Consciousness and Cognition 27 (2014) 155–167

Contents lists available at ScienceDirect

Consciousness and Cognition

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

Modelling unsupervised online-learning of artificial grammars: Linking implicit and statistical learning

Martin A. Rohrmeier [a,b,*], Ian Cross [b]

[a] Cluster Languages of Emotion, Freie Universität Berlin, Habelschwerdter Allee 45, 14195 Berlin, Germany
[b] Centre for Music and Science, Faculty of Music, University of Cambridge, United Kingdom

ARTICLE INFO

Article history:
Received 26 May 2011

Keywords:
Unsupervised learning
Online learning
Artificial grammar learning
Implicit learning
Incidental learning
Statistical learning
Computational modelling
N-gram model
Simple recurrent network
Competitive chunking

ABSTRACT

Humans rapidly learn complex structures in various domains. Findings of above-chance performance of some untrained control groups in artificial grammar learning studies raise questions about the extent to which learning can occur in an untrained, unsupervised testing situation with both correct and incorrect structures. The plausibility of unsupervised online-learning effects was modelled with n-gram, chunking and simple recurrent network models. A novel evaluation framework was applied, which alternates forced binary grammaticality judgments and subsequent learning of the same stimulus. Our results indicate a strong online learning effect for n-gram and chunking models and a weaker effect for simple recurrent network models. Such findings suggest that online learning is a plausible effect of statistical chunk learning that is possible when ungrammatical sequences contain a large proportion of grammatical chunks. Such common effects of continuous statistical learning may underlie statistical and implicit learning paradigms and raise implications for study design and testing methodologies.

1. Introduction

Humans are very efficient learners. We constantly learn without intention and without awareness, and in this respect implicit learning constitutes one powerful and fundamental root mechanism of learning (Reber, 1993) spanning domains like language and music acquisition, steering through traffic, or playing sports (e.g. Land, 1998; Masters, MacMahon, & Pall, 2004; Rebuschat, 2008; Reed, McLeod, & Dienes, 2010; Rohrmeier, 2010; Rohrmeier & Rebuschat, 2012; Rohrmeier, Fu, & Dienes, 2012; Williams, 2009). Humans learn and adapt to the environment, whilst being in the midst of ongoing activities: we pick up individual characteristics, or melodic features in a piece of music while we are listening or dancing to it, sportsmen are able to adapt to characteristics of their opponents and the environment while playing, and musicians adapt to characteristic musical patterns of other musicians while improvising together. From this perspective, a phenomenon such as online-learning seems hardly surprising. This paper studies effects of online learning in the context of the artificial grammar learning paradigm.

Participants acquire implicit knowledge about rule based structures quickly and there is converging evidence from the different experimental traditions of artificial grammar learning, statistical learning and serial reaction time experiments. Under the artificial grammar learning paradigm participants rapidly acquire rule-based structures generated from finite-state grammars (or other types of structures) often after very short exposure periods to a set of grammatical structures

* Corresponding author. Address: I^2 Intelligence Initiative, Massachusetts Institute of Technology, Cambridge, MA, United States.
E-mail addresses: mrohrmeier@cantab.net (M.A. Rohrmeier), ic108@cam.ac.uk (I. Cross).

http://dx.doi.org/10.1016/j.concog.2014.03.011
1053-8100/

(Reber, 1993; Pothos, 2007). Statistical learning studies provide strong evidence that adult and infant participants are able to acquire words, tone or timbre patterns that are building blocks of the sequence, whilst being presented a continuous stream of syllables or tones (e.g. Aslin & Newport, 2012; Newport & Aslin, 2004; Rebuschat & Williams, 2012; Saffran, Aslin, & Newport, 1996; Saffran, Johnson, Aslin, & Newport, 1999; Schön et al., 2008; Schön & François, 2011; Tillmann & McAdams, 2004). Such studies provide evidence for incidental and unsupervised learning, i.e. learning of unlabelled stimulus structures without intending to learn and without having error or reward feedback available. Similar statistical learning experiments found that adult participants acquired tone distributions and sequence structure in new musical systems (Loui, Wessel, & Hudson Kam, 2006; Loui & Wessel, 2008; Loui, Wessel, & Hudson Kam, 2010; Rohrmeier & Cross, 2009; Rohrmeier, Rebuschat, & Cross, 2011; Rohrmeier & Cross, 2013; Rohrmeier & Widdess, 2012). Similarly serial reaction time experiments employ measures of response times to examine the rapid online learning of rule-based structures (e.g. Clegg, DiGirolamo, & Keele, 1998; Fu, Fu, & Dienes, 2008; Hunt & Aslin, 2001; Jiménez, 2008; Koch & Hoffmann, 2000; Nissen & Bullemer, 1987; Shanks & Johnstone, 1999).

From an overarching perspective, these different experimental paradigms can be regarded as methodologically different approaches to the same underlying phenomenon of incidental, unsupervised learning (Perruchet & Pacton, 2006). The main focus of this article is to analyse and model one particular phenomenon found in the performance of untrained control groups in artificial grammar learning experiments and to relate this to the context of incidental, unsupervised learning and learning under noisy conditions.

Typically, an artificial grammar learning experiment involves two main parts, a learning phase and a testing phase. During the first, participants are familiarised with a set of sample sequences generated from an artificial grammar under incidental learning conditions, for instance, using an unrelated distraction task. In the subsequent testing phase, participants are presented an equal number of grammatical and ungrammatical sequences in random order and are required to respond with forced-choice ratings (for instance, familiarity) to assess their acquired knowledge. The procedure frequently combines this with additional confidence ratings which assess participants' awareness of the acquired knowledge. In contrast to experimental groups, control groups are not exposed to the same training phase. They are either trained on unrelated or random materials or receive no initial training at all (cf. Dienes & Altmann, 2003). In the latter case, there is a common presumption that controls make their choices by mere chance or response biases as they could not know which stimuli belong to an unknown set (the experimental group was exposed to) and which not (cf. Reber & Perruchet, 2003).

However, even though the testing phase contains 50% ungrammatical structures, that seem to make incidental learning implausible, several studies have found control groups performing above (or below) chance level. Such findings are particularly unexpected and interesting given that the testing phase involves exposure to half grammatical and half ungrammatical structures. For instance, Dulany, Carlson, and Dewey (1984) found that untrained controls, who had no training and completed a testing phase straight away, performed above chance. Redington and Chater (1994) have found a similar above chance performance based on the same materials. This might suggest that participants have picked up some regularity in the structures during the testing.

Redington and Chater (1996) discuss the possibility of such a learning process, whereas Reber and Perruchet (2003) argue, criticised by Dienes and Altmann (2003), that above chance performance of a control group would not stem from a learning effect but from confounding structural biases that may be easy to detect. However, some studies using musical materials (Loui, 2003; Rohrmeier et al., 2011; Tillmann & McAdams, 2004) found above chance (.5) performance levels of untrained controls (performing a generalisation task) which may reopen the question about a potential rapid online-learning effect. Further converging evidence is provided by Dienes, Broadbent, and Berry (1991, Experiment 1), which involved a group trained with a training set consisting of equal numbers of grammatical and ungrammatical items and tested with another set of grammatical and ungrammatical items. This group performed at 60%, which renders plausible an online learning such as that described above.

Rohrmeier et al. (2011) analyzed the time course of the performance of an untrained group compared with a trained group and found a power function curve that matched the data (see the plot of the human performance data (thick black line) and the power function fit (dotted line) from their experiment as shown in the centre of Fig. 1). The fact that the performance curve begins at chance level (and not above) and steadily rises to a level of .62, suggests that participants gradually pick up some knowledge that enables them to distinguish the structures, with little prior bias. In particular, this study found the performance of the untrained group to be significantly above chance level already after the 11th test sequence. These findings suggest that participants gradually acquire information about the stimuli as they move through the (randomised) testing phase. Unlike experimental groups in related artificial grammar learning studies, serial reaction time, or statistical learning studies, this result is unexpected as a case of incidental, unsupervised learning because the materials do not consist of only grammatical structures. This suggests that participants pick up information about regularities in materials within a noisy environment.

There are not many studies with untrained control groups and they do not provide a comprehensive picture. Above chance-level performance was found by Dulany et al. (1984), Redington and Chater (1994), Dienes (reported in Redington & Chater (1996), Rohrmeier et al. (2011), Loui (2003); below chance-level performance was found by Meulemans & Van der Linden (1997) and in one experiment by Reber & Perruchet (2003). In turn, the studies by Altmann, Dienes, & Goode (1995), Kuhn & Dienes (2005) and Reber & Perruchet (2003) found performance at chance-level. Even though this body of experimental evidence is not yet large, it suggests that there may be an effect of learning that happens during the course of the testing phase.