CHAPTER 1

WHY NONLINEAR METHODS?

Claudia Carello

Center for the Ecological Study of Perception and Action
University of Connecticut
406 Babbidge Road
Storrs, CT 06269-1020
U. S. A.
E-mail: claudia.carello@uconn.edu

Miguel A. Moreno

Haskins Laboratories
300 George Street
New Haven, CT 06511-6695
U. S. A.
E-mail: miguel.moreno@uconn.edu

As behavioral scientists, we're in the business of understanding pieces of behavior. Everyone has his or her favorite types of behaviors, oftentimes things that simply derive from personal interest—we are fascinated by language or sport or animals and somehow contrive to make those into experimental topics. Apart from our idiosyncratic preferences we also bring our intellectual preferences, our assumptions about the kinds of explanations we expect to work. And although our theoretical positions are reasonably explicit, we also have metatheoretical positions that are typically somewhat hidden. Nonetheless, they sit implicitly behind what we do. Whether we are connectionists or computationalists or direct realists, the inherent philosophies of those positions dictate the kinds of problems we study and the kinds of variables by which we choose to define them. Our metatheories tell us what we think ought to be important.

But that's not the end of how we frame problems. In designing our studies, we still have a number of choices to make. Some of those choices are dictated by the requirements of the analyses that we'll use—repeated measures or factorial designs, randomized or blocked trials. Unlike our theoretical and metatheoretical positions, we tend to think of analyses as objective and benign with respect to intellectual assumptions. To be sure, all analyses assume criterial characteristics of the data that render the analysis in question legitimate. But we tend to think of those assumptions as mathematical. An important lesson of the chapters in this volume, however, is that our statistical analyses buy into intellectual assumptions as well. As you'll see, what we analyze and how we analyze it entails assumptions about the kinds of things that exist and assumptions about how those things can fit together. The chapters show us how we might begin to change the way we understand the very nature of those pieces of behavior that interest us. At the very least, these chapters show that analyses that acknowledge the dynamical nature of certain behaviors reveal a good deal of rich structure that cannot be extracted with more familiar analyses.

Consider Figure 1.1, which illustrates a few of the many different kinds of patterns of data that behavioral scientists encounter. Panel A shows distributions of the kind that we assume are typical of our experiments. Some manipulation increased the likelihood of larger responses, although, in this case, with greater variability among the responses. Consequently, the means of the two distributions are numerically different but the variances are such that there is considerable overlap between the two distributions. Conventional analyses allow us to assess the extent to which the variability seems to be systematic (i.e., due to the manipulation) or random (e.g., due to the vagaries of individual differences among people) in order to determine whether those means are different enough to be reliable.

The remaining panels show data of the kind considered in these chapters. Whether they fit this conventional characterization is an issue. Panel B, for example, shows two distributions that appear to be of the same general sort as panel A. Distribution 2 is a little more variable than Distribution 1, but in this case their peaks are in the same location. A closer look, however, suggests a subtle difference. The mean of Distribution 2 is larger than the mean of Distribution 1, and by the same amount as in panel A. But this time the increase is not due to a straightforward, overall addition. There appears to be a stretching of the high end of the distribution so that more large values get included

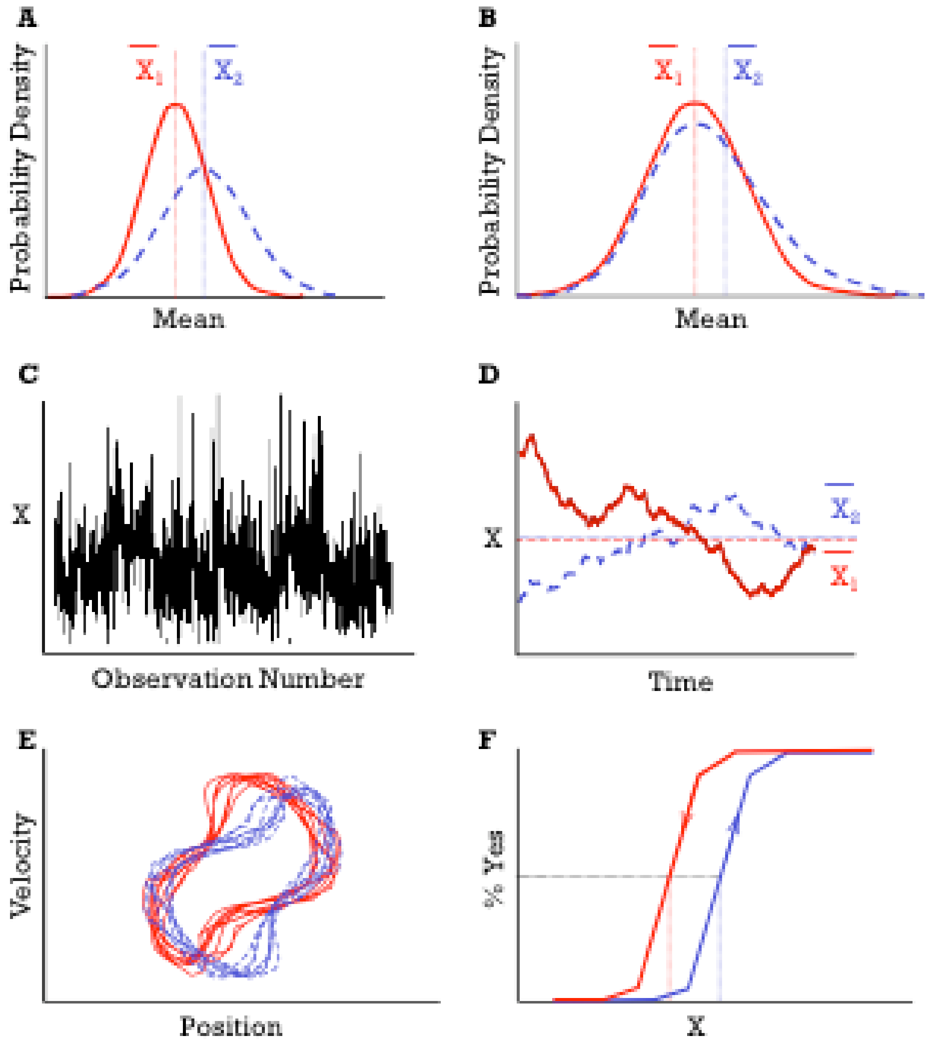

Figure 1.1. (A) Two normal distributions with different means and variances. (B) A normal distribution (solid line) and a distribution with a stretched "tail" (dashed line). (C) A time series. (D) Two time series with identical summary means. (E) Two time-ordered velocity × position profiles. (F) Categorical responses with different orders of presentation (indicated by the arrows).

in the calculation of the mean (Moreno, 2001). Is it appropriate to say that the means of these two distributions differ? The central tendency is at the same value. Eliminating values of the dependent variable that are larger than a certain cut-off might eliminate this tail but would that be an accurate depiction of the consequences of this particular experimental manipulation? Panels C-F show data for which the mean may be an even less appropriate measure. They are all time-series of one kind or another, depictions of individual responses being tracked over time or, at least, over order of presentation. In panel C you can see that the data are very noisy trial by trial. But there also appears to be some kind of large-scale wavy pattern overlaid on this noise. Summary statistics such as the mean and variance would be hard-pressed to capture this (see Chapter 6 by Holden). Panel D plots two time series together. The average position of the two series can be said to be the same if we simply add up the values and divide by the number of observations. But it is quite apparent that this single value is not an appropriate characterization of either time series or of the differences between them. Indeed, there is not really a single mean for either series; the mean of each changes over time, making it "illegal" to conduct conventional analyses. Panel E suggests that tracking the coincident changes in two variables might be informative. There is more to these data than a correlation could reveal. Whether small values of X_1 go along with large or small values of X_2 depends on when the observations occurred in the series. Finally, in panel F, two parallel functions are displaced from one another, not as a function of the X variable but as a function of whether that variable was encountered in an ascending series or a descending series (indicated by the arrows). There might be a temptation to average over the two presentation orders so as to identify "the" transition point, or to have ignored order altogether in a randomized presentation (but see Chapter 8 by Tuller).

In this opening chapter, we have two goals. First, we'll take one measure that is common in cognitive psychology research and use it to illustrate the kinds of intellectual assumptions that standard data analyses embrace. Second, we'll provide an overview of the issues that are treated in detail in the individual chapters.

A brief history of decomposing performance into components

Reaction (or response) time is the workhorse for exploring the nature of cognitive systems. Traditional approaches have tried to understand responses as the sum of component effects. While such approaches allow that the intrinsic dynamics of components may be complex, they severely restrict the kinds of interactions that can occur between components. In particular, it is common to assume that interactions between components must be linear. Traditional approaches have gambled that the effect of each cognitive component combines additively with the effects of other components, which together define the shape of response time distributions. This brief history tracks the payoff, so far, of this gamble.

Linear interactions mean that the effect of an unobservable component can be recovered in an overall measure like response time because each component effect spans a sub-interval of response time. The overall finishing time of the same component doing the same job will vary from occasion to occasion, however. Thus, the overall time course of all components would appear, to an experimenter, as a distribution of finishing times (like one of those shown in Figure 1.1A).

The enterprise of decomposing response time performance into component processes is an old one. In 1868, Donders proposed a *subtractive method* for identifying stages of information processing. The subtractive method was based on the idea that a stage could be inserted into (or deleted from) a sequence of stages. Donders hypothesized that a new stage would be added to accompany specific modifications to an experimental task. Comparing response times from two tasks could estimate the duration of the added stage.

The subtractive method was the preferred procedure for revealing mental stages for decades (Wundt, 1874; Cattell, 1886; Jastrow, 1890). It fell out of favor for several reasons. One criticism was that task modifications are more likely to alter the entire sequence of stages than to insert or delete individual stages (Külpe, 1895). Devising an experimental manipulation that unambiguously introduced a new processing stage proved to be the downfall of the subtractive method.

A more contemporary effort to identify component processes adopted Donders' assumption of additive finishing times plus assumptions geared to the asymmetrical shape of response time distributions. Empirical response time distributions typically have a hyperbolic shape with an elongated, slow tail, much like the dashed-line distribution in Figure 1.1B. The slow tails of response time distributions resemble exponential distributions (Christie & Luce, 1956; McGill, 1963) and the fast tails resemble the left half of Gaussian distributions. Christie and Luce hypothesized that empirical response time distributions are the intertwining of an exponential distribution and a base distribution of an unspecified form.

Hohle (1965) suggested that the form of the base distribution was Gaussian and indeed the convolution of an exponential and a Gaussian distribution can approximate very closely the shape of empirical response time distributions (Luce, 1986). Based on this idea, response time distributions are the sum of numerous component distributions with similar variances plus an exponentially distributed component with a much greater variance.

Hohle's assumptions include Donder's core assumption of additivity, that the interactions between components are linear. Consequently the shape of response time distributions should reduce to three parameters. Two parameters, μ and Σ, summarize the shape of an underlying Gaussian distribution. μ describes the location of the Gaussian distribution along the time axis and Σ describes the extent of the distribution's spread. A single parameter, τ, summarizes the location and spread of the exponential distribution.

Different component processes can be inferred if the parameter estimates systematically dissociate across experimental manipulations. Some manipulated factors should selectively influence one distribution (e.g., the exponential distribution) without affecting the other (e.g., the Gaussian distribution). This strategy for identifying component processes avoids one of the pitfalls of Donders' subtractive method, the requirement that two different tasks add (or delete) a stage of processing. Hohle's method requires that different conditions of the same task influence the exponential and Gaussian parameters independently.

As part of a model of response time performance, Hohle assumed that the exponential distribution is the effect of a response-choice process and the Gaussian distribution is the sum of all other processes (see also, Christie & Luce, 1956). The mapping of parameters to component processes was merely intuitive, however. For example, McGill (1963) had opposite intuitions. He assigned response-choice and other decision processes to the Gaussian sum of processes and suggested that the exponential distribution represents motor processes.

Unfortunately these *ex-Gaussian* strategies (combining exponential and Gaussian distributions) fared no better than Donder's subtractive method (Sternberg, 1969). Different factor manipulations did not systematically discriminate among different component parameters across experiments (Hohle, 1967; Gholson & Hohle, 1968a, 1968b). This outcome presents a problem because the ex-Gaussian hypothesis combines so many assumptions. When results are inconclusive, it is difficult, or impossible, to decide which assumption is false. Another problem, pointed out by Sternberg, is that combinations of many other distributions also approximate response time distributions (see also Van Zandt & Ratcliff, 1995).

Sternberg (1969) realized that the core assumption that cognitive systems are composed of successive stages could be isolated from supplemental assumptions such as specifying the form of component distributions. Sternberg stripped the assumptions regarding the nature of cognitive systems down to their core and asked, "How do component processes interact?"—the recurring question in this brief history. Sternberg proposed that if the component processes interact linearly then there must exist some factors that when manipulated will selectively influence different component distributions. If components interact linearly, then component distributions selectively influenced by separate factors will combine additively.

Sternberg's strategy of testing for linear interactions requires experimental manipulations of two or more factors. If the influence of one factor on overall performance is completely independent of the influence of another factor—a statistically additive interaction—then the two experimental factors relate to two different component processes. Alternatively, if the influence of one factor is modulated by another factor—a non-additive interaction—then both factors influence at least one common component process.

Assessments of additive interactions between component processes require estimates of component distributions that combine additively. Appropriate estimates of component finishing times, according to Sternberg (1969, p. 286), are arithmetic means. The mean of a sum of component distributions is the sum of component distribution means. Response time means, therefore, can be treated as the sum of component means.

Unfortunately, Sternberg's additive factors method has yet to identify any component process unequivocally. Additive interactions are the exception in cognitive studies. This situation could imply that the right set of factors has yet to be identified. The right set of factors could provide the necessary context to discover fundamental additive interactions. However, there is no guarantee that such a set of factors exists. Moreover, it may not be feasible to prove that such a set of factors does not exist. Sternberg provided an elegant and scientifically conservative test of the traditional assumption regarding the nature of cognitive systems. Unfortunately, the results have yet to answer Sternberg's question, "how do components interact?"

Following Sternberg's lead, in some sense, the chapters in this volume focus on the question, "how do components interact?" Nonlinear dynamical systems provide another way to explore this question. Nonlinear dynamical systems do not exclude the possibility of linear interactions; linear interactions are special circumstances within the range of possibilities of a dynamical system. Thus, modeling response times, among other things, as a dynamical system is a very general and conservative approach. There are fewer a priori assumptions regarding the components of the system. Furthermore, there are fewer restrictions on how components may interact.

THE CHAPTERS

The workshop focused on two types of analyses—recurrence quantification and fractals—that seem particularly fruitful for behavioral research. The general premises of these techniques are summarized in the chapters by Webber and Zbilut ("Recurrence Quantification Analysis of Nonlinear Dynamical Systems") and by Liebovitch and Shehadeh ("Introduction to Fractals") and each is followed by particular experimental implementations. Also included is an illustration of what can be gained by treating an established phenomenon dynamically from the start.

Recurrence Quantification Analysis

Much of the behavior of living systems is complex and seemingly non-predictable. Nonetheless, aspects of this behavior can be counted on to repeat. The bits that repeat may do so over long stretches, perhaps producing a pattern, or the recurrences can be quite short-lived. Consider an activity like a square dance. Much of a dancer's time is spent synchronized with the group in a large and obvious pattern, say, concentric circles alternately moving clockwise and counterclockwise. Only occasionally and briefly does one dancer get back together with his or her original partner. Both levels of recurrence—the circular patterns of the group and the momentary contact between partners—can be quantified and tell us something additional about the activity.

The levels of recurrence in the "RQA Dance"[1] as executed by other kinds of particles may not be as obvious, particularly the rare recurrences, but they are just as informative. And, it seems, the more complex the behavior the rarer and less obvious the recurrences, and the greater the need for ways to discover them. As noted by Webber and Zbilut, "the degree to which those systems exhibit recurrent patterns speaks volumes regarding their underlying dynamics." Even if we don't have a recurrent behavior as obvious as dancing partners holding hands, a system's underlying dynamics are accessible. Picking up on the theme that everything is connected to everything else, Takens (1981) introduced a theorem allowing a behavior space to be reconstructed from any measured variable. To be sure, a complex system is ultimately characterized by a number of participating variables. But these variables are necessarily coupled to one another and, therefore, each reflects the behavior of the system. In keeping

[1] The metaphor of RQA as describing a dance of particles was first illustrated by the duo of M. T. Turvey and Nobuhiro Furyama during a typically staid Turvey lecture at the University of Tokyo in May, 2004.

with our intuitions that behavior evolves over time, Takens' reconstruction is accomplished through time-delayed copies of some nominated variable. That is, some variable *x* is chosen as a preliminary index of the system's behavior and we track what happens to *x* over time. But we also want to know how *x* behaves relative to itself at later points in time, say, *t* plus a delay of δ or *t* plus a delay of 2δ So the original variable *x* becomes a dimension of the system in question and each time-delayed copy becomes another dimension of the system. Trajectories are traced through this multi-dimensional space and recurrences are measured: Do the trajectories come together at a point, do they travel together for a sequence of points, and so on? Each of these becomes an objective indication of some aspect of the system's dynamics. An advantage here is that the analysis allows you to characterize the dynamics of the system from the measurement of any variable, not necessarily a variable that seems like it ought to be right (what standard analyses refer to as *face validity)*.

This technique is illustrated in Shockley's chapter "Cross Recurrence Quantification of Interpersonal Postural Activity." He exploits RQA in a line of research aimed at quantifying the synchronization between two people who are engaged in a conversation. You can appreciate the challenge this behavior poses. What do you measure? The history of interpersonal synchrony is to treat it as a phenomenon of social coordination and to look for overt signs of that coordination. This means that the problem has been addressed fairly subjectively. For example, researchers might examine videotapes and look for signs of synchrony (e.g., similar gestures by a talker and a listener). As rigorously objective as researchers try to be, they must still interject themselves into the process of identifying an occasion of synchrony. Shockley and his colleagues have instead chosen a behavior—the postural sway of the participants—that is not an overt part of the act of conversing and used RQA to sift through the trajectories and extract the recurrent patterns. This is unlikely to be a behavior that people are controlling consciously. It changes "for free," pushed around by the fact that our heads are pretty massive sitting up on top of the relatively skinny sticks that are our bodies. So when we breathe and talk and gesture, those big heads move around, moving the body's center of mass, CM, along with them. The trajectory of the CM is tracked over time, time-delayed copies of its trajectory can be generated, and you're on your way to generating a behavior space. The subtle measures of RQA allow Shockley and colleagues to manipulate the constraints on just how coordinated the joint behavior is. They have people talk to each other or to someone else, at the same time or in the course of taking turns, using words that differ in their similarity, and so on—in order to uncover influences on the degree of coordination. This means that motor behavior, a level of behavior that some might like to relegate to the bin of basic behaviors that we can take for granted, can be used as an index of language, something we take to be one of our fanciest behaviors.

In the chapter on related work by Pellecchia and Shockley "Application of Recurrence Quantification Analysis: Influence of Cognitive Activity on Postural Fluctuations," RQA is applied to a single postural trajectory. This time, your big head is moved around not by coordinating with another person but simply by standing while directing some of your attention to another task. Here the emphasis is not so much on the objectivity to be gained—postural sway indexed by the excursions of the center of pressure (COP) is a common measure in some domains. But Pellecchia and Shockley point out that more traditional postural measures are summary measures, for example, of COP path magnitude and variability. But these summary measures turn out to be insufficiently sensitive to the varied ways in which cognitive attentional load can influence certain aspects of postural stability but not others. They "do not reflect the dynamical properties of postural control." In particular, they miss the temporal structure of a COP time series. One challenge is that posture data are what is called *nonstationary*. No single summary measure adequately captures them because the mean changes over time and the variability changes over time (see Figure 1.1C). The notion of an average postural location doesn't make sense. But this nonstationarity also makes posture data inappropriate for analyses that assume stationarity. RQA, in contrast, makes no such assumptions about the way the data are distributed.

Pellecchia and Shockley deal with some of the technical aspects of carrying out RQA. When you're evaluating the RQA dance, what qualifies as an instance of the two original partners having come together? Do they actually have to touch hands or can they simply slap hands or wave in the vicinity of each other? When you're generating your behavior space through time-delayed copies, what should the delay be? How many dimensions define your hyper-space? Pedagogically, our preferences is to illustrate it with three because we can visualize three-space. But there are no such mundane constraints on RQA. Pellecchia and Shockley suggest a kind of exploratory strategy in which the RQA quantities are calculated for a range of parameter values (numbers of dimensions, size of time-delay), and then settle on a value of the latter from a range that doesn't have dramatic consequences for the RQA measures. The upshot in the research they describe is that RQA promotes a different understanding of what attention does. In the particular experiment they described, for example, summary measures of COP were all affected in the same way by attentional demands, suggesting that attention causes a decrement in postural control. The RQA measures, in contrast, yielded differences that suggest a nuanced understanding of the ways in which postural components (e.g., the front to back vs. side to side movements) might be modulated by a stander in order to meet attentional demands. Variability is something to be harnessed by the system to achieve a goal.

Fractal Analyses

Most of the things that we need to measure tend to be irregular. This is no less true when the things we measure are behaviors rather than objects. Geometrically, this means that things are more like coastlines than rectangles. As Liebovitch and Shehadeh point out in "Introduction to Fractals," this not only makes them hard to measure, it makes the ruler that we use important. Quite surprisingly, the measured size of the coastline depends on the size of the ruler. Smaller rulers get into more of the nooks and crannies, thereby including more stretches of coastline than would be the case with a large ruler that is forced to bridge those gaps. So the level of resolution that we choose to achieve in our measurement affects the values we get. Consider one classic value that is typically used to characterize the behavior of a system, its mean. The mean is a measure derived from a collection of property or performance values that all tend to distribute around this more or less central value. It is considered a value that typifies the thing being measured. But with irregular objects like coastlines, there really is no typical value. The mean depends on the resolution of the measurement and isn't all that meaningful.

Subjectivity in measurement is certainly a problem to be reckoned with. You can appreciate how standard assumptions about normal distributions (most notably, more samples should lead to more precision, not more stuff) will be inappropriate in such circumstances. But so-called fractal objects, which are defined, in part, by this dependency on measurement resolution, have an additional property that makes them especially interesting. Fractal objects are *self-similar*, that is, structure at the large scale (or the structure of behavior at the large scale) is duplicated at the small scale: "…the statistics of the small pieces are similar to the statistics of the large pieces." This is so whether we are talking about the structure at different scales of space or at different scales of time. What happens within a square centimeter mimics what happens within a square meter; what happens within a one second window mimics what happens within a one minute window. Self-similarity can reveal much about the dynamics of a system. Depending on how the variability relates to the size of the window—a relationship indexed by what is called the Hurst exponent—it tells you whether an increase in the value of a measure taken now is likely to be followed by an increase or a decrease in that measure taken later. In essence, continuous dynamical processes manifest a kind of memory without the logical attributions and storage metaphors we gravitate to in the behavioral sciences.

A respect for unfolding dynamics encourages the treatment of order effects as entities of interest in an experiment rather than as sources of contamination. This is illustrated by Holden in "Gauging the Fractal Dimension of Cognitive Performance." He uses a simple repetitive time-estimation task to demonstrate the options for analyzing a time series of response times. A participant attempted to produce a sequence of equal intervals to mimic a presented target interval. Here our interest is not so much in accuracy as in the kind of process that produced the performance that was obtained. The way in which the variability changes over time is the source of the insight. Holden provides a nice intuitive metaphor here. If the interval estimates did not vary, then a graph (of value over time) would yield a one-dimensional straight line. To the extent that the time series is messy, it more closely resembles a two-dimensional plane. The *fractal dimension* of the time series can be calculated from the time series, with a value of 1.5 indicating randomness and 1.2 indicating pink noise, a time-scale dependent variability. It is, of course, more complicated than this, since order matters. Shuffling the data yields white noise because you've destroyed the fractal structure, the pattern of variability at short time scales that is echoed at ever-longer time scales.

Holden's chapter also contains several caveats for conducting these kinds of nonlinear analyses. Caveats are necessary because there are options and we are sensitive to the fact that options often entail assumptions. Let's consider a few that routinely arise due to the practical finiteness of data collection. Ordinary statistics tell us that outliers are a problem and should be eliminated. Sympathy for dynamics burdens us with the knowledge that outliers are not necessarily produced by an aberrant procedural hiccup. We have to assess the extent to which they would dominate the analysis. You also have to be careful that the analysis is not dominated by spurious trends (e.g., a stretch of linearity) in the finite data set that might be part of something else if the data collection had continued. So the longer the time series the better in order to see the cascading structure of varied time scales. But since our interest is in emergence over time, we cannot pretend that data collected over several days are the same as data collected in one sitting. Finally, with all of these options, you're bound to come up with varied characterizations depending on which choices you made. So you need to conduct more than one analysis type (e.g., spectral and dispersion analyses) as converging operations.

In her chapter "1/f Dynamic in Complex Visual Search: Evidence for Self-Organized Criticality in Human Perception" Aks applies this perspective to visual search behavior, the eye movements that people engage in when looking for a specific small detail amidst a clutter of distracting detail. These movements appear quite haphazard. Our eyes dart back and forth, up and around, in a mix of short jumps and long, often landing on the same places time and again. We do not follow a systematic path, say, from upper left to lower right that would seemingly guarantee that the target would be encountered. Our behavior doesn't show in any obvious way that we remember where we looked in vain before. Yet the repetitive, jerky movements are surprisingly effective—we find our friend in the crowd; we select the perfect bolt from a stash of culch.[2] Indeed, the ordinary conditions under which visual search happens don't really favor the tidy, thorough search. We don't have all the time in the world; our targets or our goals are often on the move.

At issue in this research domain is what guides visual search. If eye movements truly were random, we would have to conclude that they were driven by something other than what the system had done before, that there was no real role for memory. But it is in the superficially random noisiness that a fractal analysis uncovers subtle structure. By quantifying how the noise changes over time we gain insight into the kind of system that has produced that noise. Presenting subjects with a difficult search task (looking for a target with a conjunction of features, not just one) reveals that a bout of scanning has its own internal history. What we do early in the bout does indeed influence what we do later in that same bout. But this dynamic history is quite different from more standard characterizations of memory that entail a certain degree of address-specific tagging. It implicates more subtle contingencies, indexed by what is called 1/f behavior.

The final chapter by Tuller does not incorporate either recurrence or fractal analyses. Instead, she illustrates the advantage of a general dynamical attitude in designing an experiment, with subsequent opportunities for new interpretations of seemingly well-understood results. The generality is especially apparent in that this is strictly a perception experiment. The relevant series is not of the *timing*

[2] In the New England vernacular, culch refers to items that may (or may not) come in handy someday.

of the responses but of their *content*—why should a speech token sound like one syllable versus another? Tuller tackles the classic phenomenon of categorical perception. Syllables are synthesized to vary incrementally on some acoustic property. Even though each syllable token is defined by a different value taken from a wide range of the acoustic property, each is heard as either one syllable or the other; they belong to one syllable category or the other. Categorical perception has historically been treated statically: You hear a token and its maps better onto one representation than another. But Tuller shows that thinking of the process as dynamical—with stable, attractive states that change abruptly—focuses the experimenter on finding what encourages those nonlinear shifts from one stability to another.

Tuller's major caveat is that some standard methodological choices, most notably, randomizing the order in which stimuli are presented, obscure the dynamics of a system. In a now-familiar refrain for this volume, she notes that far from being a nuisance that has to be controlled, order effects allow a system's dynamical signature to emerge. A dynamical perspective requires that the stimuli be presented in order, for example, alternating increasing and decreasing levels of the acoustic property. This allows an interpretation of the acoustic property as a *control parameter* rather than as a cue. Now the categorical shift can be mined. Does it happen at the same value of the control parameter in one order as the other (see Figure 1.1F)? Does the syllable you first hear persist? Or do you switch to the other syllable at a low level on the way up but a high level on the way down? Tuller notes that all three of these patterns have been observed and one or the other can be encouraged. Most notably, the abrupt change from one syllable to another, from one stable state to another, comes about during situations of instability. What had been unanimity of responses becomes somewhat mixed. Such instability links this perceptual phenomenon to the general phenomenon of self-organized pattern formation. This linkage allows one to write the differential equations that "define systems with attractor properties that fit the observed experimental data." That is to say, the linkage is far from metaphorical. A perception system is modeled in the same way, with the same ontological status, as an action system, clearly a different direction from the traditional treatment of categorical perception.

Conclusions

Behavioral scientists study the actions of humans and animals. Some of these actions make sense at the level of the individual and some emerge only in a social setting. As we noted at the outset, we make a number of choices in conducting our studies. The practical issue of what kind of equipment we have, how many participants are available, and so on, are supplemented by the kinds of analyses we know how to do and the kinds of data we collected to put into those analyses. We have suggested that such choices are not necessarily as benign and objective as we would like to believe. Do we try to avoid data like those shown in panels B-F of Figure 1.1 because they are messy? Or do we try to contrive our manipulations to produce those kinds of data because they allow the richness of dynamical systems to be seen? The following chapters are of a mind that we should not be afraid of variability. It may well be the driving force of nature.

AUTHOR NOTE

Preparation of this work was supported by grants from the National Science Foundation (SBR 00-04097) and the National Institutes of Health (HD-01994).

REFERENCES

Cattell, J. McK. (1886). The time taken up by the cerebral operations. *Mind*, *11*, 220-242.

Christie, L S., & Luce, R. D. (1956). Decision structure and time relations in simple choice behavior. *Bulletin of Mathematical Biophysics, 18*, 89-112.

Cramer, H. (1946). *Mathematical methods of statistics*. Princeton: Princeton University Press.

Donders, F. C. (1868). Over de snelheid van psychische processen. Onderzoekingen gedaan in het Physiologisch Laboratorium der Utrechtsche Hoogenschool, 1968-1869, Tweede reeks, II, 92-120. Transl. By W. G. Koster, *Acta Psychologica, 30*, 412-431.

Gholson, B., & Hohle, R. H. (1968a). Verbal reaction time in hues and hue names and forms vs. form names. *Perception & Psychophysics, 3*, 191-196.

Gholson, B., & Hohle, R. H. (1968b). Choice reaction time to hues printed in conflicting hue names and nonsense words. *Journal of Experimental Psychology, 76,* 413-418.

Hohle, R. H. (1965). Inferred components of reaction time as functions of foreperiod duration. *Journal of Experimental Psychology, 69,* 382-386.

Hohle, R. H. (1967). Component process latencies in reaction times of children and adults. In L.P. Lipsitt & C.C. Spiker (Eds.), *Advances in child development and behavior, Vol III*. New York: Academic Press.

Jastrow, J. (1890). *The time-relations of mental phenomena*. New York: N.D.C. Hodges.

Külpe, O. (1895). *Outlines of psychology*. New York: MacMillan, Secs. 69, 70.

Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.

McGill, W. J. (1963). Stochastic latency mechanisms. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology, Vol I*. New York: Wiley.

Moreno, M. A. (2001) A nonlinear dynamical systems perspective on response time distributions. Unpublished doctoral dissertation, Arizona State University, Tempe, AZ.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica, 30*, 276-315.

Van Zandt, T., & Radcliff, R. (1995). Statistical mimicking of reaction time data: Single-process models, parameter variability, and mixtures. *Psychonomic Bulletin & Review, 2, 20-54.*

Wundt, W. (1874). *Grundzuge der Physiologischen Psychologie*. Leipzig: Engelmann.