Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research 45 (2017) 59–69

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Deep learning and punctuated equilibrium theory

Simon Hegelich

Bavarian School of Public Policy, Technical University of Munich, Richard-Wagner-Str. 1, D-80333 Munich, Germany

Received 21 June 2016; received in revised form 11 December 2016; accepted 28 February 2017
Available online 11 April 2017

Abstract

Deep learning is associated with the latest success stories in AI. In particular, deep neural networks are applied in increasingly different fields to model complex processes. Interestingly, the underlying algorithm of *backpropagation* was originally designed for political science models. The theoretical foundations of this approach are very similar to the concept of *Punctuated Equilibrium Theory* (PET). The article discusses the concept of *deep learning* and shows parallels to PET. A showcase model demonstrates how deep learning can be used to provide a missing link in the study of the policy process: the connection between attention in the political system (as inputs) and budget shifts (as outputs).

Keywords: Deep learning; Neural networks; Punctuated equilibrium; Policy process; Backpropagation

1. Introduction

Deep learning is associated with the latest success stories in AI. From autonomous cars to AI beating a Go-master: deep learning is the method of choice to construct machine learning models that are useful in many complex situations. Taking this success-story into account, it seems obvious that political science could profit from this method as well. A whole branch of political science approaches sees the *policy process* as a kind of cognitive system that transforms political inputs from society into outputs. *Punctuated Equilibrium Theory* (PET) is a very successful concept in political science that is grounded in the theoretical works of Herbert Simon on *bounded rationality* (Simon, 1955). What separates this approach from rational choice theories is - amongst others - the understanding of organizations (Jones, 2003). While theories that rely solely on market mechanisms can see organizations only as individuals (maximizing their utility function) or as markets, where individuals meet. In bounded rationality organizations are seen as cooperations of individuals identifying with the organization. In many cases, this makes organizations much more effective than markets, especially because parallel processes can be organized with less information costs. But this makes organizations quite complex, as well.

> If complex systems must operate in a constantly changing environment [. . .] they must modify their structures at a corresponding pace. The need for close coordination, even in the presence of strong identification with the organization's goals, places a very heavy burden on a system's capacity to evolve toward greater effectiveness under changed conditions. For although identification reduces the need to police self-interest and to ensure its compatibility with organizational objectives, it also causes excessive influence of existing organizational practices and identifications upon decisions that should be adapting to a changing world (Simon, 2000, p. 753).

E-mail address: simon.hegelich@hfp.tum.de

http://dx.doi.org/10.1016/j.cogsys.2017.02.006
1389-0417/

Deep neural networks are very suitable for dealing with this complexity. Taken into consideration the strong interest of Herbert Simon in machine learning, it seems surprising that the connection between PET and deep learning has not yet been tapped for the benefit of political science, at least to the best of my knowledge. This is even more astounding against the background that the underlying algorithm of *backpropagation* was originally designed for political science models. Unfortunately, the formerly mutual connections between computer science and political science seem today to have eroded. Cognitive system research seems to be an ideal place, to bridge this gap between the two disciplines by highlighting the parallels of deep learning and PET. Hopefully, this attempt will work as a humble contribution to re-establish the interdisciplinary field of *political data science*.

The article proceeds as follows: First, the concept of deep learning is introduced with a focus on neurons as the building blocks of neural nets. Second, the idea to understand the policy process as information processing is recapitulated and linked to the problems of complexity and noise. On this basis, third, a showcase of the implementation of deep learning in PET is presented. It is demonstrated, that deep learning is capable of linking attention signals in the political system to policy outcomes in the form of budget changes. Fourth, the theoretical relevance of this demonstration is discussed. Finally, the article provides an outlook explaining how more advanced deep learning models could push the development of PET even further.

2. What is deep learning?

Deep learning as a machine learning approach that is based on *neural networks*. A very good description of neural networks is given by Pat Langley and Herbert Simon:

> One major paradigm, associated with the area of neural networks, represents knowledge as a multilayer network of units that spreads activation from input nodes through internal units to output nodes. Weights on the links determine how much activation is passed on. The activations of output nodes can be translated into numeric predictions or discrete decisions about the class of the input. [...] One common learning algorithm, among the many that have been explored, carries out gradient descent search through the space of weights, modifying them in an attempt to minimize the errors that the network makes on training data (Langley & Simon, 1995).

There are different opinions which conditions make a neural network actually "deep". A very basic idea of deep neural networks is the combination of *multiple hidden layers* of non-linear transformations. In the show-case model used in this article for demonstration this basic definition of deep learning is applied. "Real" deep learning goes far beyond the simple addition of layers in neural networks but alters the underlying algorithms to create feedback loops and reinforcement learning. These aspects will be discussed in Section 6.

In an analogy to biological processes in the brain the building blocks of neural networks are called "neurons". A neuron collects different inputs and transfers them into a non-linear output. From a mathematical point of view, a neuron combines two functions: a summation function $f(s)$ and an activation function $f(a)$. "Given a sample of input attributes $x_1, \ldots, x_n$ a weight w_{ij} is associated with each connection into the neuron" (Lewis, 2016, p. 16). All inputs are summed up according to:

$$f(s) = \sum_{i=1}^{n} w_{ij} x_j + b_j \tag{1}$$

The parameter b_j is the bias and can be interpreted like an intercept in regression. It allows to the activation function to be shifted upwards or downwards.

The activation function takes the result of the summation as an input and transfers it in a non-linear way, usually to values between 0 to 1 or −1 to +1. There are many different functions that can be used as activation function but the s-shaped sigmoid function is very common (Friedman, Hastie, & Tibshirani, 2001, p. 393):

$$f(a) = sig(t) = \frac{1}{1 + e^{-t}} \tag{2}$$

The reason for the popularity of the sigmoid function is that it can be differentiated very easily, which is important for the optimization process in neural networks.

A neural network is a combination of single neurons. In a deep neural network the output from a layer of neurons functions as input for the next layer (see Fig. 1).

It is import to note that the strength of the connections of the neurons is determined by the weights assigned in the summation function of the following neuron. An output from a neuron (or the input layer) may have weight 0 in the summation function of one following neuron - i.e. it does not count for activation at all. But for another neuron the same output may have a higher weight that makes it very important for activation. Two things should become clear at this point:

1. A deep neural net is able to represent a very complex non-linear prediction space.
2. The final result of the model is determined by the assigned weights of the neurons.

This leads to the question: How are the weights assigned?

The basic algorithm to calculate the weights is called *backpropagation*. "The first practical application of backpropagation was for estimating a dynamic model to predict nationalism and social communication in 1974" (Werbos, 1994, p. 270) by Paul J. Werbos in his dissertation. The connection between political science and deep learning