Poetics xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Poetics

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

Culture and computation: Steps to a Probably Approximately Correct theory of culture

Jacob G. Foster

Department of Sociology, UCLA, 264 Haines Hall, 375 Portola Plaza, Los Angeles, CA 90095, United States

ARTICLE INFO

Keywords:
Culture
Cognition
Computation
Theory
Science

ABSTRACT

This paper outlines some provisional steps toward a theory of culture grounded in computational thinking. I begin by describing computational thinking, drawing on Marr's hierarchy for the analysis of information processing systems. I then address the definition of culture, arguing that culture is a property of causal chains, rather than a thing-in-the-world. I briefly address contemporary debates over the nature of culture—embodied versus embedded—and argue for an ecological approach in which culture-in-action unfolds as embodied schemas recognize (and produce) "handles" in the environment. When schemas are "objectively adapted" to the handles, they generate action that is ecologically rational. To explain ecologically rational culture-in-action, I outline a formal approach to cultural learning based on Probably Approximately Correct (PAC)-learning theory. I illustrate my approach throughout with examples drawn from the sociology of science.

1. Introduction

Culture has become a central concept in contemporary sociology (Patterson, 2014). Since the 1990′s (DiMaggio, 1997), sociologists have sought to clarify this key explanatory resource by connecting conceptions of culture in contemporary social theory (Bourdieu, 1990, 2000; Sewell, 1992; Swidler, 1986) to recent work in cognitive science (Brubaker, Loveman, & Stamatov, 2004; Lizardo, 2004, 2009, 2017, 2014; Lizardo & Strand, 2010; Martin, 2010, 2011; Vaisey, 2008). The targets of interdisciplinary synthesis have ranged from concrete, middle-range theories—schemas (DiMaggio, 1997; Rumelhart et al., 1980), dual-process models of "thinking fast and slow" (Evans, 2008; Kahneman, 2011), or the taxonomy of declarative and non-declarative memory (Lizardo, 2017)—to more comprehensive accounts of cognition as embodied (Clark, 1997; Slingerland, 2008), situated (Goodwin, 2000; Hutchins, 1995), distributed (Hutchins & Klausen, 1996), or all of the above (Clark, 2008; Goodwin, 2017). Debates over the cognitive foundations of culture matter. Explanations positing unrealistic cognitive processes should be discarded (Vaisey, 2008), and empirical work should probe the relevant cognitive processes (Martin, 2011; Pugh, 2013; Vaisey, 2009, 2014).

Fast on the heels of this cognitive turn came the "data deluge." The unprecedented availability of computer-readable cultural traces (Evans & Foster, 2011), combined with the proliferation of sophisticated computational methods, allowed sociologists to puzzle out patterns in cultural data of unprecedented scale and complexity (Evans & Aceves, 2016). Multiple special issues have explored the promise and peril of computational methods in the study of culture (Mohr & Bogdanov, 2013; Wagner-Pacifici, Mohr, & Breiger, 2015), and this issue of *Poetics* contains outstanding examples of computational cultural analysis.

I argue, however, that computation can do more than supply new methods for studying culture. It can actually change the way we think about culture. Previous approaches to the cognition-culture nexus have been limited by their neglect of what might be called *computational thinking*. Computational thinking is a theoretical stance that encourages detailed analysis of the "procedures for arriving

E-mail address: foster@soc.ucla.edu.

https://doi.org/10.1016/j.poetic.2018.04.007
Received 15 September 2017; Received in revised form 18 April 2018; Accepted 26 April 2018

at adaptive decisions[1] based on approximate and noisy information," i.e., of computation[2] (Krakauer, 2014).

Computational thinking about culture is urgently needed; sociological explanations make implicit assumptions and even claims about the underlying computational problems faced by social actors, as well as potential algorithmic and representational solutions (in)consistent with the "hardware" (i.e., human bodies, brains, artifacts, and material environments) in which they must be implemented (Marr, 1983; Peebles & Cooper, 2015). Computational thinking also leads naturally to formalization, with its attendant advantages (clarity) and disadvantages (loss of nuance). Grounding our formal models in computational thinking (e.g., contemporary ideas in machine learning and artificial intelligence) links our discipline directly to important interdisciplinary conversations, and promotes a deeper connection between the computational *methods* we use to analyze culture and the *theories* we develop about it.

This paper outlines some provisional steps toward a sociological theory of culture grounded in computational thinking. These steps address three key questions[3]: "How should we study culture?" (Section 2); "What is culture?" (Sections 3–4); and "How should we explain culture?" (Sections 5–6). Approaching these questions from a computational perspective exposes a key gap in current theories of culture—namely, an account of cultural learning—and suggests entirely different lines of attack on classic questions, e.g., whether culture is coherent or fragmented (DiMaggio, 1997) and "whether and how some cultural elements control, anchor, or organize others" (Swidler, 2013). To keep the discussion concrete, I use illustrative examples from the sociology of science throughout.

Because this paper is programmatic in nature, it covers more ground that is typical (or desirable) in a short essay. In some cases, the treatment of specific themes will be cursory, and important problems are consigned to footnotes or omitted entirely. I offer the paper in an exploratory spirit: to introduce the reader to promising directions and provocative ways of thinking, rather than follow a particular path to its conclusion.

My broader aims in this paper are threefold. First, to translate some key sociological ideas into computational language, so that they might be more easily digested by readers outside our theoretical tradition who are converging on similar questions (Hassabis, Kumaran, Summerfield, & Botvinick, 2017). Second, to identify the opportunities for *theoretical* synthesis between the sociological literature on culture and cognition and literatures on machine learning and computational learning theory (Goodfellow, Bengio, & Courville, 2016; Hassabis et al., 2017; Marblestone, Wayne, & Kording, 2016; Valiant, 2013). Third, to sketch out a research agenda that I hope will be attractive to others working in the theory, culture, science, and computational sociology communities.

2. What is computational thinking?

How should we study culture from a computational perspective? Turing's formal model of computation (Turing, 1937, 1950) epitomizes a research strategy: develop an abstract, implementation-independent model of a computational process (Valiant, 2013). Building on that strategy, the computational neuroscientist David Marr noted that "understanding the nature of the problem being solved" is often more important than "examining the mechanism (and the hardware) in which it is embodied" (Marr, 1983). I argue that we should apply Turing's strategy to the theory of culture: specifically, how people acquire and deploy culture, broadly construed as any **shared regularity** in the organization of experience or the generation of action **acquired through social life**. Understanding the computations involved in learning, deploying, and creating culture is an essential prerequisite to a complete causal account: we must know how pieces of culture get into actors, how they are recognized in social interaction, and how they are reproduced (or created) in actors' behavior.

Marr introduced a powerful heuristic for computational thinking. He articulated three "levels of explanation" for information processing systems: the *computational* ("what is the goal," "why is it appropriate" to the situation, and "what is the logic of the strategy"); the *representational and algorithmic* (how can this strategy be implemented, how are input and output represented, and what algorithms might connect them); and the *hardware* ("how can the representation and algorithms be realized physically") (Marr, 1983; Peebles & Cooper, 2015)[4]. These levels of analysis are, in principle, separate; in practice, they direct and constrain one another. Adequate description of the computational problem suggests potential representational and algorithmic solutions; such solutions can (and should) be rejected if they cannot be implemented in the relevant hardware (which in our case may include one or more human bodies, as well as the surrounding material environment).

In many cases, sociologists operate at a particular level of analysis without explicitly identifying that level or highlighting implications or constraints from other levels[5]. For example, considerations at the hardware level, e.g., from cognitive neuroscience (Cerulo, 2010; Lizardo, 2007, 2009) or embodied and situated cognition (Lizardo & Strand, 2010), should not be invoked to *close* the analysis of a particular phenomenon; instead, they *open it up* to further analysis. The claim that some phenomenon is embodied *or* situated does not address the two other levels of explanation.

Consider two examples where "thinking across Marr levels" has proved useful (without being labeled as such). First, consider the

[1] Note that "decision" does not imply deliberation, rationality, or conscious awareness. Note also that *all* information encountered by social actors is approximate and noisy. Hence humans, like all organisms, *generically* face computational problems.

[2] "Computation" should not be taken to imply the exactitude or determinism that it commonly connotes. Krakauer's definition emphasizes approximation and noise—messiness.

[3] I thank Mario Small for suggesting these questions as an organizing scheme.

[4] See Rescorla (2017) for a recent critique of Marr's levels when applied to natural (as opposed to artificial) computing systems. Krafft and Griffiths (In press) develop a Marr-ian approach to the analysis of social systems.

[5] We sociologists are happy to divide our analysis of a phenomenon across multiple *substantive* levels. We may get additional analytical purchase by breaking our analysis of social action across distinct *computational* levels, following Marr. I thank John Levi Martin for this helpful point.