Journal of Contextual Behavioral Science 3 (2014) 279–290

Contents lists available at ScienceDirect

Journal of Contextual Behavioral Science

journal homepage: www.elsevier.com/locate/jcbs

On the orderliness of behavioral variability: Insights from generativity theory

Robert Epstein [a,b]

[a] School of Social Sciences, University of the South Pacific, Laucala Bay, Suva, Fiji
[b] American Institute for Behavioral Research and Technology, 1035 East Vista Way, Vista, CA, United States

ARTICLE INFO

Article history:
Received 10 December 2013
Received in revised form
23 June 2014
Accepted 19 August 2014

Keywords:
Generativity theory
Creativity
Transformation functions
Automatic chaining
Resurgence
Frequency profile

ABSTRACT

Over time, many natural phenomena that had long appeared to be disorderly have been found to be orderly and predictable under specifiable conditions. First introduced in the early 1980s, generativity theory is a formal, predictive theory of the behavior of organisms that reveals the orderliness, moment to moment in time, in apparently disorderly behavior – even the surprising behavior a community sometimes calls "creative." According to this theory, under two specific conditions – when behavior is ineffective or when stimuli present in the environment are novel, compound, or ambiguous – novel behavior emerges in a predictable way as a result of a dynamic process in which multiple behavioral processes operate simultaneously on the probabilities of multiple behaviors. The process can be represented by a series of equations called transformation functions. Instantiated in a computer model, the equations have proved useful in the moment-to-moment prediction of the emergence of novel behavior in both pigeons and people. A graphical method that generates a "frequency profile" has also helped to reveal the orderliness in the apparently disorderly behavior of individuals. Generativity theory makes no assumptions about the existence or nature of cognitive mechanisms and does not depend on the statistical analysis of aggregated data to show the orderliness in complex behavior. Although its predictive power in the laboratory is perhaps unparalleled, the full potential of generativity theory has yet to be explored.

1. Introduction

Charles Darwin's ability to detect certain patterns in the apparently disorderly physical characteristics of 26 types of birds in the Galápagos Islands helped him to formulate one of the most profoundly important theories in the history of science, the theory of evolution (Desmond & Moore, 1991). He deduced that the orderly variability in phenotypes produced by sexual reproduction in each generation of a given species, in combination with the selective survival requirements exerted by different environments, could, over time, account for the creation of new species. He knew nothing about genes, meiosis, gametes, or syngamy – about the biological mechanisms underlying the phenomena he observed – but he inferred that such mechanisms must exist.

Although the details vary from one scientific domain to another, this is generally what science is all about: finding the orderliness in apparently disorderly phenomena, often at just one level of observation. Ideally, that orderliness is eventually expressed in formal terms so the theory can be used to make specific predictions; this increases the utility of the theory and also helps to establish its validity. Einstein's general theory of relativity is an excellent case in point. First published as a set of field equations (Einstein, 1915), in the years and decades that followed, the mathematical form of the theory proved to be successful in making accurate predictions about gravitational time dilation, the bending of starlight around the sun, and other natural phenomena (Einstein, 1915). Just recently, the theory was shown to predict with remarkable accuracy the redshift of light that occurs in massive galactic clusters (Wojtak, Hansen, & Hjorth, 2011).

The behavior of organisms is one of those natural phenomena that often appears to be disorderly and unpredictable. Although it is true that hundreds of millions of human drivers manage to stay in the correct lanes on roadways every day, when human behavior is not constrained by salient stimuli (curbs, signs, and lines on the road), the consequences of previous actions (traffic tickets and warnings), or rules and laws (descriptions of how one must drive to avoid future tickets), it often appears to be quite disorderly. In new or challenging situations especially, it can be difficult to predict what people will do, think, or say – even for *them* to predict what they will do, think, or say. People also do genuinely new things; virtually every sentence that we speak or write is new in some respects, our dreams are sometimes extraordinary, and occasionally an individual will do something so new and interesting that a community will label the action or its product "creative."

E-mail address: re@aibrt.org

http://dx.doi.org/10.1016/j.jcbs.2014.08.004
2212-1447/

Although the language of creativity is applied erratically by a community (depending on what it happens to consider interesting at the time), there is also no question that occasionally someone behaves in ways that are profoundly new, producing the kind of output Csikszentmihalyi calls "Big C" creativity (e.g., Csikszentmihalyi, 1996; Csikszentmihalyi & Epstein, 1999).

We face a range of predictability challenges here, which can be said to exist on a continuum from fairly easy (predicting which turns someone will take on the way home from work) to quite difficult (predicting which turns someone will take when he or she is lost in a new city) to probably impossible (predicting the specifics of an amazing new mobile computer application someone devises while being lost in a new city, which will someday help millions of people to avoid getting lost).

In the early 1980s, inspired by observations I made over the course of conducting a series of somewhat irreverent experiments with pigeons, I formulated a theory of how behavior is generated moment to moment in time, possibly in a wide range of higher organisms (Epstein, 1985b, 1991, 1996a, 1999). The theory – generativity theory – can be expressed as a series of equations called "transformation functions." Instantiated in a computer model, the equations have proved useful in predicting fairly complex behavior moment to moment in time in the laboratory. Over the years, I have become increasingly confident that generativity theory, or at least something like it, can help us understand how behavior is generated across that entire range of difficulty – from the reappearance of an old, well-established behavior to the occurrence of profoundly new behavior. Along the way, I also developed a type of graph called a "frequency profile" which reveals the orderliness in certain types of performances that appear through direct observation to be disorderly in nature.

Before I present the basics of generativity theory, I will attempt to put the theory into a broader context of some contemporary thinking about the orderliness of behavior.

2. Some models, ideas, and approaches

2.1. How random is behavior?

In recent years, a growing number of experts have come to view some aspects of behavior to be truly random in nature – unpredictable, by definition. In a review of the relevant literature as of 2005, Paul Glimcher identified a number of researchers who not only have concluded that "indeterminacy" is a basic feature of both human animal behavior but even that evolutionary forces may have favored organisms who can behave unpredictably under certain conditions. An animal is less likely to killed, certainly, if a pursuing predator cannot easily anticipate its next move. Entire classes of individual behavior have been studied which, according to Glimcher, are "as fully random as can be measured" (Glimcher, 2005, p. 28). Accordingly, neuroscientists, he says, are uncovering evidence for the existence of "apparently indeterminate processes within the architecture of the mammalian brain" (p. 28). In a broad overview of literature that overlaps fairly little with the studies cited by Glimcher, Allen Neuringer (2002) draws similar conclusions: that behavioral variability is a "stochastic process" (p. 697) and that "organisms have evolved to behave unpredictably" (p. 701).

But there is a problem here. These conclusions, as well as those of many other investigations of behavioral variability (e.g., Emonet & Cluzel, 2008; Hopkinson & Neuringer, 2003; Johansen, Killeen, & Sagvolden, 2007; Machado, 1997; McIntosh, Kovacevic, & Itier, 2008; Reichert, 1978; Shimp, Froehlich, & Herbranson, 2007; Tatham, Wanchisen, & Hineline, 1993), are based on aggregated data – aggregated over trials with individual subjects or, more commonly, across organisms. Do aggregated data tell you anything meaningful about the orderliness of the behavior of organisms?

If, over time, you keep track of whether I drink water or Diet Pepsi with my meals (which, in truth, are virtually the only liquids I ever drink from a glass or bottle, even with my breakfast), you will likely conclude that there is a 50/50 chance that I will drink one or the other with my next meal. As you continue to keep count over time, you will become increasingly confident about this prediction; you might even conjecture that a neural random number generator – the proverbial "mental coin toss" – governs my choice. But if, *on a single occasion*, you observe that I enter the kitchen, then prepare a peanut-butter-and-jelly sandwich – which you know from previous observations that I consume only with water – then open the refrigerator door where a water bottle and a Diet Pepsi bottle stand side by side, will you have any trouble predicting my next move?

Take this a step further. Say you begin to keep careful records of the specific foods I consume with water versus the specific foods I consume with Diet Pepsi, along with specific stimulus conditions and behaviors that reliably precede the consumption of each type of drink. Over time, wouldn't not you be able to predict my choice of drinks *on any single occasion* with increasing accuracy, perhaps with nearly 100% accuracy?

Although I am a heavy user of statistics in most of my current research projects (because I'm working with data sets obtained from thousands of people – see, for example, Epstein, McKinney, Fox, & Garcia, 2012), when it comes to understanding the moment-to-moment behavior of a single organism, the statistical analysis of aggregated data can be misleading (cf. Barlow & Nock, 2009; Sidman, 1960; Skinner, 1966, 1976). Is it even meaningful to say, based on aggregated data, that there is a 50/50 chance that I will select one drink or the other when, with the right data in hand, it would be easy to make an accurate prediction about my behavior on any single occasion? Considered in this context, the statistical approach to understanding ongoing behavior in the natural environment may be of questionable value.

Yet this approach is quite common. Evolutionary biologist John Maynard Smith's (1982) "hawk-dove model" is a case in point. Applying concepts from game theory, Maynard Smith shows that under certain conditions – specifically when the value of a territory is high (implying an animal should protect it aggressively) and the cost of an injury is also high (implying that an animal should retreat to avoid being hurt), the only sensible strategy for that animal is to be aggressive on some occasions (a hawk) and passive on others (a dove), assuming one role or the other unpredictably from one occasion to the next. Survival is enhanced by this strategy; the math is clear.

But, again, by observing a particular animal for a long period of time, wouldn't it be fairly easy to predict which role it will assume on a single occasion? Its behavior on a particular occasion will be the net result of its recent environmental history and the particular stimulus conditions it faces. As the attacker grows near, the appearance, odor, movements, and sounds of that attacker will, from one moment to the next, make all the difference in your prediction. Imagine, in fact, that you could speed up your perception so that the scene will appear to unfold in slow motion, giving you time to analyze all aspects of what is occurring. Fifteen minutes before the attack, could you make a reasonably good prediction? One minute before the attack, could you make a better one? One second before the attack, is there any doubt that you could make an accurate prediction about which role the animal will assume? What's more, by altering the variables of which you know the behavior is a function, couldn't you perhaps guarantee that an animal will *always* behave as a hawk or *always* behave as a dove?

Let us take this idea even further. If, in a particular setting, from one occasion to the next you are highly adept at predicting