symmetry gets better for decreasing n, which is also what OCW found using a similar stimulus type, whereas for another stimulus type, Tapiovaara found that symmetry gets worse. In other words, as vdHL96 already discussed, the data are inconclusive for $n < \sim 20$, and agree with HA for $n > \sim 20$.

Furthermore, for the SvP condition, OCW claimed to have constructed perturbed symmetries that all have an HA-predicted goodness of $W = .33$, no matter the number n of contour segments. If so, HA would predict no effect of n in the discrimination from

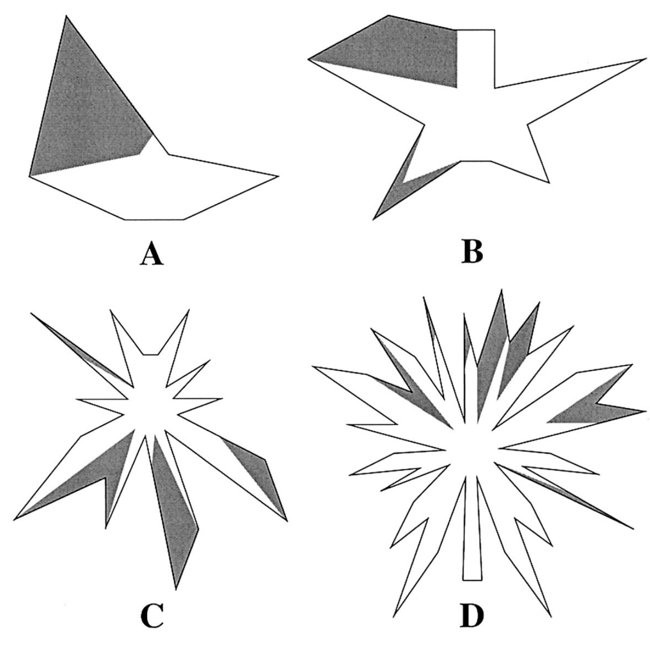

Figure 3. Sample from Olivers et al.'s (2004) Experiment 3 stimuli, with 6 (A), 12 (B), 24 (C), and 48 (D) contour segments. The white areas form the remaining symmetry, and the shaded areas form the noise (in the actual stimuli, all areas were of the same color). Unlike Olivers et al.'s suggestion, the noise-to-symmetry ratio in D is much smaller than in A, B, and C.

perfect symmetries, whereas OCW did report such a number effect (worse discrimination for larger n). However, OCW's (see their Figure 10) error scores (from which one may compute d' as a measure of discriminability) show that the number effect is completely due to the data for $n = 48$. That is, for $n = 6$, 12, and 24, OCW found hardly any number effect. Furthermore, for OCW's stimuli, the outlier $n = 48$ is an HA-predicted outlier, as we show next.

OCW were unclear about what their perturbation manipulation entailed, and they did not present examples of its result. On request, however, we received more information about the manipulation and a sample of 12 of their perturbed symmetries (3 for each value of n; see our Figure 3). OCW's perfect starlike symmetries can be said to consist of n inward or outward pointing spikes, and the perturbation manipulation implied that $n/6$ spikes were "pulled" inward or outward. The HA-predicted goodness of a perfect symmetry is $W = .50$, and OCW claimed that their manipulation implies a reduction of $1/6$ in W-load. However, OCW again used a too simplistic signal plus noise picture (see also the previous section). Apparently, they did not realize that a perturbed symmetry is not to be analyzed as if it were the original symmetry plus some perturbation. For instance, even knowing their perturbation manipulation, there is no way of telling for sure what the original symmetries were in our Figure 3. Hence, perturbed symmetries have to be analyzed anew, which we did for the sample we received from OCW, in the following two ways.

First, we applied a semantic mapping involving a contour scan. This is what OCW probably had in mind, but they did not seem to realize that apart from the contour segments, the angles between them also count as pattern elements. Nor did they seem to realize that pulling one spike inward or outward may perturb up to five pattern elements (two contour segments and three angles). For each of the values $n = 6$, 12, and 24, our reanalysis results in a predicted goodness of $W \approx .20$; therefore, there is no number effect predicted. For $n = 48$, however, it results with $W \approx .25$; therefore, there is a predicted outlier. This higher W-load arises, first, because some of the angular perturbations were well below perceptual resolution and, second, because some of the perturbed spikes were adjacent or symmetrically related, which reduces the effect per spike (see Figure 3D).

Second, we applied a more precise semantic mapping involving a (retinal) pixel scan of the area inside the contour. In this case, it is insightful to rewrite our formula $W = E/n$ into the equivalent formula $W = 1/(2 + P/S)$, in which P/S forms a noise-to-signal ratio (see vdHL96, p. 450). That is, S is the number of pixels that have a symmetrical counterpart (50% of the white areas in Figure 3), and P is the number of noise pixels (the shaded areas in Figure 3). This results, for each of the values $n = 6$, 12, and 24, in a predicted goodness of $W \approx .30$; therefore, again no number effect is predicted. For $n = 48$, however, it results in $W \approx .40$; therefore, again an outlier is predicted. Hence, according to both reanalyses, HA predicts for OCW's stimuli that discrimination is about equally good for $n = 6$, $n = 12$, and $n = 24$ but worse for $n = 48$. This is exactly what OCW found.

Scale Effect in Repetition but Not in Symmetry

To broaden the foregoing discussion, we consider Csathó, van der Vloed, and van der Helm's (2003) study on regularity detection in coarse-scaled versus fine-scaled blob stimuli (see Figure 4, first and second columns). Using similar stimuli, Oomes (1998) and Dakin and Watt (1994) looked at symmetry only and found no scale effect in discrimination from random patterns; this is also what Csathó et al. found. For repetition, however, Csathó et al. did find a scale effect (coarse is better). Because the number of blobs decreases as the scale gets coarser, this result agrees with HA's number effect in repetition but not in symmetry. Furthermore, implementationally, this result may be understood as follows.

Detection of both repetition and symmetry seems to occur in the middle occipital gyrus (MOG; Tyler & Baseler, 1998). The MOG is fed by information from the lateral geniculate nucleus (LGN), which is believed to mediate lower spatial frequencies (i.e., coarser structures) by a smaller number of larger receptive fields (see, e.g., Palmer, 1999). Logically, the MOG may exploit this LGN property

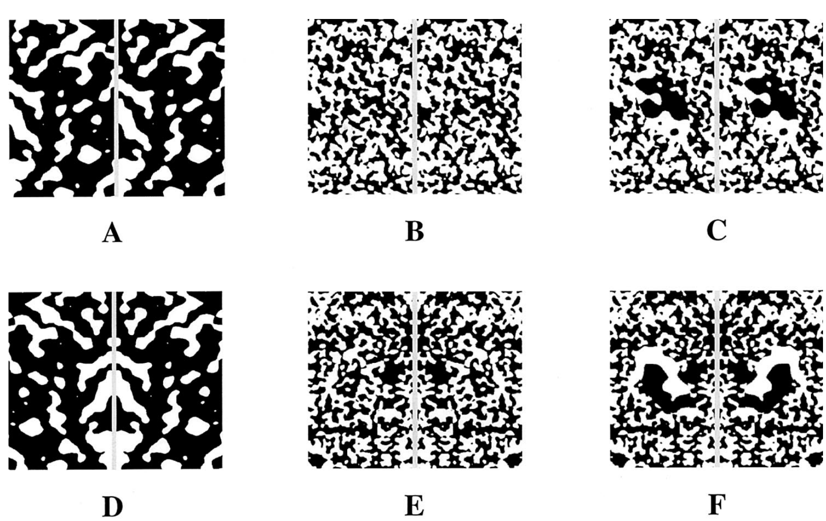

Figure 4. Sample from Csathó et al.'s (2003) stimuli. As predicted by the holographic approach, detection of repetition (top row) in A and C was found to be better than in B, and detection of symmetry (bottom row) was found to be equally good in D and E but worse in F.

for coarser-scaled repetition but not for coarser-scaled symmetry because, as vdHL99 argued, the symmetry of blob shapes can only be assessed at a fine scale. This agrees with Dakin and Watt's (1994) empirical finding that symmetry detection matches the performance of a fairly fine-scale filter.

Hence, HA's computational effect of the number of blobs in repetition but not in symmetry seems to correspond, implementationally, to a variable number of MOG fields involved in repetition detection but a constant number of MOG fields involved in symmetry detection. A second implementationally relevant implication of HA is discussed next.

Salient Blobs Help Repetition but Hinder Symmetry

As we argued vdHL96, HA predicts that extra local regularity facilitates the detection of a global regularity, and more so for a global repetition than for a global symmetry, which agrees with Corballis and Roldan's (1974) data. First, OCW objected to our computational analysis in a way that is rather incomprehensible to us: As we show in the next section, our analysis agrees with data from OCW's own Experiment 4. Second, OCW seemed to approve of Wagemans's (1999) implementational proposal that local regularity helps because it may form salient substructures that are picked up rapidly at low-spatial-frequency levels via the fast magnocellular pathway. For symmetry, however, this proposal is supported by neither empirics nor logic (see vdHL99 and the previous section). In fact, next, we discuss evidence not only against this proposal but also in favor of HA.

If salience is the reason that extra local regularity helps, then salient substructures that do not exhibit extra local regularity should help too. Csathó et al. (2003) tested this by means of stimuli as depicted in the third column of Figure 4. These novel stimuli exhibit, just like those in the second column of Figure 4, only one global regularity, except that they contain salient blob areas. Csathó et al.'s experiment involved discrimination from imperfect regularities in which either the blob areas were regular and the rest random, or vice versa. The correct scores for both perturbation manipulations showed, for imperfect repetition as well as imperfect symmetry, that the blob areas become more decisive as their coarseness increases, which confirms their salience. Furthermore, Csathó et al. found indeed a positive scale effect for perfect repetition (i.e., better for coarser blob areas). For perfect symmetry, however, they found an inverse scale effect (i.e., worse for coarser blob areas)—this goes against the proposal by Wagemans (1999) that OCW approved.

At the implementational level, Csathó et al.'s (2003) finding could be described as the result of lateral inhibition between spatial frequency levels (cf. Hughes, Nozawa, & Kitterle, 1996). That is, in repetition, the coarse-scaled blob areas seem to give a strong (low-spatial-frequency) signal that dominates the much weaker (high-spatial-frequency) signal from the rest of the pattern. In symmetry, however, the two signals are probably about equally strong (this can be concluded from Tyler & Hardage, 1996, and Dakin & Herbert, 1998), which apparently triggers a competition that hinders the detection process. In any case, Csathó et al.'s finding for symmetry negates the idea (by, e.g., Oomes, 1998) of an easy integration of information across spatial frequency levels (cf. Dakin & Watt, 1994). This implementational description, however, is in our view still rather scanty, and we feel that Csathó et al.'s finding needs further explaining. As we discuss next, a complementary explanation is provided by HA, from which Csathó et al. actually started.

As mentioned, vdHL99 translated HA's computational number effect in repetition but not in symmetry into a speed difference within HA's algorithmic holographic-bootstrap model: Detection advances linearly in repetition but exponentially in symmetry. Now, imagine a two-stage process in which at the first stage a relatively small part of a global regularity is processed and at the second stage the rest of the regularity. As vdHL99 showed, the speed difference implies that such a temporally split process hinders symmetry but not repetition. For details, see vdHL99, but by way of analogy, one may think of a slow car (repetition) for which it does not matter whether the road is busy versus a fast car (symmetry) for which it does matter. To be clear, HA does not predict when such a split process occurs, and vdHL99 predicted only what happens if it occurs.

In Csathó et al.'s (2003) stimuli, it is plausible that the coarse blob areas are processed predominantly by the fast magnocellular pathway and the rest of the pattern by the slower parvocellular pathway, thus establishing a temporally split process in the MOG. Hence, for their repetition stimuli, the HA-predicted presence of a number effect implies a positive effect of the smaller number of elements inside the blob areas but no effect of the split process as such. For their symmetry stimuli, however, the HA-predicted absence of a number effect implies a negative effect of the split process. Csathó et al.'s data confirm these implications of the HA-predicted number effects.

Hence, these last two subsections show, contrary to OCW's (p. 245) assessment, that HA's computational and algorithmic descriptions do relate to factors such as attention and spatial frequency and, thereby, to functional and neural mechanisms in vision.

Glass Patterns Can Destroy or Boost Symmetry

OCW argued that the data from their Experiment 4 contradict HA's predictions on the goodness of combinations of Glass patterns and bilateral symmetry. Here, however, we show first that their alleged HA predictions are far off the actual HA predictions and second that their data are, in part, undecisive because of empirical flaws but, for the rest, straightforwardly in support of the actual HA predictions.

OCW presented subjects with stimuli that exhibited various combinations of translational Glass patterns and bilateral symmetry (see Figure 5), in four different orientations. For all stimuli, but with blocked symmetry-detection and Glass-detection tasks, subjects had the speeded task to respond as soon as they had detected the relevant regularity (yielding reaction time data), immediately followed by the unspeeded task to specify the orientation (yielding correct scores). A "no orientation" response would be correct for Figure 5A in the symmetry-detection block and for Figure 5B in the Glass-detection block. The correct scores were highly correlated to the reaction time data and were not analyzed further by OCW.

OCW were unclear about exactly what regularity the subjects were supposed to detect. This is problematic not only for Figure 5F, which does not exhibit a global Glass pattern (as OCW acknowledged), but also for Figure 5E, which does not exhibit a mirror symmetry. Furthermore, OCW's baseline mirror symmetries were composed of dipoles (see Figure 5B). This undermines their intended comparisons because in HA dipoles are the perceptual substructures of Glass patterns but not of mirror symmetry (see Figures 1 and 5 and Table 2). Finally, in the absence of random or perturbed stimuli, OCW's dual task allowed subjects in the Glass-detection block to respond on the basis of an arbitrary pair of dipoles, whereas in the symmetry-detection block, they could respond only once the symmetry axis had emerged. This response bias does not justify OCW's (p. 256) conclusion that

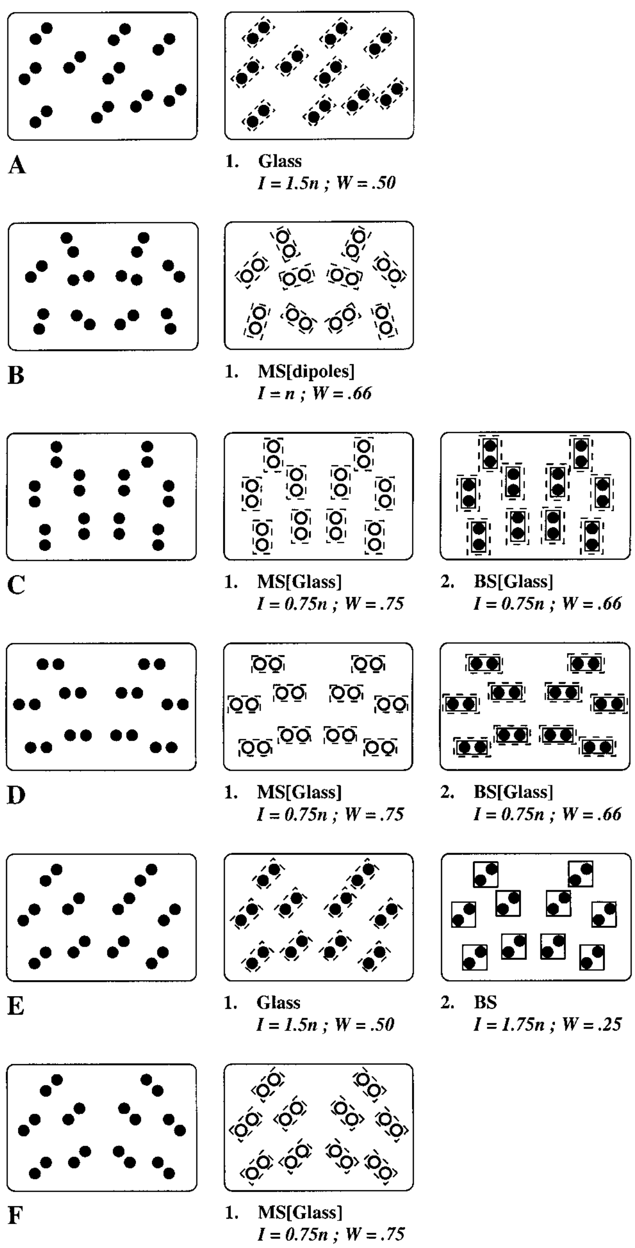

Figure 5. The left column shows the types of stimuli (consisting of n dots) in Experiment 4 by Olivers et al. (2004). The middle column visualizes simplest interpretations with complexity I and predicted goodness W. The right column visualizes alternative interpretations. Open dots, solid rectangles, and dashed rectangles indicate substructures of mirror symmetry (MS), broken symmetry (BS), and Glass patterns, respectively. See the text for further explanation of A–F.

Glass patterns are processed at an earlier perceptual stage than symmetries. It also implies that only their symmetry-detection data are useful. Diametrical to OCW's assessment, these data in fact support HA's analysis. This is discussed next.

Originally, Glass (1969) constructed his patterns by superimposing two random dot stimuli, one being a slightly translated, rotated, or dilated copy of the other. This results, proximally, in a stimulus consisting of many randomly positioned but coherently oriented dipoles. The distance between the two dots in a dipole (i.e., the dipole length) and the size difference between these two dots vary according to the transformation applied. Virtually all later studies we know on Glass patterns (including OCW), however, started from stimuli composed of identical dipoles with identical dots (see Figure 5A). For such an otherwise-random Glass pattern, SIT predicts a dipole interpretation with an HA-predicted goodness of $W = .50$ (see Figure 5A-1 and Table 2-4).

OCW assessed correctly that $W = .50$ also holds for an otherwise-random mirror symmetry (see Table 2-2). However, they incorrectly claimed that it also holds for their baseline symmetries: These were composed of identical dipoles (see Figure 5B) that form an extra regularity that enables a simpler interpretation with, for the symmetry, a higher HA-predicted goodness of $W = .66$ (see Figure 5B-1 and Table 2-5). Unfortunately, because of the earlier mentioned flaws, OCW's data are undecisive about this prediction.

Furthermore, OCW argued that SIT can capture (in one code) the combination of a Glass pattern and a bilateral symmetry for Figure 5E but not for Figure 5F. In fact, however, it is precisely the other way around, as we assess next.

First, the bilateral symmetry in Figure 5E is not a mirror symmetry, but it is a broken symmetry (see Figure 5E-2). The difference may be illustrated as follows in terms of symbol strings. The string *abcdeedcba* contains a mirror symmetry described by the code $S[(a)(b)(c)(d)(e)]$, with $I = 5$ sip and $W = .50$. The string *abcdedecab*, however, contains a broken symmetry described by the code $S[(ab)(c)(de)]$, with $I = 7$ sip and $W = .30$. This illustrates that a broken symmetry does not have a point structure and that it generally has a higher complexity and a lower predicted goodness than a mirror symmetry. Now, in the case of Figure 5E, the broken symmetry can be described but, then, the Glass pattern cannot also be described, and vice versa. The reason is that capturing the Glass pattern and capturing the broken symmetry require different descriptions of the spatial relations between the dipoles, as hinted at by the different rectangles in Figures 5E-1 and 5E-2. Furthermore, the Glass interpretation is simpler than the broken symmetry interpretation. Hence, when OCW argued that in this case Glass patterns destroy bilateral symmetry, they were in fact perfectly in line with HA. Moreover, their data confirm HA's prediction that this bilateral symmetry has a very low detectability.

Second, the mirror symmetry in Figure 5F can be captured in a code that retains one symmetry half in which, subsequently, a Glass pattern can be described (see Figure 5F-1 and Table 2-6). This yields the simplest interpretation, with a predicted high goodness of $W = .75$ for the mirror symmetry. Hence, when OCW (p. 251) argued that in this case Glass patterns boost bilateral symmetry, they were in fact again perfectly in line with HA. Moreover, their data confirm HA's prediction that this mirror symmetry is better detectable than their "dipoled" baseline symmetry with $W = .66$.

Table 2
Codes for Stimuli Consisting of n *Identical Nonoverlapping Dots*

Interpretation type and code	I	W
1. Dots at random positions $\alpha_1 p_1 \bullet \ldots \alpha_n p_n \bullet$	$3n$	.00
2. Mirror symmetry with Type 1 halves $S[(\alpha_1)(p_1)(\bullet)\ldots(\alpha_{n/2})(p_{n/2})(\bullet)]$	$\pm\ 3n/2$	.50
3. Identical dipoles at random positions and random orientations $<(\alpha_1 p_1 \beta_1)\ldots(\alpha_{n/2}\ p_{n/2}\ \beta_{n/2})>/<(S[(\bullet), (d)])>$	$\pm\ 2n$	.33
4. Glass pattern with parallel dipole orientations (see Figures 5A and 5E) $<(S[(\alpha_1), (p_1)])\ldots(S[(\alpha_{n/2}), (p_{n/2})])>/<(S[(\bullet), (d)])>$	$\pm\ 3n/2$	.50
5. Mirror symmetry with Type 3 halves (see Figure 5B) $S[<((\alpha_1)(p_1)(\beta_1))\ldots((\alpha_{n/4})(p_{n/4})(\beta_{n/4}))>/<(S[((\bullet)),((d))])>]$	$\pm\ n$	.66
6. Mirror symmetry with Type 4 halves (see Figures 5C, 5D, and 5F) $S[<(S[((\alpha_1)), ((p_1))])\ldots(S[((\alpha_{n/4})),((p_{n/4}))])>/<(S[((\bullet)),((d))])>]$	$\pm\ 3n/4$	.75
7. Broken symmetry with Type 4 halves (see Figures 5C and 5D) $S[<(S[((\alpha_1)), ((p_1))])\ldots(S[((\alpha_{n/4})),((p_{n/4}))])>/<((S[(\bullet),(d)]))>]$	$\pm\ 3n/4$	.66
8. Glass pattern with broken symmetry structure (see Figures 5C and 5D) $<S[((S[(\alpha_1), (p_1)]))\ldots((S[(\alpha_{n/4}),(p_{n/4})]))]>/<(S[(\bullet),(d)])>$	$\pm\ n$	.58

Note. For all stimuli, the semantic mapping yields strings of $\pm\ 3n$ symbols (see Figure 2). These symbols represent the so-called entry-level elements, that is, the elements that enter the perceptual organization process. In the codes, the chunks (anything in parentheses) represent the resulting exit-level elements, that is, the perceptual substructures.

Finally, for Figures 5C and 5D, OCW correctly felt that a special situation arises when the dipoles are either parallel or orthogonal to the symmetry axis. They were wrong, however, by claiming that HA cannot recognize this special situation. First, they incorrectly assessed that in HA the hierarchical relation between the symmetry and the Glass pattern can be reversed without consequence. That is, the Sym[Glass] code with a mirror symmetry is actually simpler than the Glass[Sym] code with a broken symmetry (see Table 2-6 and 2-8). Second, just as simple as the Sym[Glass] code above is another Sym[Glass] code with a broken symmetry in which each dipole forms a substructure (see Figures 5C-2 and 5D-2, and Table 2-7). Thanks to the special dipole orientation in Figures 5C and 5D, this broken symmetry can be captured by way of the same description of the spatial relations between the dipoles as the one required for capturing the Glass pattern.

Thus, in Figures 5C and 5D, SIT predicts perceptual ambiguity: Either the dots or the dipoles may form the substructures of the bilateral symmetry. This ambiguity and that both options yield a simpler code than the Glass[Sym] code agree with OCW's (p. 256) observation that at lower spatial frequency levels, the Glass dipoles blur away into one blob each, leaving the therefore dominant symmetry. The predicted goodness of the broken symmetry ($W =$ .66) is lower than that of the mirror symmetry ($W = .75$). This, plus the fact that perceptual ambiguity distracts subjects, makes it understandable that, as OCW found, performance is not better than for their baseline symmetry with $W = .66$. Hence, in summary, OCW's symmetry-detection data actually agree well with HA.

Symmetry Effects Are Context Dependent

Lastly, OCW disputed our analysis of symmetry effects as found by Freyd and Tversky (1984). Freyd and Tversky showed subjects a reference pattern (a perturbed symmetry) and two targets (a slightly more symmetrical and a slightly less symmetrical version of the reference). Subjects had to judge which target was more similar to the reference. When reference and targets had a relatively high level of symmetry, subjects tended to choose the more symmetrical target (a symmetry effect), but when reference and targets had a relatively low level of symmetry, subjects tended to choose the less symmetrical target (an asymmetry effect).

In vdHL96 (p. 450), we suggested that the overall level of symmetry may not have been the decisive factor in triggering these opposite effects. For similar triadic comparisons, we showed that HA's goodness measure $W = E/n$ predicts a symmetry effect (W_{ref} closer to W_{more}) if the targets are created by manipulating the amount of symmetry in the reference but an asymmetry effect (W_{ref} closer to W_{less}) if the targets are created by manipulating the amount of noise in the reference.

OCW disputed our analysis by alleging that we calculated W_{more} and W_{less} in different ways, depending on the condition at hand. This, however, is nonsense. We simply calculated $W = E/n$ for every pattern separately, that is, no matter whether it was a reference or a target and no matter the condition to which it belonged. It then turned out to depend on the specific triadic comparison whether W_{more} or W_{less} was closer to W_{ref}. To be clear, later analyses revealed that HA also predicts an effect of the overall level of symmetry, yet recently finished experiments confirmed that at a constant overall level of symmetry, both symmetry and asymmetry effects do occur as predicted.

Finally, Freyd and Tversky's (1984) research was triggered by observations that humans seem to perceive more symmetry than there really is. Our analysis suggests that in some contexts such a symmetry effect indeed occurs but that in other contexts even an asymmetry effect may occur. Contrary to OCW's (p. 258) allegation, this context dependency is precisely what vdHL96 (p. 452) pointed at in their ecological argument that both symmetry effects and asymmetry effects have survival value. To be clear, symmetry effects probably have survival value for the perceivers only, whereas asymmetry effects probably have survival value for the perceived only. This is not awkward: We think that the evolution presents organisms with "package deals." That is, an internal cognitive system plus an external ecological niche may have good and bad points, but the organism can survive as long as the good points are good enough to compensate for the bad points.

Summary and Conclusion

In this article, we discussed the HA to goodness (i.e., to the detectability of visual regularities) and its application to noise resistance, number effects, nested regularities, and symmetry effects. Whatever the extent of HA's explanatory power eventually may be, this article shows that HA is much stronger than OCW suggested. HA may not provide all the answers OCW wanted to know, but it provides more and better answers than OCW seemed to realize.

The main flaws in OCW's comment were, first, that they mistook the explicitly informal clarifications in vdHL96 for vdHL91's formal definitions and, second, that they showed a poor understanding of basic coding principles. For instance, they persistently violated the coding principle that a stimulus after a change has to be encoded anew because it does not automatically inherit the visual properties of the original stimulus.

Furthermore, OCW repeatedly pointed at the danger for HA of becoming a redundant theory: According to OCW, one can think of implementational factors that might be decisive without needing HA any longer. However, the risk of becoming redundant is neither frightening nor high for a theory that plays a fruitful part in the empirical cycle. For instance, the transformational approach and the original bootstrap model have played fruitful parts in the empirical cycle, and their underlying ideas continue to have relevance—even though they have to resort to implementational factors to account for even the most basic goodness effects.

In view of the above, it is strange that OCW (p. 259) nevertheless suggested that those two approaches might be better than HA. After all, without having to resort to implementational factors, HA accounts for much more goodness phenomena. Also OCW's (p. 259) alternative suggestion to take the complexity of simplest codes as measure of goodness is strange because for many goodness phenomena it implies that one has to resort to implementational factors that detract from simplicity. This option clearly is worse than HA's goodness measure *W* that accounts for much more goodness phenomena without detracting from simplicity; after all, *W* is a property of simplest codes.

Finally, HA is primarily a theory at the computational level, but this article shows that it also complies with OCW's wishes to gain deeper insights at the algorithmic and implementational levels, without, at least so far, becoming redundant itself. That is, by providing more insight in the structures to be detected, HA has provided and still continues to provide more insight in the actual detection of these structures.

References

Barlow, H. B., & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research, 19,* 783–793.

Baylis, G. C., & Driver, J. (1994). Parallel computation of symmetry but not repetition within single visual shapes. *Visual Cognition, 1,* 377–400.

Corballis, M. C., & Roldan, C. E. (1974). On the perception of symmetrical and repeated patterns. *Perception & Psychophysics, 16,* 136–142.

Csathó, A., van der Vloed, G., & van der Helm, P. A. (2003). Blobs strengthen repetition but weaken symmetry. *Vision Research, 43,* 993–1007.

Dakin, S. C., & Herbert, A. M. (1998). The spatial region of integration for visual symmetry detection. *Proceedings of the Royal Society of London, Series B, 265,* 659–664.

Dakin, S. C., & Watt, R. J. (1994). Detection of bilateral symmetry using spatial filters. *Spatial Vision, 8,* 393–413.

Freyd, J., & Tversky, B. (1984). Force of symmetry in form perception. *American Journal of Psychology, 97,* 109–126.

Glass, L. (1969). Moiré effect from random dots. *Nature, 223,* 578–580.

Hughes, H. C., Nozawa, G., & Kitterle, F. (1996). Global precedence, spatial frequency channels, and the statistics of natural images. *Journal of Cognitive Neuroscience, 8,* 197–230.

Hulleman, J., & Boselie, F. (1999). Perceived shape regularity does not depend on regularities along the contour. *Perception, 28,* 711–724.

Leeuwenberg, E. L. J. (1969). Quantitative specification of information in sequential patterns. *Psychological Review, 76,* 216–220.

Leeuwenberg, E. L. J. (1971). A perceptual coding language for visual and auditory patterns. *American Journal of Psychology, 84,* 307–349.

MacKay, D. (1969). *Information, mechanism and meaning.* Boston: MIT Press.

Maloney, R. K., Mitchison, G. J., & Barlow, H. B. (1987). Limit to the detection of Glass patterns in the presence of noise. *Journal of the Optical Society of America A, 4,* 2336–2341.

Marr, D. (1982). *Vision.* San Francisco: Freeman.

Olivers, C. L. N., Chater, N., & Watson, D. G. (2004). Holography does not account for goodness: A critique of van der Helm and Leeuwenberg (1996). *Psychological Review, 111,* 242–260.

Olivers, C. L. N., & van der Helm, P. A. (1998). Symmetry and selective attention: A dissociation between effortless perception and serial search. *Perception & Psychophysics, 60,* 1101–1116.

Oomes, S. (1998). *Human visual perception of spatial structure.* Unpublished doctoral thesis, University of Nijmegen, Nijmegen, The Netherlands.

Palmer, S. E. (1983). The psychology of perceptual organization: A transformational approach. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 269–339). New York: Academic Press.

Palmer, S. E. (1999). *Vision science: Photons to phenomenology.* Cambridge, MA: MIT Press.

Pothos, E. M., & Ward, R. (2000). Symmetry, repetition, and figural goodness: An investigation of the weight of evidence theory. *Cognition, 75,* B65–B78.

Tapiovaara, M. (1990). Ideal observer and absolute efficiency of detecting mirror symmetry in random images. *Journal of the Optical Society of America A, 7,* 2245–2253.

Tyler, C. W., & Baseler, H. A. (1998). fMRI signals from a cortical region specific for multiple pattern symmetries. *Investigative Ophthalmology & Visual Science, 39*(Suppl.), 169.

Tyler, C. W., & Chang, J.-J. (1977). Visual echoes: The perception of repetition in quasi-random patterns. *Vision Research, 17,* 109–116.

Tyler, C. W., & Hardage, L. (1996). Mirror symmetry detection: Predominance of second-order pattern processing throughout the visual field. In C. W. Tyler (Ed.), *Human symmetry perception and its computational analysis* (pp. 157–172). Zeist, The Netherlands: VSP.

van der Helm, P. A. (1988). *Accessibility and simplicity of visual structures.* Unpublished doctoral dissertation, University of Nijmegen, Nijmegen, The Netherlands.

van der Helm, P. A. (2000). Simplicity versus likelihood in visual perception: From surprisals to precisals. *Psychological Bulletin, 126,* 770–800.

van der Helm, P. A., & Leeuwenberg, E. L. J. (1991). Accessibility, a criterion for regularity and hierarchy in visual pattern codes. *Journal of Mathematical Psychology, 35,* 151–213.

van der Helm, P. A., & Leeuwenberg, E. L. J. (1996). Goodness of visual regularities: A nontransformational approach. *Psychological Review, 103,* 429–456.

van der Helm, P. A., & Leeuwenberg, E. L. J. (1999). A better approach to goodness: Reply to Wagemans (1999). *Psychological Review, 106,* 622–630.

van der Helm, P. A., van Lier, R. J., & Leeuwenberg, E. L. J. (1992). Serial pattern complexity: Irregularity and hierarchy. *Perception, 21,* 517–544.

van Tuijl, H. F. J. M., & Leeuwenberg, E. L. J. (1980). Perceptual interpretation of complex line patterns. *Journal of Experimental Psychology: Human Perception and Performance, 6,* 197–221.

Wagemans, J. (1999). Toward a better approach to goodness: Comments on van der Helm and Leeuwenberg (1996). *Psychological Review, 106,* 610–621.

Wagemans, J., Van Gool, L., Swinnen, V., & Van Horebeek, J. (1993). Higher-order structure in regularity detection. *Vision Research, 33,* 1067–1088.

Weyl, H. (1952). *Symmetry.* Princeton, NJ: Princeton University Press.

Received April 4, 2002
Revision received April 16, 2003
Accepted June 23, 2003 ■