Cognition 157 (2016) 156–173

Contents lists available at ScienceDirect

Cognition

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

Incremental implicit learning of bundles of statistical patterns

Ting Qian [a,*], T. Florian Jaeger [b,c,d], Richard N. Aslin [b]

[a] Department of Biomedical and Health Informatics, The Children's Hospital of Philadelphia, United States
[b] Department of Brain and Cognitive Sciences, University of Rochester, United States
[c] Department of Computer Science, University of Rochester, United States
[d] Department of Linguistics, University of Rochester, United States

ARTICLE INFO

Article history:
Received 24 November 2014
Revised 2 September 2016
Accepted 8 September 2016

Keywords:
Bayesian nonparametric models
Stimulus bundles
Incremental learning
Implicit learning
Sequential input
Memory

ABSTRACT

Forming an accurate representation of a task environment often takes place incrementally as the information relevant to learning the representation only unfolds over time. This incremental nature of learning poses an important problem: it is usually unclear whether a sequence of stimuli consists of only a single pattern, or multiple patterns that are spliced together. In the former case, the learner can directly use each observed stimulus to continuously revise its representation of the task environment. In the latter case, however, the learner must first parse the sequence of stimuli into different bundles, so as to not conflate the multiple patterns. We created a video-game statistical learning paradigm and investigated (1) whether learners without prior knowledge of the existence of multiple "stimulus bundles" — subsequences of stimuli that define locally coherent statistical patterns — could detect their presence in the input and (2) whether learners are capable of constructing a rich representation that encodes the various statistical patterns associated with bundles. By comparing human learning behavior to the predictions of three computational models, we find evidence that learners can handle both tasks successfully. In addition, we discuss the underlying reasons for why the learning of stimulus bundles occurs even when such behavior may seem irrational.

1. Introduction

One of the fundamental challenges in navigating the world is to guide our own behavior appropriately by forming a representation that captures the essential features of the task environment. Understanding how people construct this representation is a central problem in the study of learning and cognition. Importantly, in most real-world circumstances, this learning process must rely on input that unfolds gradually over time. Several formalisms have been proposed to explain how such learning occurs. For example, under the framework of Bayesian belief updating, learners are assumed to represent the task environment as a probabilistic model and update their estimates of model parameters after each observation (e.g., Anderson, 1991; Sanborn, Griffiths, & Navarro, 2010). In connectionist and other associative theories, learners are assumed to represent the associative weights between the variables of the task environment (both observable and hidden), and revise them according to the degree to which their previous settings had correctly predicted a new observation (e.g., Love, Medin, & Gureckis, 2004; McClelland & Rumelhart, 1981; Sakamoto, Jones, & Love, 2008). Regardless of which broad category an incremental learning model falls into, a common assumption is that the task environment can be summarized by a *single* set of parameters. Under this view, learning is essentially a process of continuously *revising* this single set of parameters, as the average properties of the input will eventually converge onto the true properties of the task environment with more and more observations.

While this assumption holds true for many laboratory tasks that have been employed in learning experiments, many real-world situations may challenge its validity (see also Gallistel, Krishan, Liu, Miller, & Latham, 2014; Gershman, Blei, & Niv, 2010; Kleinschmidt & Jaeger, 2015; Kording, Tenenbaum, & Shadmehr, 2007; Yu & Cohen, 2008 for similar concerns). For example, consider a task where a naive learner observes a sequence of daily weather phenomena over the course of a year. For a knowledgeable learner, this sequence of stimuli will contain four subsequences that exhibit unique and localized patterns — a *spring* subsequence that mainly consists of sunny and warm days, a *summer* subsequence of hot days, a *fall* subsequence of rainy days and a *winter* subsequence of snowy days.[1] In other words, this sequence of

* Corresponding author.

[1] This example is purposefully simplified in at least two aspects. First, there are also systematic changes within a season that can be learned and predicted. Second, the learning of seasons is not a purely statistical learning problem, but one that depends on the knowledge of astrophysics. However, to a *naive* learner, the learning of seasons can be (almost) treated as a categorization problem where the concept of seasons emerges primarily from the observation of the statistical patterns in daily weather.

http://dx.doi.org/10.1016/j.cognition.2016.09.002
0010-0277/

weather patterns consists of multiple *stimulus bundles*, each of which consists of a subsequence of stimuli that are linked to the same underlying properties of the task environment (Qian & Aslin, 2014). If a learner ignores the existence of these stimulus bundles and keeps only a running average of the "daily weather", the learned representation will be of poor predictive value as it does not converge onto the properties of any latent structure that we refer to as a *season*. Instead, if a learner is sensitive to the presence of stimulus bundles (or relatedly, the latent causes underlying them), they can form the representation that encodes four different average daily weather estimates, and by doing so, predict the weather given the current season with much better accuracy.

Although much research has been conducted on the topic of incremental learning, we know surprisingly little about whether people can successfully infer the presence of stimulus bundles in an incremental learning task, and more importantly, whether they can build a representation that captures the complexities of such a task environment. Here, we focus on the case where the stimuli in the same bundle collectively defines a *statistical pattern*. There are at least three challenges in learning such stimulus bundles. First, the latent state of a task environment, such as the notion of seasons in the above example, is not observable. As a result, a learner must infer when one bundle ends and another begins from the sequentially observed input, possibly by evaluating any deviation in the observed input with respect to the consistency of the statistical pattern (Gebhart, Aslin, & Newport, 2009). With correctly inferred bundle boundaries, each observation can then be used to revise the belief about the latent properties of the environment that apply only in the corresponding state of that bundle (cf., Summerfield, Behrens, & Koechlin, 2011; Yu & Dayan, 2005). (For instance, the weather on a summer day will only be relevant to the learning of summer weather, but not the weather of other seasons.) We refer to this first challenge as **the problem of bundle boundaries**. This problem is also often referred to as *non-stationarity*, i.e., the presence of multiple statistical patterns in the input (see Aslin, 2014; Brown & Steyvers, 2009; Cohen, McClure, & Yu, 2007; Gallistel et al., 2014; Jaeger & Snider, 2013). Second, a task environment may enter the same state multiple times by returning to an old state, producing sequences of stimulus bundles that vary in content but share the same underlying characteristics (e.g., over multiple summer seasons, the exact sequences of weather patterns will differ but the general tendency of the summer weather will stay the same). Recognizing the latent state of the environment underlying a stimulus bundle and retrieving its learned representation from memory can greatly reduce the cost of relearning and predict the observations more accurately. However, doing so means that learners must be able to differentiate and identify the statistical patterns in sequential input, most likely based on only partial observations of a complete bundle. We refer to this second challenge as **the problem of bundle identity**. Finally, a task environment may contain *contextual* cues that are correlated with the *identities* of states underlying stimulus bundles (e.g., Freidin & Kacelnik, 2011; Gureckis & Love, 2009). For example, in our weather scenario, the presence of shorter periods of daylight serve as a contextual cue that the season is "winter" (note that a contextual cue is different from cues that are *causally* related to observed stimuli, see Speekenbrink & Shanks, 2010; Wasserman & Castro, 2005 for example). But, it is not necessarily clear to a naive learner what these cues are or how reliable they are in a task environment (cf., Griffiths & Tenenbaum, 2009). In addition, without learning the statistical patterns of stimulus bundles, contextual cues themselves are of little value since they are merely labels (e.g., knowing the next season is named "summer" does not help predict the weather unless one has already had experience with the summer season). We refer to this third challenge as **the problem of contextual cue validity**.

The goal of the present article is to investigate whether people are able to overcome these three challenges in an incremental learning task. Our research, as well as the perspectives outlined above, builds upon insights from several separate lines of research. Although not commonly viewed as such, the problems of stimulus bundle boundary, bundle identity, and contextual cue validity affect learning in all aspects of perception, motor control, and higher level cognition. Closely related to the first challenge — the problem of bundle boundaries — is the issue of change detection, which has been studied in the context of decision-making tasks (e.g., Behrens, Woolrich, Walton, & Rushworth, 2007; Boorman, Behrens, Woolrich, & Rushworth, 2009; Brown & Steyvers, 2009; Nassar, Wilson, Heasly, & Gold, 2010; Payzan-LeNestour & Bossaerts, 2011; Speekenbrink & Shanks, 2010; Wilson & Niv, 2012). In these studies, participants are either asked to make sequential choices among several alternatives with various reward rates, or required to make predictions about certain variables of interest that in turn yield differential rewards. Either the reward rates associated with the alternatives or the variables themselves will change unpredictably, thus requiring participants to detect the changes and update their preferences. In our terminology, a series of trials with the same mapping between task variables and reward configuration would constitute a stimulus bundle, and after a change in this mapping, a new bundle begins. People detect such changes rather successfully and swiftly. However, there are two important limitations in relating the findings from this line of work to the problem of bundle boundaries in general. First, the change detection literature has almost exclusively focused on task environments with very simple statistical structures, such as changes in a Bernoulli distribution that specifies the reward configuration (Behrens et al., 2007), or changes in the mean of a Gaussian distribution that controls the properties of the stimuli (Nassar et al., 2010). It is unclear whether this ability to detect changes extends beyond those simple scenarios that have been investigated (see Gebhart et al., 2009; Kraljic, Samuel, & Brennan, 2008 for examples where learners facing complex problems fail to detect changes as quickly). Second, previous studies in this tradition have focused on changes that are signaled by immediate and explicit changes in reward. However, in everyday tasks, external reward may not be available and is often indirect and delayed. It is thus unclear how well learners will perform in an implicit learning task when they can only rely on the statistical pattern of the input itself as the primary means of detecting bundle boundaries.

The second challenge — the problem of bundle identity — concerns the ability of learners to see commonalities between various subsequences of input and to construct a representation that compactly encodes only the unique statistical patterns of the task environment (cf., Collins & Koechlin, 2012). A related behavior is reported in the animal conditioning literature, that after a period of behavioral extinction, the previously conditioned response of an animal can spontaneously recover (e.g., Sissons & Miller, 2009), be renewed (e.g., Bouton & King, 1983) or be reinstated (e.g., Rescorla & Heth, 1975). For animals to exhibit such behavior, they must be able to detect the boundary between bundles of conditioning trials and bundles of extinction trials, and represent the distinct causal contingencies associated with each type of bundle as separate and unique states of the task environment (cf., Gershman et al., 2010; Qian, Jaeger, & Aslin, 2012). Similarly in human learners, it has been shown that language users can adapt to the statistics of phonetic categories (Eisner & McQueen, 2005; Kraljic & Samuel, 2005, 2007; Norris, McQueen, & Cutler, 2003), words (Creel, Aslin, & Tanenhaus, 2008; Yildirim, Degen, Tanenhaus, & Jaeger, 2016), prosodic patterns (Kurumada, Brown, Bibyk, Pontillo, & Tanenhaus, 2012; Kurumada, Brown, & Tanenhaus, 2012), and syntactic structures (Fine, Jaeger, Farmer,