Cognition 153 (2016) 63–78

Contents lists available at ScienceDirect

Cognition

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

Original Articles

Statistical learning is constrained to less abstract patterns in complex sensory input (but not the least)

Lauren L. Emberson [a,b,c,*], Dani Y. Rubinstein [c,d,e]

[a] Brain and Cognitive Sciences Department, University of Rochester, USA
[b] Psychology Department, Princeton University, USA
[c] Psychology Department, Cornell University, USA
[d] Department of Neuroscience, Brown University, USA
[e] Section on Integrative Neuroimaging, Clinical and Translational Neuroscience Branch, National Institute of Mental Health, NIH, Bethesda, MD, USA

ARTICLE INFO

Article history:
Received 26 October 2014
Revised 12 April 2016
Accepted 15 April 2016

Keywords:
Psychology
Learning
Statistical learning
Implicit learning
Cognitive development
Pattern recognition
Vision
Perception
Perceptual learning
Object perception
Categorization

ABSTRACT

The influence of statistical information on behavior (either through learning or adaptation) is quickly becoming foundational to many domains of cognitive psychology and cognitive neuroscience, from language comprehension to visual development. We investigate a central problem impacting these diverse fields: when encountering input with rich statistical information, are there any constraints on learning? This paper examines learning outcomes when adult learners are given statistical information across multiple levels of abstraction simultaneously: from abstract, semantic categories of everyday objects to individual viewpoints on these objects. After revealing statistical learning of abstract, semantic categories with scrambled individual exemplars (Exp. 1), participants viewed pictures where the categories *as well as* the individual objects predicted picture order (e.g., $bird_1$—dog_1, $bird_2$—dog_2). Our findings suggest that participants preferentially encode the relationships between the individual objects, even in the presence of statistical regularities linking semantic categories (Exps. 2 and 3). In a final experiment we investigate whether learners are biased towards learning object-level regularities or simply construct the most detailed model given the data (and therefore best able to predict the specifics of the upcoming stimulus) by investigating whether participants preferentially learn from the statistical regularities linking individual snapshots of objects or the relationship between the objects themselves (e.g., bird_$picture_1$—dog_$picture_1$, bird_$picture_2$—dog_$picture_2$). We find that participants fail to learn the relationships between individual snapshots, suggesting a bias towards object-level statistical regularities as opposed to merely constructing the most complete model of the input. This work moves beyond the previous existence proofs that statistical learning is possible at both very high and very low levels of abstraction (categories vs. individual objects) and suggests that, at least with the current categories and type of learner, there are biases to pick up on statistical regularities between individual objects even when robust statistical information is present at other levels of abstraction. These findings speak directly to emerging theories about how systems supporting statistical learning and prediction operate in our structure-rich environments. Moreover, the theoretical implications of the current work across multiple domains of study is already clear: statistical learning cannot be assumed to be unconstrained even if statistical learning has previously been established at a given level of abstraction when that information is presented in isolation.

1. Introduction

Throughout our lifetime, experience shapes our mental model of the world; learning from patterns, regularities, or statistics in the environment is a way that experience might powerfully and incidentally shape cognition. Rich mental models or representations of the environment in turn support the prediction of upcoming sensory input facilitating both rapid and accurate cognitive processing when these predictions are correct, or error-signals and further learning and/or adaptation when these predictions are incorrect.

* Corresponding author at: Peretsman-Scully Hall, Psychology Department, Princeton University, Princeton, NJ 08540, USA.

E-mail address: laurenle@princeton.edu (L.L. Emberson).

http://dx.doi.org/10.1016/j.cognition.2016.04.010
0010-0277/

Indeed, statistical learning,[1] or the sensitivity to statistical information in sensory input, is increasingly considered to form the foundation of numerous and diverse perceptual and cognitive abilities. For example, not only is statistical learning believed to be an essential component of the development of language (Romberg & Saffran, 2010; Saffran, Aslin, & Newport, 1996; Thiessen & Erickson, 2013), but researchers have increasingly turned to mechanisms of adaptation, or sensitivity to statistical or distributional patterns in the environment, to explain adult comprehension (e.g., Fine & Jaeger, 2013; Kleinschmidt & Jaeger, 2015). Likewise in vision, there have been many demonstrations that the visual system is sensitive to statistical information both immediately after exposure (e.g., neural demonstrations: Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008; Turk-Browne, Scholl, Chun, & Johnson, 2009; Turk-Browne et al., 2010) and cumulatively over a lifetime of visual experience (Purves, Monson, Sundararajan, & Wojtach, 2014; Purves, Wojtach, & Lotto, 2011). Indeed, even young infants have the ability to perform complex visual statistical learning tasks (as young as 2-months-old for sequential visual statistical learning tasks; Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002 for spatial statistical learning tasks with 9-month-olds).[2] The more the field searches, the more it finds that the sensory environment is filled with statistical information. Paired with the rapid and ubiquitous statistical learning abilities of both infants and adults in many different types of input, it is clear that understanding cognitive responses to statistical information is foundational to understanding many aspects of perception and cognition.

However, despite its clear importance, little is known of how statistical learning proceeds in our daily environments. A central problem is that it is unknown how statistical learning mechanism(s) operate given the richness of the statistical information that we routinely receive. In this paper, we focus on a particular type of *rich* information: **what happens when statistical information is present across multiple levels of abstraction?** How does learning proceed when an observer can learn that not only do categories of objects predict another, but also that individual objects within categories predict each other? This kind of overlapping statistical information is routinely presented to us. For example, the predictive relationship between dogs and leashes exists based on abstract categories as well as on the actual objects or exemplars seen in the world (e.g. dogs have their specific leashes). There are many other examples where both categories of objects and also the individual objects within these categories are predictive of each other: types of animals and their habitats (e.g., fish and bodies of water), chairs and tables, people and their personal affects (e.g., wallets, cell phones). Thus, when two groups of objects are paired in the world, there is visual statistical information from which one can learn, all the way from low-level visual similarities across exemplars within a category, to relationships between individual objects, to their semantic, abstract categories. While the current paper focuses on visual statistical learning, similar inputs are found in other domains; for example language input, where there is overlapping statistical information at the levels of phonology, lexical items, and syntactic information when one rehears a single utterance.

The current paper addresses what is learned when there is predictive visual information across levels of abstraction. Is statistical learning *a priori* constrained to learn particular patterns? There are two types of competing theories on this: The first proposes that statistical learning is largely unconstrained; the second proposes that statistical learning is inherently constrained to learn from particular types of regularities.

By and large, previous research has supported the view that learning mechanism(s) are largely unconstrained: Statistical learning has been demonstrated in multiple sensory modalities (e.g., Conway & Christiansen, 2005) and across a wide range of perceptual input. For example, in the visual modality, learning can occur from sequences of gestures (Baldwin, Anderrson, Saffran, & Myers, 2007) or abstract shapes (Fiser & Aslin, 2001) and from spatial and non-spatial information (Mayr, 1996). While the majority of these studies have focused on learning probabilistic relations between individual objects, there is evidence that learning can occur at higher levels of informational abstraction including based on categories of nonsense words (Reeder, Newport, & Aslin, 2013; Saffran, 2002) or familiar semantic categories (Brady & Oliva, 2008). Overall, these studies, among many others, have led to the belief that statistical learning is largely unconstrained. That is, if there is any reliable probabilistic information in the environment, humans can learn from it regardless of modality or level of abstraction.

If statistical learning were unconstrained, how would learning proceed when a learner is faced with statistical information across multiple levels of abstraction simultaneously? A logical prediction is that an observer would learn across these multiple levels of representation *simultaneously* and in parallel. While this topic has been largely unexamined, there is a handful of evidence that learning can interact across levels of representation: Onnis, Waterfall, and Edelman (2008) found that the narrative structure of successive utterances supported statistical learning of individual syllables in adult learners; similarly, Saffran (2001) found that syntactic information in the form of phrase structures aided learning at lower levels (i.e., learning across individual elements), and Koch and Hoffmann (2000) suggested that different levels of information can compete in a serial reaction time (SRT) task.

However, the evidence that statistical learning can occur for all types of input and all levels of abstraction has been obtained under very constrained experimental conditions which may not reveal how learning operates over the rich statistical input encountered "in the wild". To illustrate, Brady and Oliva (2008) use a paradigm where the categories of scenes are predictive of picture order but individual scenes are not (e.g. beaches predict kitchens as categories of scenes but $beach_1$ does not predict $kitchen_1$). Participants show evidence of learning the relationship between abstract categories of scenes, but it is not possible to learn based on individual scenes because only category-level regularities are present. While studies such as Brady and Oliva (2008) provide essential existence proofs, statistical learning has not been investigated in the context of the rich input that an everyday learner encounters, where multiple types of statistical regularities are present simultaneously.

The alternative theory is that statistical learning does have some constraints and that these constraints bias what is learned when presented with rich input. These theoretical views come in two types: The first type directly concerns the question of

[1] This paper focuses on "statistical learning," which is typically studied with participants passively learning from stimuli endowed with probabilistic, distributional or statistical information (though also see Hunt & Aslin, 2001; Perruchet & Pacton, 2006 for related behavioral paradigms). However, many, and perhaps all, types of learning can be considered *statistical* (i.e., diverse learning tasks can be commonly characterized by the patterns or statistical information presented). In addition, the ability to incidentally learn from statistical regularities in the environment has been found to recruit multiple classically-defined learning and memory systems (see Emberson & Amso, 2012; Karuza et al., 2013; Turk-Browne, Scholl, Chun, & Johnson, 2009; Turk-Browne, Scholl, Johnson, & Chun, 2010: the hippocampus, basal ganglia and inferior frontal cortex are involved in learning from statistical regularities). Thus, we consider statistical learning to be representing a general and essential aspect of learning: the ability to pick up on statistical information in the environment and use it to shape one's behavior.

[2] Unlike the numerous studies linking statistical learning to language development, additional studies are needed to more directly link visual statistical learning to visual development.