Cognitive Psychology 103 (2018) 85–109

Contents lists available at ScienceDirect

Cognitive Psychology

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

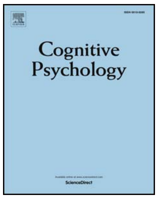

Subjective randomness as statistical inference

Thomas L. Griffiths[a,*], Dylan Daniels[b], Joseph L. Austerweil[c], Joshua B. Tenenbaum[d]

[a] *Department of Psychology, University of California, Berkeley, United States*
[b] *Department of Statistics, University of California, Berkeley, United States*
[c] *Department of Psychology, University of Wisconsin-Madison, United States*
[d] *Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, United States*

ARTICLE INFO

Keywords:
Randomness
Bayesian inference
Algorithmic complexity

ABSTRACT

Some events seem more random than others. For example, when tossing a coin, a sequence of eight heads in a row does not seem very random. Where do these intuitions about randomness come from? We argue that subjective randomness can be understood as the result of a statistical inference assessing the evidence that an event provides for having been produced by a random generating process. We show how this account provides a link to previous work relating randomness to algorithmic complexity, in which random events are those that cannot be described by short computer programs. Algorithmic complexity is both incomputable and too general to capture the regularities that people can recognize, but viewing randomness as statistical inference provides two paths to addressing these problems: considering regularities generated by simpler computing machines, and restricting the set of probability distributions that characterize regularity. Building on previous work exploring these different routes to a more restricted notion of randomness, we define strong quantitative models of human randomness judgments that apply not just to binary sequences – which have been the focus of much of the previous work on subjective randomness – but also to binary matrices and spatial clustering.

1. Introduction

Imagine you are driving and approach a stoplight. You see eight cars in front of you, all of which are black, and begin to wonder why all of these vehicles are the same color. At the next stoplight you see eight cars all of varying colors and perceive the situation as happenstance, thinking nothing of it. These two events strike us differently because of our intuitions about randomness. The second event seems clearly a result of chance, but the first event seems to suggest another explanation. But where do these intuitions about randomness come from?

One naïve explanation purports that non-random events are simply low-probability events, but multiple studies have confirmed that humans do not judge equally likely events as equally random (Kahneman & Tversky, 1972; Tversky & Kahneman, 1974). One of the classic examples concerns flips of an fair coin: Asked to choose which of the following coin flip sequences of length eight is more likely to occur, `HHHHHHHH` or `HTHTTHTT`, most people will choose the latter, despite the fact that each sequence has the same probability of occurring, $\left(\frac{1}{2}\right)^8$. This intuition is surprisingly strong. Indeed, even trained statisticians would likely be surprised to see a coin turn up heads eight times in a row. Despite knowing that the two coin flip sequences are both equally probable, we still question whether we are actually witnessing random flips of fair coin; it takes

---

[*] Corresponding author.
E-mail address: tom_griffiths@berkeley.edu (T.L. Griffiths).

https://doi.org/10.1016/j.cogpsych.2018.02.003
Accepted 24 February 2018

cognitive effort to believe the coin is truly unbiased.[1]

For this reason, randomness has emerged as a central and persistent topic in the cognitive sciences. Many studies (e.g., Falk & Konold, 1997; Lopes & Oden, 1987) suggest that what people usually mean by randomness is the absence of certain detectable patterns. Because this notion is at odds with ideas of randomness used in probability and statistics, people seem to reason poorly about chance. As Nickerson (2002) puts it, "The general conclusion that the results of these experiments in the aggregate seem to support is that people are not very good at these tasks – that they find it hard to generate random sets on request and to distinguish between those that have been produced by random processes and those that have not" (p. 71). Bar-Hillel and Wagenaar (1991) conclude "People either acquire an erroneous concept of randomness, or fail to unlearn it" (p. 448).

But are people really bad at reasoning about chance? Or are they solving a different problem from merely judging the probability of different outcomes under a random process? If subjective randomness is about detecting patterns, then it is intimately related to intelligent action: "random" stimuli provide no help in predicting outcomes (e.g., what information do eight cars of all different colors confer?), whereas "non-random" stimuli aid our thinking (e.g., if you see eight black cars in a row you might infer that a foreign dignitary is visiting your town). Consistent with this idea, several papers have connected subjective randomness to formal frameworks for characterizing the amount of structure in a stimulus, such as algorithmic complexity (Falk and Konold, 1997; Feldman, 2004; Gauvrit et al., 2014, 2016; Griffiths and Tenenbaum, 2003, 2004) and Bayesian inference (Griffiths & Tenenbaum, 2001; Hsu, Griffiths, & Schreiber, 2010; Williams & Griffiths, 2013).

In this paper, we present a formal framework that unifies these previous approaches, demonstrating that subjective randomness can be explained as a form of statistical inference about the process that generated a stimulus. The key challenge for this approach is characterizing the kinds of structure that people might identify in a stimulus – a problem that algorithmic complexity theory solves by considering all regularities generated by short computer programs (e.g., Li & Vitányi, 2008). This notion is too general to capture human subjective randomness judgments, but by reformulating algorithmic complexity as a statistical inference we identify two ways in which it can be adapted: simplifying the kinds of computing machines considered and restricting the set of possible regularities. We use these two approaches to develop models of human randomness judgments for three different kinds of stimuli: sequences of coin flips, binary matrices, and dot patterns in spatial arrays.

2. Phenomena and theories of subjective randomness

Much of the research regarding humans' perceptions of randomness has concerned sequences of numbers or binary outcomes. Reichenbach (1934/1949) is credited with having made the original suggestion that mathematical novices will be unable to produce random sequences of numbers, instead showing a tendency to overestimate the frequency with which numbers alternate. Subsequent research has provided strong support for this claim (reviewed in Bar-Hillel & Wagenaar, 1991; Tune, 1964; Wagenaar, 1972), with both sequences of numbers (Budescu, 1987; Kareev, 1995a; Rabinowitz, Dunlap, Grant, & Campione, 1989) and two-dimensional black and white grids (Falk, 1981). Most strikingly, people believe that subsequent coinflips are more likely to alternate than to stay the same. Basic probability tells us that we should expect the probability of alternation from `H` to `T` (and vice versa) to be about 0.5, yet Falk and Konold (1997) provide a vast body of evidence which shows that people perceive sequences with a probability of alternation around 0.6 or 0.7 as most random. In addition to alternation, people are sensitive to symmetries (e.g., `HHHHTTTT`) and duplications (e.g., `HHTTHHTT`) within sequences, and deem sequences with these properties less likely to have been created by random processes (see, e.g., Hahn & Warren, 2009; Lopes & Oden, 1987).

A number of theories have been presented in an attempt to account for the accuracy of Reichenbach's conjecture. One of the earliest such theories can be traced to Skinner (1942), involving the suggestion that people develop a concept of "randomness" that differs from the true definition of the term. More recently, a theory of this kind was endorsed by Falk (1981). The claim that individuals hold a concept of randomness is supported by the finding that people consistently apply the same criteria in their judgments of randomness (Budescu, 1987; Falk, 1981). However, some studies have failed to find this kind of consistency (e.g. Wiegersma (1982)), and the difficulty of finding a useful definition of randomness to which this subjective concept might be compared is viewed as a deficiency of the theory (Lopes, 1982).

Other explanations have focused on the limitations that the human information processing system places upon the generation of random numbers. Baddeley (1966) and Wiegersma (1982) suggested that restrictions in short term memory span might account for the high levels of alternation present in generated random sequences. This claim is attractive, but fails to provide an account of why similar effects are obtained in grids and other stimulus displays (Falk, 1981). Furthermore, Neuringer (1986) showed that biases in the judgment of randomness could be removed by educating participants about the statistical properties of random sequences. This result suggests that memory alone may not provide a full explanation of the observed phenomena.

One account that has had a strong influence upon the wider literature of cognitive psychology is Kahneman and Tversky's (1972) suggestion that subjects may be attempting to produce sequences that are "representative" of the output of the generating process. For sequences, representativeness means that the number of elements of each type appearing in the sequence will correspond directly to the overall probability with which these elements occur. Kahneman and Tversky suggested that random sequences are also subject

---

[1] A statistically-minded reader might be concerned that getting a sequence of the kind "eight heads in a row" *is* less probable than getting a sequence of the kind "three heads and five tails." However, randomness isn't just about the probability of different kinds of events. Griffiths and Tenenbaum (2007) discuss this issue in detail, and provide counter-examples. For instance, `HHHHTHTTHHHTHTHHTHTTHHH` is an instance of the kind "fifteen heads, eight tails" which is a less probable outcome of tossing a fair coin 23 times than `HHHH` – "four heads in a row" – is for tossing a fair coin 4 times, even though the former would presumably be considered more random.