COGNITIVE NEUROPSYCHOLOGY, 1997, *14* (1), 131–165

Complex Dynamic Systems also Predict Dissociations, but They Do Not Reduce to Autonomous Components

Guy C. Van Orden and Marian A. Jansen op de Haar
Arizona State University, Tempe, USA

Anna M. T. Bosman
University of Nijmegen, The Netherlands

Dissociations, according to the target articles, are due to damaged autonomous phonologic (or spelling) representations. However, a damaged recurrent network model may also produce dissociations. Recurrent networks do not entail autonomous components. They are strongly nonlinear dynamic systems that self-organise through recurrent feedback. A simple model with these properties that produces both regularisation errors (PINT named to rhyme with MINT) and semantic errors (BUSH named as TREE) is described. It may also produce dissociations between "spoken" responses and "written" responses. The mathematical basis of this model is motivated by contemporary neurobiological accounts that also derive from dynamic systems theory. The mathematical basis may also predict *multistability* and *metastability*. These are indicated by *hysteresis* and *1/f noise*, respectively, and we review recent reports of these phenomena in speech perception and word recognition. In addition, feedback has been corroborated in the feedback consistency effect. Reported generic behaviours of a complex system, the simulated dissociation of errors, and the established bidirectional nature of perception all demonstrate the utility of a cognitive systems approach to cognitive phenomena.

Requests for reprints should be addressed to Guy Van Orden, Cognitive Systems Group, Department of Psychology, Arizona State University, Tempe, AZ 85287–1104, USA (email: guy.van.orden@asu. edu).

Preparation of this commentary was supported by a FIRST Award (CMS 5 R29 NS26247) and an Independent Scientist Award (1 K02 NS 01905) to Guy Van Orden from the National Institute of Neurological Disorders and Stroke. We thank Peter Killeen for a brief tutorial on *metastability.*

INTRODUCTION

This special issue of Cognitive Neuropsychology showcases four careful and detailed case studies. All four converge on a common conclusion: Autonomous spelling representations may affect behaviour associated with literacy. By inference, automotous spelling representations must be included in a proper account of intact reading and writing. The following performance profiles motivate this conclusion.

PS, a patient described in Hanley and McDonnell (this issue), produces the correct written response when presented with the picture of a BEAR, but cannot produce the correct spoken name of the picture. Generally, PS had great difficulty with tasks that required spoken responses relative to tasks that allowed written responses. Similarly, his performance was better when he could point at the correct alternative, as in matching a picture with a written word. His difficulty with spoken responses is due to a general deficit in phonology. For example, he is sometimes unable to produce the sound-alike alternative of a homophonic word (although most errors in this and other tasks are constrained by the spelling and phonology of the correct responses). A similar case is described by Shelton and Weinrich (this issue). The patient EA's picture naming was better when the task allowed a written response than when it required a spoken response. However, EA had great difficulty writing words to dictation, and was essentially incapable of writing nonwords to dictation. As in the previous case, EA performs poorly on many tasks that require a spoken response, but unlike the previous case, EA can correctly repeat spoken words.

The patients PW (Rapp, Benzing, & Caramazza, this issue) and WMA (Miceli, Benvegnù, Capasso, & Caramazza, this issue), produced overall profiles that were similar to each other, but different from the previous two patients. Both PW and WMA were better at reading words aloud than at writing words to dictation, an asymmetry that is also found in intact reading and writing. Their overall picture naming was also better when the task allowed spoken responses compared to written responses. However, like PS and EA, both patients occasionally produced a correct written response when they did not produce a correct spoken response. This occasional pattern is the key dissociation. Consider WMA, for example. When asked to name a picture, WMA may provide one name if the response is spoken and a different name if the response is written. Two different responses to the same picture! Sometimes the correct response is the written response, sometimes it is the spoken response, and sometimes neither is correct.

Each of the previous patients produces a dissociation between written responses and spoken responses. A traditional logic licenses a formulaic interpretation of dissociations. For example, writing one name and saying a different name to the same picture indicates that the separate responses have separate causal origins. The causal basis of a spoken name is a phonologic representation. By implication, if phonologic representations affect the written response, then the written response should be the same as the spoken response. Consequently, when the written name diverges from the spoken name it cannot have originated in phonology. And yet a written response is produced, which implies that some intact representation is still present, perhaps an orthographic representation. Thus, by default, we infer the presence of spelling representations that are independent of phonologic representations. This inference is then extended backward in time. We project back in time to before these men had brain damage and infer that the same orthographic representations were autonomous in an intact specialised reading (writing) process.

We were invited to describe a different perspective on the previous patients' data. Our account is not more correct than those offered in the target articles. The framework that we work within cannot be discriminated from traditional computational models of cognitive performance on the basis of correspondence to data (see Stone & Van Orden, 1993, 1994). Instead, the value of our account derives from two simple propositions: Things can look a bit different from a different perspective, and multiple perspectives on phenomena may yield more general understandings of those phenomena.

The difference in perspective is fundamental. To understand why, we must back up the traditional logic and explore its root assumptions—i.e. how it is that one infers cognitive structures from observed behaviour. Next we discuss these root assumptions of the standard practice of cognitive neuropsychology. Past that point, we briefly review contemporary hypotheses concerning the neurological basis of behaviour. Our goal in that review is to motivate an alternative neurobiological metaphor that does not entail the root assumptions of conventional cognitive neuropsychology. Following that, we describe a simple recurrent network model of intact word perception that is congruent with contemporary neurobiology, and explain how this simple network begins to account for key findings from the target articles. The simple model is strictly grounded in the theoretical basis of the neurobiological metaphor. Finally, we review several recent findings that demonstrate the utility of our approach.

THE EFFECT = STRUCTURE ASSUMPTION

Standard computational models are usually laid out in a flow chart of processes that transform one cognitive structure into another, as when a spelling representation is transformed into a phonologic representation. Two of the target articles include flow chart illustrations that track the cognitive structures of reading and writing. In this section, we look closely at the theoretical method that licenses the discovery of these cognitive structures.

Orthographic representations and other cognitive structures are induced from reliable features of human behaviour. We observe specific patterns in behaviours, but our goal is to induce general structures of a cognitive architecture. However, no guaranteed formula exists for this generalisation, because an objective God's-eye-view of cognition is not possible. Instead we rely on plausible a priori assumptions, which we trust as though they were true. In contemporary cognitive psychology, for example, we assume that careful laboratory studies can reveal the presence of cognitive structures. Observed performance in laboratory tasks (e.g. the overall variability in response times or errors) is divided into component effects using linear statistical methods (e.g. ANOVA), and these component effects originate in causal components of mind. Thus, behaviour is assumed to be the *sum* of strictly separable pieces, *plus* some noise.

For example, in a categorisation task, subjects miscategorise homophones like ROWS as a *flower*, more often than control words like ROBS. This main effect of homophone phonology is treated as a separate piece of overall behaviour. The isolated piece is thought to originate in a distinct causal structure of the cognitive architecture, namely, a representation of /roz/. The contrast between the experimental condition (ROWS) and the control condition (ROBS) reduces behaviour (categorisation errors) to pieces (effects) that, in turn, indicate the pieces of mind (causal structures) in which the behaviour originates. In other words, *the presence of an effect equals the presence of a structure*.

In cognitive psychology, this logic usually ends as we left it in the previous paragraph—a positive demonstration of a reliable effect ends in the inference of a cognitive structure. However, it is just as important for this logic that the opposite side of this inference is reliable, namely, *the absence of an effect equals the absence of a structure* (Mackie, 1974; Mill, 1974). This logical entailment is not always made explicit by cognitive psychologists; it is more prominent in neuropsychology. Cognitive neuropsychologists must infer the nature of cognition prior to brain damage from the shape of behaviour after brain damage. They take careful note of the missing pieces of behaviour and use these observations to reconstruct the previously intact system. This is the basis for the standard dissociation logic of cognitive neuropsychology.

When the patient PW fails to name a picture of a PEAR correctly (Rapp et al., this issue), his error deviates from the response that would be expected from intact naming. Apparently, a causal structure that would be present in intact naming is absent in PW. The standard dissociation logic licenses the inference that a particular causal structure is missing and that this causal structure would have had isolable effects in the intact architecture. The failure to produce the spoken word *pear* thus indicates the absence of output phonology or some other linking structure in a causal chain with output phonology. All discussion in the target articles concerns precisely which structures are absent from a patient's behaviour. Cognitive neuropsychologists must always infer which or what kind of cognitive structures were present, prior to a lesion, from behaviour's missing pieces after the lesion (Patterson, 1981). Thus, they require a reliable basis to infer that *the absence of an effect equals the absence of a structure*. We refer to the *presence* and *absence* sides of this inference together as *effect = structure* (Van Orden, Holdend, Podgornik, & Aitchison, submitted; cf. Lakoff, 1987).

Assuming that *effect = structure*, the case studies described in the target articles supply compelling evidence that the complex writing and naming behaviour of these patients may be reduced to classes of behaviour or functions (e.g. writing versus naming), and that these functions may be reduced further to more elementary causal structures (e.g. orthographic and phonologic representations). Thus behaviour originates in isolable cognitive structures or *single causes*). Single causes entail the familiar notion of "domino causality." Push the first domino in a chain of standing dominos and each will fall in its turn. The input to this causal chain, a shove on the first domino, is linked to the output, the force of the first domino as it falls. In turn, this output becomes the input to the second domino, and so on, for each trailing domino down the chain. It is this notion of causality that is assumed in flow charts of cognition. A stimulus input is linked by causal rules through a hierarchical chain of representations, and the final output of this causal chain is the observed datum in a laboratory task.

Now comes the tricky part. Given that we may only observe the final output datum, how do we get inside this causal chain to discover its components? The solution is to choose tasks and manipulations that differ from each other by the causal equivalent of one domino (or one branch off a forked chain of dominos). The extensive test batteries in the target articles are designed for exactly this purpose. The problem that arises, however, is how to decide which tasks differ by exactly one single cause, or by one branch of single causes. One requires objective knowledge of how tasks are accomplished to know reliably which or how many components each task entails. Thus, we face an inescapable problem of circularity. Our goal is to induce general cognitive components entailed in a specific task from observed behaviour, but the method by which we induce these components requires reliable a priori knowledge of the self-same cognitive components.

No theoretical approach escapes this problem. Every act of induction derives from a set of a priori assumptions, and no act of induction can validate the assumptions from which it derives (cf. Duhem, 1954; Quine, 1961). No matter how compelling the apparent confirmation of single causes, we must resist accepting this confirmation as conclusive. Objective knowledge of single causes is required, a priori, to induce single causes reliably from observed behaviour. Consequently, the successful induction of single causes cannot be turned around to validate the root assumption. The simple danger here is that scientists may be seduced by their own success. The danger of this seduction is complacency, by which we mean the exclusion of other possibilities. The history of science is littered with previously successful paradigms that, in their time, were practised to the exclusion of other possibilities. With respect to the target articles, the conclusion that the observed pattern of dissociations demonstrates autonomous or independent representations (single causes) simply affirms the inevitable consequent of assuming there were autonomous representations in the first place (Shallice, 1988; Van Orden, Pennington, & Stone, submitted).

To this point, we have not questioned the utility of continuing in the standard practice of cognitive neuropsychology. We merely suggest that it is wise to maintain a sceptical stance with respect to a priori assumptions such as *effect = structure*.

NEUROBIOLOGY AND BEHAVIOUR

So why are we so concerned with a version of causality that, on the surface, seems so plausible? After all, causality, at our natural scale of experience, seems to agree with this intuition—dominos do knock each other down. Although intuitive causality may serve us in ordinary circumstances, it may nevertheless distort our view of cognitive systems. By comparison, many other areas of contemporary science have been reframed by taking an alternative perspective that does not simply entail single causes (e.g. see Cohen & Stuart, 1994; Freeman, 1995; Goodwin, 1994). This alternative also invokes a reciprocal form of causality in which every part of a system is always present in each behaviour of that system. Each of these parts continuously affects every other part, to the point that their independent contributions cannot be sorted out in the behaviour of the whole. Most important, this perspective has usefully been applied to neuroscience, which is pertinent to our present concerns. The educated guesses that we make concerning brain damage are informed by our knowledge of how brains work in the first place. We will illustrate this notion of reciprocal causality by describing contemporary hypotheses concerning the neurobiological basis of behaviour.

At one time, Hubel and Wiesel's classic experiments seemed to provide a reasonable basis in neurobiology for single causes. They demonstrated reliable correlations between stimulus events and individual neuronal activity (Hubel & Wiesel, 1962, 1965, 1968). Flow-chart theories often extrapolate from Hubel and Wiesel's findings. "Feature detectors" are extended metaphorically to cognitive systems where they become hierarchies of representations. In perception, hierarchies of stimulus features, and combinations of features, culminate in explicit representations of whole stimulus forms ("grandmother cells"). This qualitatively linear scheme assumes a causal chain between real-world objects, their stimulus forms, and their consequent representations. Access to representations may depend on weakly nonlinear mechanisms such as thresholds, but they are laid out in a linear chain of single causes from proximal stimulus to intermediate representations to observed behaviour.

The previous causal chain was composed of isolable representations that correspond to behaviourally meaningful information. In its simplest form, this metaphor implies representations with binary states analogous to neural detectors that are either above or below their thresholds (alternatively, a continuous change in neural response amplitude eventually yields a discontinuous change in representation). Recent findings, however, suggest that neural activity undergoes a more complex qualitative change than crossing a threshold. This qualitative change entails the reciprocal causality that we mentioned earlier.

Presently, behaviour is thought to originate in neural activity that self-organises through recurrent feedback into interdependent, context-sensitive, dynamic patterns. Observable behaviour derives from complex coordinated activity among populations of *sensory* and *motor* neurons (Bressler, Coppola, & Nakamura, 1993; Freeman, 1991a, 1991b, 1995; Singer, 1993; Skarda & Freeman, 1987; von der Malsburg & Schneider, 1986). Although, initially, a stimulus pattern may activate specific neural ensembles (Livingstone & Hubel, 1988), this local activation is subsequently transformed into a global pattern as the system self-organises in recurrent feedback dynamics (Freeman, 1995; Skarda & Freeman, 1987). Input conditions of local activation become complex oscillating patterns in which the character of "input" activity is a strongly nonlinear function of "output" activity. These complex patterns are not strictly tied to the local tissue that serves as their excitable medium. Global patterns of neural activation are related to neurons in (very) roughly the same way that global patterns of ocean waves are related to water molecules. (See Goodwin, 1994, for more precise but less mundane analogies.) Nevertheless, each global pattern maintains an identifiable but highly context-sensitive profile in the amplitude "waves" of neural activity.

We may find utility in the contemporary neurobiological metaphor for understanding cognitive performance. However, we must accommodate the *nonlinear* qualitative transformation of *sensory* and *motor* activation into subsequent, behaviourally meaningful, *sensorimotor* dynamics. By implication, the performance of cognitive systems is not reduced to local chains of independent feature hierarchies; behaviour emerges in coordinated activity among interdependent *sensorimotor ensembles*. This strongly nonlinear, qualitative transformation leaves no practical possibility of reducing cognitive performance to singularly causal neural ensembles, nor even singularly causal component oscillations. Within this metaphor, sensory and motor activation are combined and transformed as their dynamic trajectory traces a path of *metastable states*.

Metastability implies that a nervous system never settles fully on a dominant percept or action (attractor). In a sense, this means that the system is always entertaining alternative possibilities, albeit only one alternative may be expressed in perception and action. Consequently this system is always slightly unstable. This instability makes it more flexible. By (very) rough analogy, a person who *always* has an alternative "plan B," or plans B, C, and D, at *every* juncture, would appear more flexible (metastable) than a single-minded person who cannot deviate from plan A. Crucially, however, if metastable neural states are the basis of observable behaviour, then no basis exists in observable behaviour for inducing causal chains of representations. The metastable basis of behaviour is antithetical to discovering implacable single causes that run from input to output.

In such strongly nonlinear complex systems, it is impossible to track backwards from observable, contextually embedded, stimulus-response, attractor states to initial conditions. Consequently, it is not possible to discover a causally distinct "input" component. Performance with few degrees of freedom emerges in patterns of sensorimotor activity with moderate degrees of freedom, which originate in prior patterns of sensory and motor activation with vast degrees of freedom. Each reduction in the degrees of freedom is a loss of information about previous states of the system. The loss of information creates an impenetrable *barrier of uncertainty* (Abraham & Shaw, 1992; Prigogine & Stengers, 1984). This barrier blocks any possibility of discovering linear causal chains running between proximal stimuli and performance (Uttal, 1990). Rather, "stimulus" and "response" are better viewed as an irreducible whole that self-organises through recurrent feedback (Freeman, 1995; Van Orden & Goldinger, 1994, 1996; Varela, Thompson, & Rosch, 1991). This self-organising gestalt allows fluid continuity between action and perception, and organism and environment (Gibson, 1986; Turvey & Carello, 1981). Coincidentally, the striking competence of behaviour is fluid accommodation of continuous changes in the environment (Gibbs & Van Orden, submitted; Stone & Van Orden, 1993; Thelen & Smith, 1994; Van Orden, Holden, et al., submitted).

Contemporary models of nervous systems are specific applications of mathematical dynamic systems theory, the most general formal framework available to scientists. Dynamic systems theory concerns how behaviour of complex systems changes over time. This mathematical theory promises someday to provide a common metalanguage for diverse areas of science (Abraham & Shaw, 1992; Haken, 1984). Today, it provides new theoretical and methodological tools for scientific enquiry. This framework is so general as to include previous flow-chart models as a subset of possibilities. Giunti (1995) describes how computational models (Turing machines) form a narrow subset of dynamic systems. Our discussion requires a related point: The products of linear reductive analyses (reliable intercorrelations between independent and dependent variables) may always be redescribed ("regraphed") to coincide with the trajectory of a dynamic system. Consequently, behaviour thought to be characteristic of linear systems may always be reframed in terms of more general nonlinear systems. The simple point here is that dynamic systems theory is sufficient to account for any behavioural phenomenon that was previously formalised in a flow-chart model. What remains to be determined is whether it will have utility over and above the traditional approach (Carello, Turvey, & Lukatela, 1992).

In the next section, we describe a model system in which behaviour is organised by recurrent feedback. Our model system is not reducible to a neural account, although it is strictly in line with the contemporary metaphor. As Thelen and Smith (1994, p. 130) note: "While mechanisms of change for mental processes most certainly do involve changes in neurotransmission, satisfactory explanations need not reside only at this level. Nonetheless . . . explanation at every level must be consistent and ultimately reconcilable . . . the dynamics of behavioural phenomena must be consistent with the dynamics of the neural phenomena."

We propose that many behavioural and neural phenomena may be reconciled with respect to a common mathematical basis, as found in dynamic systems theory. Respective choices for models within this framework, however, are mostly determined by the observed complexity of the behaviour to be modelled. For example, we model performance in simple reading tasks. Simple reading tasks provide punctate data measured at a single point in time (response time) and scored for accuracy. Thus, we may begin with models appropriate for "fixed-point" data—i.e. models that converge on point attractors. The network that we describe is not superficially isomorphic to a neurobiological model (e.g. Freeman, 1987). At most, it traces a low-dimensional "shadow" of the vastly higher dimensional description of nervous systems.

Our goal is to illustrate how the behavioural trajectory of a fully interdependent dynamic system may simulate the behaviour that is characteristic of dissociations. In the service of this goal (Freeman, 1995, p. 53): " . . . we move conceptually from the local neural network and its clearly defined properties out to the limits of its utility, multiply the network to infinity, and then awaken into a new local network in which the infinities of components are collapsed into the emergent elements at the next higher hierarchical level." Our simple model mimics intact behaviour and the behaviour characteristic of dissociations. Eventually, we explain how this simple model could fail to name a "picture" of a pear and yet produce a "written" response, or produce different "spoken" and written responses to the same picture.

PHONOLOGIC MEDIATION IN READING AND SPELLING

The emphasis in much reading research is on the perception of single words, as this is the most predictive aspect of reading skill. A poor aptitude for word perception severely limits the development of skilled reading and reading comprehension (Pennington, 1991; Perfetti, 1985). As noted in the target articles, a perennial question in such research is whether a word's phonology always mediates visual word perception. For reading, a recent proliferation of phonology effects in laboratory paradigms suggests that phonology's role is fundamental (for overviews, see Berent & Perfetti, 1995; Carello et al., 1992; Frost, submitted; Katz & Frost, 1992; Perfetti, Zhang, & Berent, 1992; Van Orden et al., 1990). For writing, systematic misspellings, like substituting ROZE or ROWS for ROSE, are common (for an overview, see Bosman & Van Orden, in press). Phonologic constraints are also apparent in patients' spelling performance, as when a patient spells YACHT as YOT (Hatfield & Patterson, 1983). Moreover, patients who exhibit substantially disrupted phonology almost always also exhibit bizarre reading performance.

Phonology effects are found with readers and writers spanning the full range of reading skill. They are found across languages (in both alphabetic and nonalphabetic writing systems), and across laboratory tasks. All these phenomena converge on the straightforward conclusion that phonology is fundamental to intact reading and spelling. Why is phonology so involved in reading or spelling? In the next section, we describe a recurrent network model of word perception and spelling that explains phonology's fundamental role. The description pertains to a very simple model that has been implemented (Farrar & Van Orden, 1994), but the principled basis of our account is not tied to the specifics of our simulation (see Stone & Van Orden, 1994; Van Orden & Goldinger, 1994).

We do not propose a connectionist network that is causally in between a stimulus and a response. The nodes of our model cannot be localised within a mind or brain separate from its environment (see also Saltzman, 1995). At first this may sound paradoxical, but recurrent dynamics are plausibly described as occurring directly, between an organism and its environment (Turvey & Carello, 1981; Van Orden & Goldinger, 1994). Perhaps this theoretical entailment will appear less paradoxical if we take into account the pragmatic entailments of behavioural research. Behavioural data cannot apportion effects between environments and organisms because behaviour always occurs at their interface. Consequently, with respect to behavioural data, there is no reliable way to determine where the environment leaves off and the organism begins (Shanon, 1993; Varela et al., 1991). Organism and environment are always both "present" in every instance of behaviour. Nevertheless, to make our story more concrete we tell it with respect to a fictitious nervous system.

A SIMPLE MODEL

Imagine a fictitious nervous system that perceives printed words. This system consists of three families of neurons: letter neurons, phoneme neurons, and semantic neurons. Every neuron in each family is (potentially) bidirectionally connected to every neuron of the other two families. Bidirectional connectivity means that if a feedforward connection exists from neuron "x" to neuron "y," there is also a feedback connection from neuron "y" to neuron "x." Now, imagine a specific pattern of activation across the letter neurons, due to the presence of a printed word. This letter pattern feeds activation forward through a matrix of "synaptic" connections, creating patterns of activation across phoneme and semantic neurons. The phoneme and semantic neurons, in turn, feed activation back through a top-down matrix of connections, transforming their patterns back into letter patterns. Whenever the feedback patterns match the original letter pattern, top-down activation *conserves* bottom-up activation. Consequently, the "matched" letter neurons conserve their capacity to reactivate matching phoneme and semantic neurons that, in turn, reactivate the letter neurons, and so on. This feedback cycle is temporarily stable, resulting in a coherent dynamic whole—a *resonance*.

This neural network is only for exposition. It is helpful to consider word perception in terms of artificial neural activity, but the more precise analogy between nervous systems and cognitive systems is a hypothetical trajectory of sensorimotor dynamics (alternatively, a perception-action trajectory) that is correlated with cognitive performance. No claim is made concerning the "correct" architecture (see Stone & Van Orden, 1994; Van Orden & Goldinger, 1994; Van Orden et al, 1990; Van Orden, Holden, et al., submitted; Van Orden, Pennington, et al., submitted). Thus, with respect to the model's nodes, we are free to discuss word perception in cognitive system's terms. Figure 1 illustrates cognitive macrodynamics of word perception (Van Orden & Goldinger, 1994, 1996) and spelling (Bosman & Van Orden, in press), and Fig. 2 illustrates microdynamics.

Figure 1 portrays a recurrent network with three families of fully interdependent nodes (letter nodes, phoneme nodes, and semantic nodes). On average, the connections between node families differ in strength; the rank order of

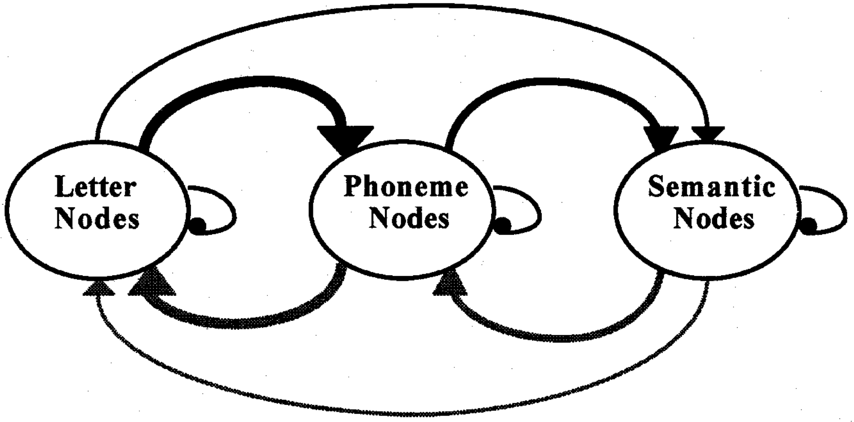

FIG. 1. Macrodynamics of reading and spelling performance emergent in a recurrent network. The boldness of the arrows indicates the overall strength of the relations between letter, phoneme, and semantic node families (see text).

overall strength is illustrated by the relative boldness of arrows in the figure. In alphabetic languages, letters and phonemes correlate quite strongly. For example, the letter B is almost always pronounced as [b], and the phoneme [b] is always written with a B. Correlations between phonemes and semantic features, or letters and semantic features, are far weaker than correlations between letters and phonemes. Knowing that a word begins with the letter B indicates almost nothing about its meaning, but much about its initial pronunciation.

Notice also that phoneme-semantic relations are depicted as stronger correlations than letter-semantic relations, primarily because we speak before and more often than we read. Once in place, this asymmetry is self-perpetuating. Reading strengthens phoneme-semantic connections, because phonology functions in every instance of printed word perception. Thus, even the exceptional condition of people who read more than they speak would support phoneme-semantic connections that are at least as strong as letter-semantic connections. Also, if a coherent positive feedback loop forms between phoneme and semantic nodes, before the feedback loop between letter and semantic nodes, then printed or spoken discourse may proceed without resolving the feedback loop between letter and semantic nodes. The absence of resonance in the latter feedback loop may preclude strengthening the connections between letter and semantic nodes (Grossberg & Stone, 1986). Thus, at this macro-level of description, node families differ in overall strength of relations with other node families. These differences in overall correlational structure are illustrated in the relative boldness of the arrows in Fig. 1.

The strong bidirectional connections between letter and phoneme nodes, as compared to those with semantic nodes, causes the letter-phoneme dynamic to cohere (resonate) before all others. This is the *phonologic coherence hypothesis*. The relatively consistent bidirectional covariance between letters (form) and phonemes (function) explains how phonology comes to be so fundamental in reading and spelling. Stated differently, it explains why sound-alike words (ROSE and ROWS) may be confused in reading (Van Orden, 1987); it explains why the majority of spelling errors (ROZE instead of ROSE) are phonologically acceptable; and it explains why patients' spoken and written errors often resemble the spelling and pronunciations of the correct response that was not produced. (Van Orden & Goldinger, 1994, 1996; Van Orden, Pennington, et al., submitted, describe various other phenomena that derive from the bidirectional covariance between spelling and phonology.)

In a model analogous to Fig. 1, presentation of a printed word activates letter nodes, which, in turn, activate phoneme and semantic nodes. Following initial activation, recurrent feedback begins among all these node families. Similarly, presentation of a spoken word activates phoneme nodes, which, in turn, activate semantic and letter nodes (and “picture naming” might be simulated with activation of semantic nodes which, in turn, activate phoneme and letter nodes, cf. Dell, Schwartz, Martin, Saffran, & Gagnon, in press). In all these cases, initial activation leads to recurrent feedback among all node families. However, the strongest recurrent dynamic is between letter and phoneme nodes, which creates the common basis of reading and spelling. The strength of a recurrent dynamic is a function of *self-consistency*—the capacity of a node or a family of nodes to conserve their own activation (Smolensky, 1986; Van Orden et al., 1990). Nodes conserve their activation when they "send" it to other nodes that "return" it in relatively exclusive recurrent feedback. This capacity to conserve activation derives from relatively consistent bidirectional covariance between nodes or node families—i.e. a history of structural coupling (cf. Varela et al., 1991). The bidirectional relation between letter and phoneme nodes is more self-consistent than other pairings of node families. Consequently, letter-phoneme dynamics supply the strongest and most generally reliable constraints on the model's performance.

Notice the difference between the description of phonologic mediation in the target articles and the phonologic coherence hypothesis. Linear flow-chart models use the term *phonologic mediation* to refer to a causally intermediate phonologic representation that is activated by spelling representations and, in turn, activates semantic representations (for example)[1]. This is why phonologic

---

[1] My (Van Orden's) view of phonologic mediation has changed in the last decade, as I have learned more of mathematical dynamic systems theory. I have moved from a representational connectionist view (e.g. Van Orden, 1987) to a "nonrepresentational" cognitive systems view (e.g. Van Orden & Goldinger, 1994). However, I do not recall proposing the straw man account that is attributed to me in the target articles. A more careful reading of my cited articles would find the following quotes:

> The extent to which phonology affects performance . . . is underscored by the simple verification model's relatively comprehensive account of [these] results . . . even though it lacked a mechanism of direct access. This is not to say that I deny the possibility of direct access (Van Orden, 1987, p. 192).
>
> A mechanism of covariant learning can also accomplish direct access in the same way that it accomplishes phonological coding . . . any linguistic features that frequently covary with orthographic features will become associated. The consequence . . . for any subsequent instance of lexical coding will be that, initially, a representation of the spelling of a word will activate most strongly those linguistic features (i.e. semantic, syntactic, and phonological features) that covary to the highest degree with its orthographic features (Van Orden, 1987, p. 194).
>
> This is not to say that we deny the possibility of direct bottom-up activation of lexical features by orthographic features. Rather, it may be useful to abandon the notion of separate, independent routes of lexical access. A potential alternative . . . is a connectionist mechanism . . . that . . . comes to reflect the covariance between all linguistic features (syntactic, semantic, and phonological) and orthographic features in its associative weights . . . (Van Orden, Johnston, & Hale, 1988, p. 382).

Continued overleaf

mediation has appeared inefficient and counterintuitive (Van Orden & Goldinger, 1994). Why should information processing traverse the same psychological distance in two steps (step one: orthography to phonology; step two: phonology to semantics), rather than one step (orthography to semantics)? By contrast, the phonologic coherence hypothesis implies that phonology's mediating effect in reading and writing is economical and efficient. Self-consistent feedback from phonology rapidly organises the system, and strongly constrains local competitions that would organise the visual stimulus. Subsequently, a coherent visual-phonologic dynamic "mediates" competitions among alternative global interpretations—their chances for survival are enhanced if they conform to extant, visual-phonologic dynamics. Phonologic "mediation" is inescapable in the simple model, due to the powerfully self-consistent relation between letter nodes and phoneme nodes.

Figure 2 illustrates microdynamics. Now we zoom in on the connectivity between letter and phoneme nodes (and ignore, for now, phoneme-semantic and letter-semantic connectivity). In Fig. 2a, reading the printed word HI includes activation of letter nodes H_1 and I_2, which activate the phoneme nodes $[h_1]$ and $[a^{I}_{2}]$, but also competing correlated nodes such as $[I_2]$ (as in [hIt]) which must be inhibited. (The subscripts refer to the positions of the letters or phonemes in words.) Figure 2b shows how, in turn, phoneme nodes feed activation back to letter nodes (illustrated for the phoneme nodes $[h_1]$ and $[a^{I}_{2}]$). The phoneme node $[a^{I}_{2}]$ activates the correct letter nodes H_1 and I_2 and also competing letter nodes, for example, the letter node Y_2 as in MY or BY. Thus, early patterns of activation are loosely structured. They include the activation of correct, but also many incorrect, candidates for resonance. This is due to *multistability*, a defining characteristic of dynamic systems. In our simple model, the dynamics from this point select a combination of nodes through cooperative-competitive dynamics.

Reliable performance emerges if the overall bidirectional configuration of connections favours mutual activation between the letter nodes H_1 and I_2 and the phoneme nodes $[h_1]$ and $[a^{I}_{2}]$. This advantage grows over time as the "strong grow stronger" and the "weak grow weaker" (cf. McClelland & Rumelhart,

---

> Name (phonologic) codes have a . . . processing advantage over other lexical codes. This advantage is not, however, an advantage in the relative time course of phonologic activation; phonologic codes are activated in parallel with other linguistic codes. Rather, relatively invariant phonologic codes are relatively stable pockets of lexical activity, and other lexical codes congeal around this relative stability (Van Orden et al., 1990, p. 513).

> Once again, we emphasize that it is not a difference in the time course of initial activation that distinguishes phonologic codes: All lexical codes are activated in parallel. Instead it is the difference in coherence times that constrains processing (Van Orden et al., 1990, p. 514).

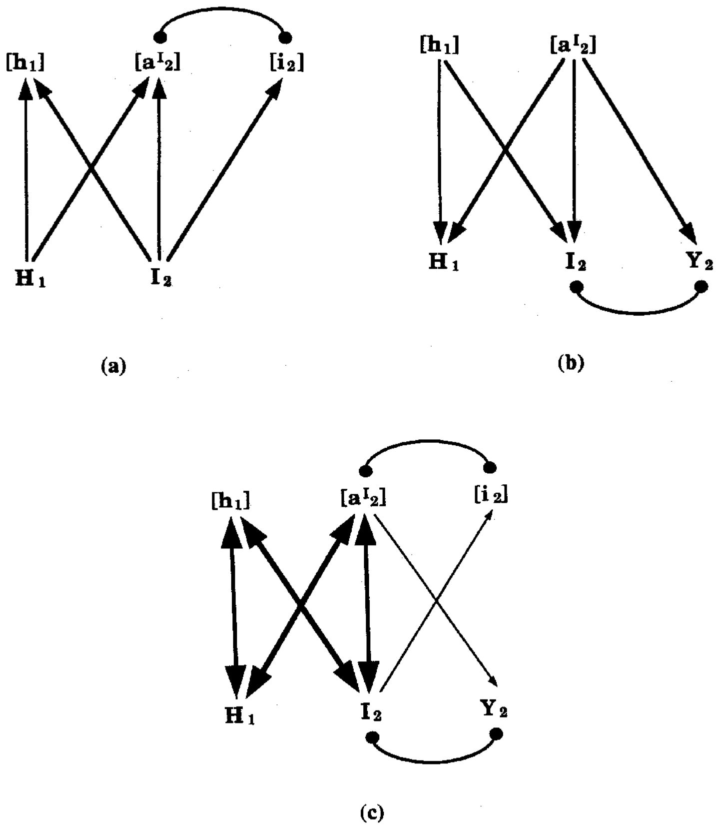

FIG. 2. A simplified illustration of microdynamics that "read aloud" the word HI. (a) Presented with HI, activation feeds forward from letter nodes to phoneme nodes. (b) In turn, phoneme nodes feed activation back to letter nodes. (c) A resonance that emerges between letter and phoneme nodes corresponding to HI. To reduce the number of lines in the figure, bidirectional connections are depicted with single double-headed arrows.

1981). This is illustrated in Fig. 2c, which combines the flow of activation from letter nodes to phoneme nodes and from phoneme nodes back to letter nodes, as assumed in a recurrent network. Presentation of the spoken word /ha^{I}/ to the network (as in a spelling task) leads to a similar dynamic between phoneme and letter nodes. Thus, activation initiated in phoneme nodes may generate a coherent pattern of activity across letter nodes.

THE DISSOCIATION BETWEEN READING AND SPELLING

Anyone who writes in English will experience occasional doubts about how to spell a word, but we almost never forget how a word should be read aloud. This dissociation of intact naming from intact spelling is evident for a variety of languages, at all levels of skill (Bosman & Van Orden, in press). This dissociation is also present, although exaggerated, in the performance of PW (Rapp et al., this issue) and WMA (Micelli et al., this issue). Both these patients correctly read aloud many more words than they correctly spell in written picture naming or to dictation. We do not suggest that their spelling performance is not affected by their lesions, just that this pattern respects the topology of intact naming and spelling. The discussion of these patients in the target articles does not take into account the topological properties of intact performance, but we eventually discuss why this concern is more salient from our perspective.

Our account explains why people find spelling more difficult than reading aloud. This dissociation illustrates one basis for dissociations in a "lesioned" recurrent network. The model may behave one way if activation of letter nodes drives a naming response, but differently if activation of phoneme nodes drives a spelling response. This dissociation may be described simply with respect to the previous illustrations of microdynamics (letter-phoneme dynamics), and macrodynamics (dynamics among node families).

Returning to Fig. 2, reading the word HI not only activates phoneme nodes $[h_1]$ and $[a^{I}_{2}]$, and the letter nodes H_1 and I_2, but also all possible pronunciations of H_1 and I_2 and all possible spellings of $[h_1]$ and $[a^{I}_{2}]$. Again, this is due to multistability. The same stimulus supports multiple percepts and actions. Multistability implies that reading a word correctly must include inhibition of incorrect phoneme nodes, and spelling a word correctly must include inhibition of incorrect letter nodes. In the case of reading, the letters are presented to the model (or reader). As a consequence, phoneme → letter ambiguity is highly unlikely to result in full activation of incorrect letter nodes, because persistent and stable environmental constraints (visible letters) accelerate correct feedback loops with phoneme and semantic nodes (as illustrated by bold arrows in Fig 2c). In the case of spelling, however, one must generate this resonant pattern from phonologic and semantic activation alone. In this case, the environment does not include explicit support for correct letter nodes.

Differences in environmental constraints strongly affect performance. For example, a patient may be much better at repeating auditorally presented pseudowords—i.e. explicit environmental support for phonology—than at reading pseudowords aloud—i.e. only implicit environmental support for phonology (Funnell, 1983). Likewise, a patient may be much better at copying printed words than at writing to dictation. This point is also more salient from a perspective that emphasises a history of covariance (structural coupling) between environmental forms and their cognitive functions. From a traditional view, the environment is equally relevant or irrelevant in both cases, because it must be represented symbolically.

The crux of spelling is that English orthography, generally, has more possible spellings for any given word than possible readings, and this is true of most writing systems (e.g. Stone, Vanhoy, & Van Orden, in press; Ziegler, Stone, & Jacobs, in press). Consider, for example, the multiple inconsistent "spelling bodies" that may correspond to the "rime" [—ûrch], —IRCH as in BIRCH, —ERCH as in PERCH, —URCH as in LURCH, and —EARCH as in SEARCH. Stone et al. (in press) estimated that 31% of low-frequency English one-syllable words are spelling → phonology inconsistent (at the grain-size of spelling-bodies and rimes), but fully 72% are phonology → spelling inconsistent (at the same grain-size). This estimate was corroborated in a larger sample, including both low- and high-frequency one-syllable words. Again, 72% of all spelling → phonology consistent words were phonology → spelling inconsistent. These linguistic analyses clearly indicate that phonology → spelling inconsistency is the rule for English (see also Ziegler, Jacobs, & Stone, in press, concerning French).

Although both reading and spelling are powerfully constrained by the strong correlational structure of letter-phoneme relations, the occasional inconsistencies in these relations are resolved by different sources of constraint. Now, we refer again to the illustration of macrodynamics in Fig. 1. When a model "reads" a low-frequency, spelling → phonology *inconsistent word* such as PINT, the more consistent letter-phoneme relation would rhyme with MINT (and HINT, LINT, TINT). Similarly, the letter-phoneme dynamic would yield two correct pronunciations for words like WIND (although it would typically favour the more regular pronunciation, Kawamoto & Zemblidge, 1992). In both these cases, relatively strong semantic-phoneme relations may supply sufficient secondary constraints to encourage the appropriate letter-phoneme dynamic. In the case of WIND, semantic constraints may also be due to context. In the model, contextual and stimulus sources of semantic constraints contribute via the relatively strong connections between semantic and phoneme nodes. Highly imageable or concrete words have stronger semantic correlations to letters and phonemes, which promotes better intact and patient performance (cf. Plaut & Shallice, 1993; Strain, Patterson, & Seidenberg, 1995). Also, added contextual support for correct performance contributes directly through semantic connections to activation of letter and phoneme nodes.

In the case of spelling, a model must resolve the inverted patterns of ambiguity in the phoneme-letter dynamic. To spell a low-frequency phonology → spelling inconsistent word such as HEAP, the rime [—ip]'s correct spelling would compete with a more strongly correlated incorrect spelling-body —EEP (as in DEEP, BEEP, KEEP, PEEP, SEEP, and WEEP). Additionally, the phoneme-letter dynamic yields two correct spellings for homophones (e.g. ROSE/ROWS). In either case, correct spelling must rely on relatively weak semantic-letter constraints (as illustrated in Fig. 1) to activate the appropriate letter nodes sufficiently—even contextual support is filtered through the weak letter-semantic connections. This weaker support for spelling, compared to the strong support for reading (i.e. phoneme-semantic constraints) is the "macrobasis" for the asymmetry between reading and spelling. Spelling is thus more difficult than reading for two reasons: microdynamic phoneme → letter relations are more inconsistent than letter → phoneme relations, and macrodynamic support for spelling (i.e. letter-semantic connections) is weaker than macrodynamic support for reading (i.e. phoneme-semantic connections).

SIMULATED BRAIN DAMAGE

The previous section demonstrates the utility of our model for understanding intact performance. In this section, we discuss how this approach may be extended to patient data. Our focus will be limited to a few theoretically important phenomena associated with acquired dyslexia, and a few key performance phenomena from the target articles. Marshall and Newcombe's (1973) classic article described distinct syndromes of acquired dyslexia: *surface dyslexia* and *deep dyslexia*. These syndromes are defined by characteristic profiles of naming errors. For example, deep dyslexics sometimes produce bizarre semantic errors (e.g. BUSH named as TREE), and surface dyslexics sometimes produce regularisation errors (e.g. PINT named to rhyme with MINT). Next, we describe how to produce similar errors in "lesioned" models that previously produced patterns of skilled naming performance.

Regularisation errors are characteristic of surface dyslexic patients, occurring when words such as PINT, with irregular pronunciations, are read aloud incorrectly to rhyme with similar regular words (e.g. MINT, HINT, and LINT). Although skilled readers also make regularisation errors in speeded naming tasks (Kawamoto & Zemblidge, 1992), surface dyslexic patients make many more. Regularisation errors are symptoms of multistable dynamics in intact naming. Multistability implies that multiple percepts and actions may arise to the same stimulus. For example, intact dynamics that lead to a correct pronunciation of PINT include the rhyme with MINT, which must be inhibited. The rhyme with MINT is locally more self-consistent because it is favoured by overall letter-phoneme covariance. In a model, the relatively late phoneme-semantic resonance strengthens the correct pronunciation of PINT and allows it to inhibit the "regularisation error" (Van Orden & Goldinger, 1994). Thus, regularisation errors in intact performance are much more likely when naming is speeded using a deadline procedure, and when a phoneme-semantic resonance is weaker as in low-frequency words. Also, regularisation errors in patient performance are correlated with semantic deficits that may imply reduced semantic constraints for inconsistent pronunciations (Patterson & Hodges, 1992; Patterson, Marshall, & Coltheart, 1985; Warrington, 1975). In a model, the reduced constraints from phoneme-semantic dynamics release the more regular pronunciation from the inhibition, resulting in a regularisation error.

Semantic errors are characteristic of deep dyslexic patients, occurring when words are read aloud incorrectly as semantically related words. For example, the word BUSH might be read aloud incorrectly as TREE. Semantic errors are also symptoms of multistable dynamics and release from inhibition. In a model, activation of semantic nodes leads, in the next time step, to activation of all letter and phoneme nodes which have previously covaried with any active semantic node. This allows the surface features of a competitor, such as TREE, to be activated by a semantically related stimulus word, such as BUSH. An intact model is saved from semantic errors because the letter nodes of BUSH are explicit in the environment and they are powerfully correlated with the phoneme nodes of BUSH. Thus, BUSH's letter-phoneme dynamic readily inhibits the surface features of TREE in intact performance.

The separate occurrence of semantic and regularisation errors has been incorrectly interpreted to be evidence against recurrent network models (see Van Orden, Pennington, et al., submitted, for a review and counter-argument). Farrar and Van Orden (1994) refuted this claim with an existence proof—i.e. a recurrent network model that produces these two error types. They began with a network very similar in structure to the simple illustrations presented in this commentary. Three families of nodes (see Fig. 1) were "taught" a sample of English words using a Hebbian-type learning algorithm (10 learning trials each for "high-frequency" words and 1 each for "low-frequency" words), until the model produced patterns of naming performance similar to those of skilled readers. The "naming response" was taken from the pattern of most active phonemes and "naming time" was defined as the number of cycles required to generate a coherent pattern of phonemes.) In particular, the model produced a frequency × consistency interaction. Low-frequency inconsistent words such as PINT were named more slowly than low-frequency consistent words such as DUCK, whereas all high-frequency words were named quickly (see Waters & Seidenberg, 1985).

To simulate the regularisation error, they added noise to the intact network. Noise was implemented as a uniform distribution of small increments of positive or negative activation added in each cycle to the activation values of randomly chosen nodes. The noise eroded the strength of phoneme-semantic attractors. In turn, this eroded the network's capacity for inhibiting regularisation errors. For example, instead of PINT's correct phonemes, the network produced activation on phoneme nodes that regularised PINT to rhyme with MINT. Specifically, noise destabilised the weakest phoneme-semantic attractors, which released from inhibition the powerful local constraints of "regular" letter-phoneme attractors. Because letter-phoneme dynamics primarily reflect the strongest correlations between letters and phonemes, they are naturally drawn into regularisation errors.

Importantly, Farrar and Van Orden (1994) could have simulated regularisation errors in several ways. A "lesion" could be implemented by reducing top-down activation from semantic nodes to phonologic nodes (Patterson et al., in press; Plaut, McClelland, Seidenberg, & Patterson, 1996). Their more subtle implementation added a uniform distribution of noise, which erodes the model's capacity to enter weakly self-consistent resonances (compare Lewenstein & Nowak, 1989). Alternatively, they could have introduced small changes in randomly chosen connection strengths. Similarly, the *locus* of noise need not be crucial. Bidirectional flow of activation means that noise introduced anywhere in the model spreads throughout the model in the next time step (although noise reduction does occur due to remote local inhibition). Because weaker "coarse-grain" phoneme-semantic relations yield less self-consistent dynamic structures, they are more vulnerable. Performances that rely on similar coarse-grain, visual-phoneme-semantic constraints include naming of irregular words, object and picture naming, and comprehension. These performances are typically deficient in surface dyslexia.

To simulate semantic errors, Farrar and Van Orden (1994) further "lesioned" the noisy network that previously produced regularisation errors. They set all of the letter-phoneme connections at zero, effectively "cutting" the connections (they could have cut fewer connections with the same effect; the minimum proportion that would produce semantic errors is interdependent with other modelling choices such as the amount of noise). Subsequently, the network produced semantic errors. When presented with BUSH, the network produced a relatively unstable pattern of activity across phoneme nodes corresponding to TREE. Setting the letter-phoneme connections to zero creates a highly unstable network, causing it to rely heavily on semantic-phoneme dynamics, the most reliable remaining source of constraints. However, in the absence of letter-phoneme constraints, semantic-phoneme dynamics are sometimes misled into a semantic error, and the relatively weaker letter-semantic dynamics cannot rescue the network from this error. Semantic errors are especially likely when semantic nodes of one word (BUSH) are strongly correlated with phoneme nodes of a different word (TREE) (see also Plaut & Shallice, 1993).

The Dissociation of Writing and Naming

So, how might we simulate dissociations of written and spoken responses in picture naming? Farrar and Van Orden (1994) did not construct a model of picture naming, but we can understand how such a model would behave by thinking of picture naming as a dynamic initiated from semantic nodes to letter and phoneme nodes. The crux of the present articles is that a dissociation between written and spoken picture names forces the inference that *intact* writing and naming include causally independent orthographic representations. Thus, any lesioned model that mimics this dissociation, but does not entail causally independent orthographic nodes, contradicts the basis of this inference. In this regard, we point out that the previous simulation of semantic errors already exhibits a dissociation of “written” letter node activation from “spoken” phoneme node activation, as we describe next.

Suppose that we presented “pictures” to a model that already produces semantic errors. An artificial lesion has disconnected letter and phoneme nodes and also generated a uniform distribution of noise. This results in a highly unstable system with no direct constraints between phoneme and letter nodes (however, they are still causally interdependent through indirect recurrent connections). Such a system may generate the same written and spoken names when a complete, relatively familiar semantic pattern resonates with phoneme and letter nodes. Alternatively, if the semantic pattern that initiates dynamics is incomplete then it may support different letter nodes and phoneme nodes. This is especially likely if the history of semantic-phoneme covariance favours a different response than the history of semantic-letter covariance (compare the previous dissociation of spelling and naming in intact performance). Of course, ablating a portion of the phoneme nodes would further degrade explicitly phonologic performance. (See Dell et al., in press, for additional modelling choices.)

We have described how an overly-simple recurrent network can dissociate written responses from spoken responses, but the intact version of this model does not include autonomous spelling representations. Moreover, it would be a mistake to imagine that the simple model is the only possible choice within the cognitive systems framework, or that “falsification” of the simple model impugnes our more general claims concerning causality. That reasoning would ignore the strong and broad theoretical basis of this framework in mathematical dynamic systems theory. Just as one may always construct a flow-chart model of performance, one may also always construct a theoretically meaningful recurrent network to mimic any performance profile of any complexity (Stone & Van Orden, 1994). We have described a minimal model that was trained to embody a subset of visual-phonologic-semantic covariant structure. Once trained, that model's behavioural trajectories suffice to mimic theoretically important intact performance, and to produce dissociations when lesioned.

The Topology of Performance

Empirical constraints on the construction of cognitive systems models are primarily derived from generic patterns in behaviour (Abraham & Shaw, 1992). This basis for rigorous qualitative analysis may be very useful for cognitive neuropsychology, which deals mostly with qualitative effects. For example, we have discussed generic predictions for intact naming and spelling that derive from self-consistency between spelling, phonology, and meaning (e.g. Bosman & Van Orden, in press; Van Orden & Goldinger, 1994). This performance topology may be tested against the behaviour of patients. Table 1 summarises the patients' performance from the target articles on six key tasks. Macro- and

TABLE 1
Representative Performance in Each of Six Tasks by the Patients PW, WMA, PS, and EA[a] (Proportion Correct)

Task	*Stimulus*	*Response*	*PW*	*WMA*	*PS*	*EA*
Copying	Printed word	Printed word	.94	.94	—	—
Repetition	Spoken word	Spoken word	.99	.97	.40	.98
Reading aloud	Printed word	Spoken word	.89	.78	.42	.40
Writing to dictation	Spoken word	Printed word	.44	.55	.76	.48
Picture naming: Spoken response	Picture	Spoken word	.72	.60	.51	.30
Picture naming: Written response	Picture	Printed word	.46	.44	.90	.98

[a] PW (Rapp et al., this issue); WMA (Miceli et al., this issue); PS (Hanley & McDonnell, this issue); EA (Shelton & Weinrich, this issue).

micropatterns of self-consistency supply qualitative predictions for several contrasts between tasks. The utility of this approach is illustrated using the overall profiles of PW (Rapp et al., this issue) and WMA (Miceli et al., this issue), so we will discuss their profiles first. Following that we discuss how this approach might accommodate the less agreeable profiles of PS (Hanley & McDonnell, this issue) and EA (Shelton & Weinrich, this issue).

In Fig. 1, we described the macrodynamics of naming and spelling; dynamics between phonology and semantics are more self-consistent than dynamics between spelling and semantics. In Fig. 2, we described the microdynamics that pertain to naming and spelling; English spelling-to-phonology is more consistent than phonology-to-spelling. We also noted several times that our approach assumes explicit direct constraint from the environment (Gibson, 1986; Turvey & Carello, 1981; van Leeuwen, Steyvers, & Nooter, submitted; Van Orden & Goldinger, 1994). Copying printed words or repeating spoken words both have relatively transparent relations between performance and environmental constraints. In the case of repetition, acoustic form and articulatory function have an ancient history of structural coupling, which explains why the two modalities have virtually isomorphic descriptive features. The practice of copying printed words doesn't go quite so far back in the history of our species, but there is a more general ancient structural coupling entailed by copying, as in drawing. Moreover, the presence of the stimulus during performance explicitly supports memory, as well as feedback for error correction. Thus, all other things equal, we expect performance on these tasks to be relatively less vulnerable to brain damage. PW and WMA demonstrate virtual ceiling performance on these tasks, but on no other tasks (see Table 1).

The relation between spelling and phonology is more consistent than the relation between phonology and spelling. Also, printed letters are explicit in the environment, but phonology is derived. Consequently, the relation that supports naming is more self-consistent, and reading aloud is superior to writing to dictation. The predicted direction of this contrast agrees with the performance of PW and WMA. They both show better performance when reading aloud than when writing to dictation.

Likewise, because we speak before and more often than we write, the history of structural coupling that supports spoken responses in picture naming is more self-consistent than the support for written responses in picture naming. The predicted direction of this contrast also agrees with the performance of PW and WMA. Additionally, spoken responses in picture naming are supported by less self-consistent relations than spoken responses in printed word naming. The predicted direction of this contrast also agrees with the performance of PW and WMA. Altogether, these outcomes weave an agreeable web of support in line with the described topology of intact behaviour.

Less Agreeable Patterns

PS (Hanley & McDonnell, this issue) and EA (Shelton & Weinrich, this issue) present us with patterns more complex than the performance topology of our simple model. It is important to understand why this does not falsify our approach, and what it means for fleshing out the topology of a more inclusive model. For example, PS is generally poorer in tasks that require spoken responses than in tasks that require written responses; even repetition performance is very poor. Specifically, PS's profile contradicts three natural predictions: (1) repetition ≥ reading aloud, writing to dictation, and both forms of picture naming; (2) reading aloud ≥ writing to dictation; (3) picture naming with spoken response ≥ picture naming with written response. However, the model is readily expanded to become more inclusive.

For example, the phoneme nodes would be more appropriately reconstituted as emergent properties of recurrent acoustic-articulatory dynamics. Then we could “lesion” the connections between acoustic nodes and articulatory nodes, or add noise to randomly chosen connections, or ablate selected articulatory nodes. We could even reconstitute the model to allow complex oscillations, as in dynamic models of speech production (Browman & Goldstein, 1995). Clearly we are not short of options, and all these options are directly in line with the principled basis of our analysis. Any of our options for damaging phonology would suffice to accommodate the overall dissociation in PS's profile. Namely, performance that requires a spoken response is worse than performance that allows a written response (see Table 1).

EA presents another interesting and complex profile. EA had extensive speech therapy, but he remains nonfluent. He only produces sentences like

"How are you" and "Good morning." He writes single words to communicate, but he cannot read them aloud. He is nonfluent and yet performs well in repetition. Except for repetition, EA's performance is low on any task that involves a spoken word response. To simulate EA's profile we would need the previous reconstituted model. We could then add weak uniform noise to acoustic-articulatory dynamics, cut the connections to letter nodes, and cut a portion of connections to semantic nodes. Subsequently, any relatively weak acoustic-articulatory attractors are isolated from sources of constraint other than the preserved semantic constraints, consonant with better repetition performance for imageable/concrete words. Like all other effects, imageability/concreteness effects are predicted in terms of self-consistency. In this case, self-consistency results from covariance of words with contexts. Highly imageable/concrete words vary less in meaning across contexts, which builds more self-consistent relations with their surface forms (Van Orden, Pennington, et al., submitted, and cf. Jones, 1985; Saffran, Schwartz, & Marin, 1979; Shallice, 1988). Nonword repetition is possible, but nonwords would not have the words' advantage of learned whole-word attractors, including phoneme-semantic attractors, so a slight deficit in nonword repetition is not surprising.

Please don't get the impression that building an actual model is a piece of cake; it is not. When all parts of a model are interdependent it can take quite a bit of work to explore the parameter space of the model and arrive at the empirical topology (just ask our friend Bill Farrar). However, success is assured for any reliable empirical topology. The explanatory power of these models does not reside in specific parameter settings; it resides in the general topological principles from which they are constructed.

The profiles of PS and EA have not exhausted our options. We may continue to reconstitute nodes, making use of reliable finer-grain relations among stimulus forms and cognitive functions, without doing any violence to the overarching framework. As we noted, it is even possible to reconstitute point attractors to accommodate complex time-varying behaviour such as that entailed by on-line articulatory gestures. Likewise, the node activation denoted as *semantic* could be expanded as emergent sensorimotor ensembles, including visual-acoustic-articulatory-postural-gestural-etc. ensembles (Allport, 1983; compare *image schemas* in Gibbs, 1994; Johnson, 1987; Lakoff, 1987). The crux of our analysis is not the discovery of correct nodes, nor correct oscillations; it is the utility of general mathematical (topological) principles at all scales of analysis (Abraham & Shaw, 1992).

The circular relation between the way we view data and the theoretical basis of our analysis presents no more problems than the circular relation between a linear reduction of data and the linear componential models that are then inferred. Data cannot decide between linear and nonlinear approaches to human performance. Consequently, claims that require the "truth" of one or the other perspective can only be supported by acts of faith. The implication is simply this: We retain an explicit and healthy scepticism toward the theories we propose, keeping one eye open for alternative workable frameworks. We shrug off tyrannical “objective” truth, exercise pluralism, and keep pragmatic concerns foremost (Lakoff, 1987).

MULTISTABILITY, METASTABILITY, AND INTACT PERFORMANCE

So far, we have described intact naming and spelling, and dissociations between naming and spelling, in ways that are as plausible as any flow-chart model, but not more plausible. One pragmatic test of a theoretical framework is whether we may learn something new about the systems we study—something that we might not have learned without the guiding framework. The flow-chart models described in the target articles have already passed this test (Carello et al., 1992). We need only track the history of reading research from the seminal articles of Coltheart (1978) and Marshall and Newcombe (1973), to validate the utility of flow-chart models for oganising and generating new findings. An explosion of studies have described important and reliable patterns in human performance. Recurrent network models and their entailed cognitive systems framework have also begun to demonstrate this utility. As we describe next, this framework has produced remarkable findings that are highly unlikely from a traditional perspective.

The phenomena we describe are not widely appreciated in cognitive psychology and neuropsychology. Consequently, it is easy to miss the fact that they converge within a cognitive systems approach. Each of the phenomena that we will describe pertain to printed and spoken word perception. However, this narrow convergence is only a small set of the large variety of reported findings that motivate this framework. From the broader perspective, an exciting possibility has taken shape. Perception and action may be generally and usefully described as the products of a self-organising complex dynamic system (Kelso et al., 1995).

Feedback and Multistability

The simple model described in the previous sections predicts a rather nonintuitive microeffect. This prediction derives from a common feature of dynamic systems—recurrent feedback. The specific prediction concerns multistability in the performance of tasks related to letters and phonemes, operationalised as ambiguity or *inconsistency*. Until recently, all discussion of consistency has concerned a classic, feedforward, spelling → phonology effect. —INT is inconsistent in PINT because it may be pronounced as in MINT; UCK is pronounced consistently as in DUCK. Inconsistent words such as PINT are named more slowly than consistent words such as DUCK. Thus, the feedforward consistency effect answers the question: Does it matter in *visual* word perception that a spelling may have more than one *pronunciation*? From the perspective of flow-chart models or feedforward connectionism, this is the only sensible question: The letter string is unambiguous to subjects (it is right in front of their eyes); the only potential ambiguity arises with respect to derived phonology. Once we consider perception as a product of recurrent feedback, however, the concept of perceptual ambiguity must be generalised—we must consider consistency in the *feedback* direction as well. Now we ask the feedback question: Does it matter in *visual* word perception that a *pronunciation* may have more than one *spelling*? From the perspective of resonant dynamics, feedback consistency should affect performance as strongly as classic, feedforward consistency.

Stone et al. (in press) tested for both feedforward and feedback consistency effects in a lexical decision task. They used a factorial design that included four types of words. In bidirectionally consistent words such as DUCK, the spelling body (___UCK) can only be pronounced one way, and the pronunciation body (/___uk/) is only spelled one way. In spelling → phonology inconsistent words such as MOTH, the spelling body can be pronounced in multiple ways (e.g. BOTH), but the pronunciation body (/___ɑ̂th/) is only spelled one way. In phonology → spelling inconsistent words such as HURL, the spelling body is pronounced in only one way, but the pronunciation body can be spelled in more than one way (e.g. GIRL). In bidirectionally inconsistent words such as WORM, the spelling body can be pronounced in multiple ways (e.g. DORM), and the pronunciation body can be spelled in multiple ways (e.g. FIRM). Stone et al. found strong evidence for perception as a "two-way street." Correct response times were equally (and strongly) slowed by both feedforward and feedback inconsistency. Additionally, they found a reliable interaction; all inconsistent words produced approximately equal response times, even those that were inconsistent in both directions. Only words that were bidirectionally consistent produced faster and more accurate performance.

The feedback consistency effect is compelling, for several reasons. First, it underscores the importance of bidirectional dynamics in perception. Second, it demonstrates that stimulus function (in this case, a word's "name function") lends perceptual structure to stimulus form. Again, note the nonintuitive nature of this phenomenon. The letter string is clearly visible to the subject, and it remains visible until a response is recorded. However, if feedback from phonology suggests that some *other* letter-string *could have* been presented, performance is slower. Third, it could only be predicted by a theory emphasising bidirectional dynamics. It is straightforward corroboration for bidirectional multistability in subword dynamics. In a simple model, multistable (ambiguous) pronunciations and spellings are resolved through successive cycles of cooperative (excitatory) and competitive (inhibitory) feedback.

Ziegler and his colleagues observed a similar counter-intuitive effect in a letter-search task (Ziegler & Jacobs, 1995; Ziegler, Van Orden, & Jacobs, in press). Subjects in this experiment were briefly presented with a letter string such as BRANE (a pseudohomophone of the word “brain”), followed by a pattern mask (#####). The subjects were instructed to respond whether a predesignated letter was present in the masked letter string, for example the letter “i”. In the case of BRANE, they (mis)reported having seen the letter “i” more often than in a control stimulus. Similarly, they failed to report the letter “i” in the letter string TAIP (a pseudohomophone of the word “tape”), more often than in a control stimulus. Presumably, the phonology of the pseudohomophones BRANE or TAIP suggested that “brain” or “tape” were presented, causing subjects to misreport the presence or absence of the letter “i.” These results also corroborate the description of word perception as a multistable dynamic system.

Hysteresis and Multistability

Betty Tuller and her colleagues have demonstrated hysteresis in speech perception (Tuller, Case, Ding, & Kelso, 1994). Hysteresis effects are a well-defined signature of multistability in nonlinear systems. In the present example, this means there are multiple perceptions of the same spoken stimulus. Tuller et al. focused on the classic phenomenon of categorical speech perception (Liberman, Harris, Hoffman, & Griffith, 1957). In one experiment, they manipulated the presentation order of speech stimuli. These stimuli were constructed to morph between the words SAY and STAY—their acoustic properties changed incrementally along a continuum from SAY to STAY (cf. Best, Morrongiello, & Robson, 1981; Hodgson & Miller, 1992). Each run of Tuller et al.'s experiment presented a subject with this continuum running from SAY to STAY and back again (or vice versa). Hysteresis was observed on 41% of runs across subjects and conditions. Specifically, for some intermediate range of stimuli, a subject perceived this range of stimuli as SAY if it had been preceded by SAY stimuli, but they perceived the identical range as STAY if it had been preceded by STAY stimuli. The intermediate range is thus multistable.

As noted, identical stimuli were perceived as SAY or STAY depending on preceding stimuli. Tuller et al. (1994) were not concerned with which reported stimulus identity is the “true” representation. They explored instead the interdependence of context and stimulus. They asked: What is the pattern of interaction between perceivers and contexts that characterises categorical speech perception? The hysteresis pattern is a generic pattern that is observed widely in physical, chemical, biological, and cognitive systems. Historically, in psychology, hysteresis has been considered a nuisance effect. For example, it motivated Fechner's *method of limits* in classical psychophysics—effectively, a statistical technique to make hysteresis disappear. Currently, hysteresis is better understood with respect to multistability. Thus, although Tuller et al.'s results were not derived from the simple model we described, hysteresis is convergent on the theoretical basis of this model. This fundamental construct provides a general and natural basis for understanding ubiquitous multistability in language, e.g. feedforward and feedback inconsistency, homography (LEAD or WIND), homophony (ROWS vs. ROSE), polysemy (Pete Rose is OVER the hill. vs. There were flies all OVER the ceiling—from Lakoff, 1987), and syntactic/semantic ambiguity (The church pardons very few people. vs. The church pardons are difficult to obtain—adapted from Rayner & Pollatsek, 1989).

1/f Noise and Metastability

David Gilden and his colleagues have demonstrated a complex interdependence between trial-by-trial response times in several cognitive tasks, including a word recognition task (Gilden, in press; Gilden, Thornton, & Mallon, 1995). The source of this interdependence is fully cognitive in affiliation (it does not arise in a "noncognitive" simple reaction time task, for example—Gilden et al., 1995). Also, mundane sources of priming such as DOCTOR → NURSE semantic priming, MINT → HINT form priming, or successions of identical responses were ruled out as potential sources (Gilden, in press). Instead, the source may be a metastable complex dynamic process. Metastability implies that a system never settles fully in a dominant attractor, and is thus more flexible. Remember the rough analogy from the earlier section on neurobiology: A person who *always* has alternative plans B, C, and D at the ready, is more flexible (metastable) than a person who is stuck in plan A. Metastability has been proposed to explain the smooth flexibility of perception and action.

Metastability in response time (and neural activity) is revealed in a rather esoteric phenomenon—*1/f noise*—observed in the "error" variance of response time (the variance left over when treatment effects are partitioned out). This phenomenon can be very difficult to grasp because it goes so strongly against the grain of typical psychological analyses. After all, we're used to discarding error variance, not analysing it for structure. 1/f noise is a mathematically generic pattern expressed here in trial by trial response time. If we graph each response time in a sequence, the data points will oscillate between fast and slow response times throughout the series of trials. If we "connect the dots," they form a complex waveform. In turn, this complex waveform may be viewed as a composite of waves that span a large range of frequencies. 1/f noise is a weak inverse relation between "power" (amplitude of change in response time) and the frequency of composite waves.

This correlation exists between changes in response time separated by small, intermediate, or large intervals. These nested correlations comprise a well-defined mathematical object from fractal geometry. Fractal objects are self-similar. So far, 1/f noise is only seen in the time domain, but self-similarity is easier to grasp in terms of ordinary objects that are extended in space. For example, a coastline is self-similar; it has a ragged structure of the same complexity when viewed from outer space, from an aeroplane, or from a cliff above the shore.

So what do we make of this rarefied phenomenon? On the down side, the presence of 1/f noise in response time is inconsistent with the conventional logic of partitioning response time into independent sources of variance, as in ANOVA. Virtually every study in which response time is the dependent measure averages response times across trials, items, and subjects. Comparisons are always conducted between means. One key assumption of this practice is that the response time in each trial is independent of the response times in other trials. The assumption of independence is at the heart of the linear statistical models that are used to discover independent sources of variance. The presence of 1/f noise contradicts this assumption. Thus, generally, we may have to rework the relation between data, method, and theory in response time studies. And, specifically, we are justified in our scepticism concerning the *effect = structure* assumption. It requires that data may be strictly carved at their joints to yield independent effects.

On the up side, 1/f noise may be a signature of metastability. If so, then it converges with the previous phenomena in this section to corroborate the utility of our approach. 1/f noise accounts for a stunning 70% (or more) of subjects' variance in response time for each of the cognitive tasks that were evaluated (Gilden, in press). By comparison, the best conventional models of word naming account for 3–12% of variance in average naming times (Besner, in press).

We wish to stress two caveats before we leave this section. First, 1/f noise does not confirm our simple model of naming. In fact, our simple model does not produce this phenomenon. The corroboration is for the cognitive systems framework that we work within. The simple model produces lower-dimensional behaviour than a more inclusive model that would produce 1/f noise. Once again, each trajectory produced by the simple model is a low-dimensional shadow of a trajectory that would be produced by a higher-dimensional model.

Second, our main point in reviewing these phenomena was to add some meat to the utility of our approach. The previous findings do not falsify the more conventional linear analyses of cognitive systems. Nonlinear accounts cannot be distinguished from linear accounts on the basis of correspondence to data —even data as compelling as those of Gilden (in press; Gilden et al., 1995) and Tuller et al. (1994). A linear componential model could be constructed for any data set, given enough components (Stone & Van Orden, 1993). Even the analytic difficulties raised by Gilden's data could be overcome if one wished to ignore the nonlinear perspective. Data can always be "corrected" to eliminate interdependence between successive trials (West & Hepworth, 1991). This practice can be very useful depending upon the goals of the analysis. However, in the case of Gilden's data, it would effectively send to the trash can 70% (or more) of theoretically meaningful variance.

Our general point is simply this: The basis for applying any particular framework is exclusively pragmatic. The practical utility of strongly nonlinear models is demonstrated as they become successful guides to a more general and inclusive understanding of cognitive phenomena. The practical utility of linear models is questioned if they would fail to discover theoretically meaningful phenomena. In that regard, the phenomena reviewed in this section question the utility of the *effect = structure* assumption. These phenomena are antithetical to that assumption, and they could never have been anticipated from the standard approach to patient data.

SUMMARY AND CONCLUSIONS

Dissociations cannot be trusted to isolate independent representations. The existence of plausible alternatives undermines their reliability for reducing performance phenomena to single causes (Van Orden, Pennington, et al., submitted). Induction of single causes required a priori the truth of single causes. Still, it is natural to prefer the familiar relation between componential methods, theory, and data. Moreover, there seems to be a general utility in such a pursuit, at least initially. Bechtel and Richardson (1993) review historic analyses in which, initially, scientists find it useful to assume linear *independence* (single causes) as a working hypothesis. Be that as it may, behaviour characteristics of linear independence are a small subset of the generic behaviours characteristic of complex dynamic systems (Abraham & Shaw, 1992; Giunti, 1995); they cannot be trusted to validate the *effect = structure* assumption.

Some readers might be tempted to claim that letter, phoneme, and semantic nodes are themselves independent representations. But that would imbue these notational distinctions with causal properties that they cannot bear (cf. Anderson, 1991; Perrone & Basti, 1995; Putnam, 1981). At most, nodes denote emergent form-function dynamics (structural couplings) that are not explicit in a model. As we noted, the node activation denoted as *phonologic* could emerge in recurrent acoustic-articulatory dynamics, and the node activation denoted as *semantic* could be effected in metastable sensorimotor trajectories, including visual-acoustic-articulatory-postural-gestural-etc. trajectories. Thus, the activation trajectories denoted as semantic emerge from the same stuff—i.e. the same excitable medium—as letter and phoneme "nodes." The simple model system is not real, except in its capacity to produce a trajectory that mimics human performance. It is plausibly interpreted as a low-dimensional projection—again, the shadow—of a vastly higher-dimensional description of readers, writers, texts, and laboratories.

No reliable basis exists for the discovery of primitive causal structures in human performance (Uttal, 1990; Van Orden, Holden, et al., submitted; Van Orden, Pennington, et al., submitted). Complex natural systems may not give up the true bases of their behaviour (Cohen & Stewart, 1994; Goodwin, 1994; Lindley, 1993). We accept this complexity and refuse to reify modelling notations such as nodes. Nodes are chosen to accommodate the grain-size of organism-environment coupling that predicts behavioural phenomena (Van Orden et al., 1990; Van Orden & Goldinger, 1994, and cf. Varela et al., 1991). Again, the explanatory isomorphism runs from described performance to a model's dynamic trajectory. Data make reference to such trajectories, exclusively; there are no data leftovers from which to deduce static atomic structures in an observed behaviour. Consequently, from a cognitive systems' perspective, the *effect = structure* assumption is not wrong; it is simply impracticable.

REFERENCES

Abraham, R.H., & Shaw, C.D. (1992). *Dynamics: The geometry of behavior*. Redwood City, CA: Addison-Wesley.

Allport, D.A. (1983). Distributed memory, modular subsystems and dysphasia. In S. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia* (pp. 32–60). London: Churchill Livingstone.

Anderson, J.R. (1991). The place of cognitive architectures in a rational analysis. In K. VanLehn (Ed.), *Architectures for intelligence: The twenty-second Carnegie symposium on cognition* (pp. 1–24). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Bechtel, W., & Richardson, R.C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research.* Princeton, NJ: Princeton University Press.

Berent, I., & Perfetti, C.A. (1995). A rose is a REEZ: The two-cycles model of phonology assembly in reading English. *Psychological Review, 102*, 146–184.

Besner, D. (in press). Basic processes in reading: Multiple routines in localist and connectionist models. To appear in P.A. McMullen & R.M. Klein (Eds.), *Converging methods for understanding reading and dyslexia*. Cambridge, MA: MIT Press.

Best, C., Morrongiello, B., & Robson, R. (1981). Perceptual equivalence of acoustic cues in speech and nonspeech perception. *Perception and Psychophysics, 29*, 191–211.

Bosman, A.M.T., & Van Orden, G.C. (in press). Why spelling is more difficult than reading. In C.A. Perfetti, M. Fayol, & L. Rieben (Eds.), *Learning to spell*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Bressler, S.L., Coppola, R., & Nakamura, R. (1993). Episodic multiregional cortical coherence at multiple frequencies during visual task performance. *Nature, 366*, 153–156.

Browman, C.P., & Goldstein, L. (1995). Dynamics and articulatory phonology. In R.F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 175–193). Cambridge, MA: MIT Press.

Carello, C., Turvey, M.T., & Lukatela, G. (1992). Can theories of word recognition remain stubbornly nonphonological? In R. Frost & L. Katz (Eds.), *Orthography, phonology, morphology, and meaning* (pp. 211–226). Amsterdam: North-Holland.

Cohen, J., & Stewart, I. (1994). *The collapse of chaos: Discovering simplicity in a complex world.* New York: Viking.

Coltheart, M. (1978). Lexical access in simple reading tasks. In G. Underwood (Ed.), *Strategies in information processing* (pp. 151–216). London: Academic Press.

Dell, G.S., Schwartz, M.F., Martin, N., Saffran, E.M., & Gagnon, D.A. (in press). Lexical access in aphasic and nonaphasic speakers. *Psychological Review.*

Duhem, P. (1954). *Aim and structure of physical theory.* York: Antheneum. (Originally published in 1906.)

Farrar, W.T., & Van Orden, G.C. (1994). *Simulation of surface and deep dyslexia in a unified network.* Poster presented at the Annual Meeting of the Psychonomic Society, St. Louis, November.

Freeman, W.J. (1987). Simulation of chaotic EEG patterns with a dynamic model of the olfactory system. *Biological Cybernetics, 56*, 139–150.

Freeman, W.J. (1991a). The physiology of perception. *Scientific American, 264*, 78–85.

Freeman, W.J. (1991b). What are the state variables for modeling brain dynamics with neural networks? In H.G. Schuster (Ed.), *Nonlinear dynamics and neuronal networks* (pp. 243–255). New York: VCH Publishers.

Freeman, W.J. (1995). *Societies of brains: A study in the neuroscience of love and hate.* Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Frost, R. (submitted). *Toward a strong phonological model of reading: True issues and false trails.* Manuscript submitted for publication.

Funnell, E. (1983). Phonological processes in reading: New evidence from acquired dyslexia. *British Journal of Psychology, 74*, 159–180.

Gibbs, P., & Van Orden, G.C. (submitted). *Pathway selection's utility for control of word recognition.* Manuscript submitted for publication.

Gibbs, R.W. (1994). *The poetics of mind: Figurative thought, language, and understanding.* New York: Cambridge University Press.

Gibson, J.J. (1986). *An ecological approach to visual perception.* (Original work published 1979.) Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Gilden, D.L. (in press). Fluctuations in the time required for elementary decisions. *Psychological Science.*

Gilden, D.L., Thornton, T., & Mallon, M.W. (1995). 1/f noise in human cognition. *Science, 267*, 1837–1839.

Giunti, M. (1995). Dynamical models of cognition. In R.F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 549–571). Cambridge, MA: MIT Press.

Goodwin, B. (1994). *How the leopard changed its spots: The evolution of complexity.* New York: Scribner.

Grossberg, S., & Stone, G.O. (1986). Neural dynamics of word recognition and recall: Priming, learning, and resonance. *Psychological Review, 93*, 46–74.

Haken, H. (1984). *The science of structure: Synergetics.* New York: Van Nostrand Reinhold.

Hatfield, F.M., & Patterson, K.E. (1983). Phonological spelling. *Quarterly Journal of Experimental Psychology, 35A*, 451–468.

Hodgson, P., & Miller, J.L. (1992). Internal phonetic category structure depends on multiple acoustic properties: Evidence for within-category trading relations (abstract). *Journal of the Acoustical Society of America, 92*, 2464.

Hubel, D.H., & Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology (London), 160*, 106–154.

Hubel, D.H., & Wiesel, T.N. (1965). Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. *Journal of Neurophysiology, 28*, 229–289.

Hubel, D.H., & Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London), 195*, 215–243.

Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination, and reason.* Chicago, IL: University of Chicago Press.

Jones, G.V. (1985). Deep dyslexia, imageability, and ease of predication. *Brain and Language, 24*, 1–19.

Katz, L., & Frost, R. (1992). The reading process is different for different orthographies: The orthographic depth hypothesis. In R. Frost & L. Katz (Eds.), *Orthography, phonology, morphology, and meaning* (pp. 67–84). Amsterdam: North-Holland.

Kawamoto, A., & Zemblidge, J. (1992). Pronunciation of homographs. *Journal of Memory and Language, 31*, 349–374.

Kelso, J.A.S., Case, P., Holroyd, T., Horvath, E., Raczaszek, J., Tuller, B., & Ding, M. (1995). Multistability and metastability in perceptual and brain dynamics. In P. Kruse & M. Stadler (Eds.), *Ambiguity in mind and nature* (pp. 159–184). Berlin: Springer-Verlag.

Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind.* Chicago, IL: University of Chicago Press.

Lewenstein, M., & Nowak, A. (1989). Recognition with self-control in neural networks. *Physical Review A, 40*, 4652–4664.

Liberman, A.M., Harris, K.S., Hoffman, H.S., & Griffith, B.C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology, 54*, 358–368.

Lindley, D. (1993). *The end of physics: The myth of a unified theory*. New York: Basic Books.

Livingstone, M.S., & Hubel, D.H. (1988). Segregation of form, colour, movement and depth: Anatomy, physiology and perception. *Science, 240*, 740–749.

Mackie, J.L. (1974). *The cement of the universe.* Oxford: Oxford University Press.

Marshall, J.C., & Newcombe, F. (1973). Patterns of paralexia: A psycholinguistic approach. *Journal of Psycholinguistic Research, 2*, 175–199.

McClelland, J.L., & Rumelhart, D.E. (1981). An interactive-activation model of context effects in letter perception. Part 1: An account of basic findings. *Psychological Review, 88*, 375–407.

Mill, J.S. (1974). A system of logic. J.M. Robson (Ed.), *Collected works of John Stuart Mill, Vol. VII.* Toronto: University of Toronto Press.

Patterson, K.E. (1981). Neuropsychological approaches to the study of reading. *British Journal of Psychology, 72*, 151–174.

Patterson, K.E., & Hodges, J.R. (1992). Deterioration of word meaning: Implications for reading. *Neuropsychologia, 30*, 1025–1040.

Patterson, K.E., Marshall, J.C., & Coltheart, M. (Eds.) (1985). *Surface dyslexia.* London: Lawrence Erlbaum Associates Inc.

Patterson, K.E., Plaut, D.C., McClelland, J.L., Seidenberg, M.S., Behrman, M., & Hodges, J. (in press). Connections and disconnections: A connectionist account of surface dyslexia. To appear in J. Reggia, R. Berndt, & E. Ruppin (Eds.), *Neural modeling of cognitive and brain disorders*. New York: World Scientific.

Pennington, B.F. (1991). *Diagnosing learning disorders: A neuropsychological framework.* New York: Guilford Press.

Perfetti, C.A. (1985). *Reading ability.* New York: Oxford University Press.

Perfetti, C.A., Zhang, S., & Berent, I. (1992). Reading in English and Chinese: Evidence for a “universal” phonological principle. In R. Frost & L. Katz (Eds.), *Orthography, phonology, morphology, and meaning* (pp. 227–248). Amsterdam: North-Holland.

Perrone, A.L., & Basti, G. (1995). Neural images and neural coding. *Behavioral and Brain Sciences, 18*, 368–369.

Plaut, D.C., McClelland, J.L., Seidenberg, M.S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review, 103*, 56–115.

Plaut, D.C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology, 10*, 377–500.

Prigogine, I., & Stengers, I. (1984). *Order out of chaos: Man's new dialogue with nature.* New York: Bantam Books.

Putnam, H. (1981). Reductionism and the nature of psychology. In J. Haugeland (Ed.), *Mind design* (pp. 205–219). Cambridge, MA: MIT Press.

Quine, W.V.O. (1961). Two dogmas of empiricism. In W.V.O. Quine (Ed.), *From a logical point of view* (pp. 20–46). (Originally published in 1953.) New York: Harper & Row.

Rayner, K., & Pollatsek, A. (1989). *The psychology of reading*. Englewood Cliffs, NJ: Prentice Hall.

Saffran, E.M., Schwartz, M.F., & Marin, O.S.M. (1979). *Neuropsychological evidence for mechanisms of reading: 1. Deep dyslexia.* Paper presented to the Seventh International Neuropsychology Society Meeting, New York.

Saltzman, E.L. (1995). Dynamics and coordinate systems in skilled sensorimotor activity. In R.F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 149–173). Cambridge, MA: MIT Press.

Shallice, T. (1988). *From neuropsychology to mental structure.* New York: Cambridge University Press.

Shanon, B. (1993). *The representational and the presentational: An essay on cognition and the study of the mind.* New York: Harvester Wheatsheaf.

Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology, 55*, 349–374.

Skarda, C.A., & Freeman, W.J. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences, 10*, 161–195.

Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 1. Foundations* (pp. 194–281). Cambridge, MA: MIT Press.

Stone, G.O., Vanhoy, M., & Van Orden, G.C. (in press). Perception is a two-way street: Feedforward and feedback phonology in visual word recognition. *Journal of Memory and Language*.

Stone, G.O., & Van Orden, G.C. (1993). Strategic processes in printed word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 744–774.

Stone, G.O., & Van Orden, G.C. (1994). Building a resonance framework for word recognition using design and system principles. *Journal of Experimental Psychology: Human Perception and Performance, 20*, 1248–1268.

Strain, E., Patterson, K., & Seidenberg, M.S. (1995). Semantic effects in single-word naming. *Journal of Experimental Psychology: Learning, Memory and Cognition, 21*, 1140–1154.

Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge, MA: Harvard Press.

Tuller, B., Case, P., Ding, M., & Kelso, J.A.S. (1994). The nonlinear dynamics of speech categorization. *Journal of Experimental Psychology: Human Perception and Performance, 20*, 3–16.

Turvey, M.T., & Carello, C. (1981). Cognition: The view from ecological realism. *Cognition, 10*, 313–321.

Uttal, W.R. (1990). On some two-way barriers between models and mechanisms. *Perception and Psychophysics, 48*, 188–203.

van Leeuwen, C., Steyvers, M., & Nooter, M. (submitted). *Stability and intermittency in high-dimensional dynamical models for perceptual grouping*. Manuscript submitted for publication.

Van Orden, G.C. (1987). A ROWS is a ROSE: Spelling, sound, and reading. *Memory and Cognition, 15*, 181–198.

Van Orden, G.C., Holden, J.G., Podgornik, M.N., & Aitchison, C.S., (submitted). *When a ROWS is not a ROSE: Null effects and the absence of mental structure.* Manuscript submitted for publication.

Van Orden, G.C., & Goldinger, S.D. (1994). Interdependence of form and function in cognitive systems explains perception of printed words. *Journal of Experimental Psychology: Human Perception and Performance, 20*, 1269–1291.

Van Orden, G.C., & Goldinger, S.D. (1996). Phonologic mediation in skilled and dyslexic reading. In C. Chase, G. Rosen, & G. Sherman (Eds.), *Developmental dyslexia: Neuro, cognitive, and genetic mechanisms* (pp. 185–223). Timonium, MA: York Press.

Van Orden, G.C., Johnston, J.C., & Hale, B.L. (1988). Word identification in reading proceeds from spelling to sound to meaning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*, 371–385.

Van Orden, G.C., Pennington, B.F., & Stone, G.O. (1990). Word identification in reading and the promise of subsymbolic psycholinguistics. *Psychological Review, 97*, 488–522.

Van Orden, G.C., Pennington, B.F., & Stone, G.O. (submitted). *What do double dissociations prove? Modularity yields a degenerating research program.* Manuscript submitted for publication.

Varela, F.J., Thompson, E., & Rosch, E. (1991). *The embodied mind.* Cambridge, MA: MIT Press.

von der Malsburg, C., & Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics, 54*, 29–40.

Warrington, E.K. (1975). The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology, 27*, 635–657.

Waters, G.S., & Seidenberg, M.S. (1985). Spelling-sound effects in reading: Time course and decision criteria. *Memory and Cognition, 13*, 557–572.

West, S.G., & Hepworth, J.T. (1991). Statistical issues in the study of temporal data: Daily experiences. *Journal of Personality, 59*, 609–662.

Ziegler, J.C., & Jacobs, A.M. (1995). Phonological information provides early sources of constraint in the processing of letter strings. *Journal of Memory and Language, 34*, 567–593.

Ziegler, J.C., Jacobs, A.M., & Stone, G.O. (in press). Statistical analysis of the bidirectional inconsistency of spelling and sound in French. *Behavioral Research Methods, Instruments, and Computers*.

Ziegler, J.C., Stone, G.O., & Jacobs, A.M. (in press). What's the pronunciation for OUGH and the spelling for /u/?: A database for computing feedforward and feedback inconsistency in English. *Behavioral Research Methods, Instruments and Computers.*

Ziegler, J.C., Van Orden, G.C., & Jacobs, A.M. (in press). Phonology can help or hurt the perception of print. *Journal of Experimental Psychology: Human Perception and Performance.*