



AFRL-AFOSR-VA-TR-2019-0275

Sticky Reasoning within Learning Representations

Yann LeCun
NEW YORK UNIVERSITY
70 WASHINGTON SQUARE S
NEW YORK, NY 10012-1019

09/18/2019
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/RTA2

DISTRIBUTION A: Distribution approved for public release.

Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE		Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services, Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.			
1. REPORT DATE (DD-MM-YYYY) 18-09-2019		2. REPORT TYPE Final Performance	
		3. DATES COVERED (From - To) 15 Sep 2015 to 14 Sep 2018	
4. TITLE AND SUBTITLE Sticky Reasoning within Learning Representations		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER FA9550-15-1-0441	
		5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) Yann LeCun		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NEW YORK UNIVERSITY 70 WASHINGTON SQUARE S NEW YORK, NY 10012-1019 US		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AF Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203		10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR RTA2	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-VA-TR-2019-0275	
12. DISTRIBUTION/AVAILABILITY STATEMENT A DISTRIBUTION UNLIMITED: PB Public Release			
13. SUPPLEMENTARY NOTES			

14. ABSTRACT

representations in which the explanatory factors of variations are disentangled. We first proposed an architecture and learning procedure that automatically learns representations of images in which the content is separated from the pose. The method is completely unsupervised, not requiring that the pose or the identity of the object in the image be provided [Mathieu NIPS 2016]. While this work used generative adversarial networks (GAN), we found the concept wanting and proposed a new formulation of GANs called Energy-Based GAN. Our method was the first to allow the generation of realistic images at 128x128 resolution. Generative models, such as GANs are a key component of unsupervised procedures that can generate complex patterns from latent variables drawn from a simple distribution [Zhao ICLR 2017] (over 550 citations to date).

Next, we developed an unsupervised representation learning method based on the concept of regularized auto-encoder. An auto-encoder is a deep learning architecture that is trained to reproduce its input on its output, while funneling the representation through an information bottleneck. In previous methods, the information bottleneck is implemented by making the dimension of the representation small (e.g. PCA, hourglass auto-encoders), by making the representation sparse (e.g. sparse auto-encoder with L1 regularization of the code), or by normalizing the code vector and adding noise to it (e.g. Variational Auto-Encoders).

15. SUBJECT TERMS

Convolutional Networks, Deep Learning, Machine Learning, Machine Perception

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			RIECKEN, RICHARD
Unclassified	Unclassified	Unclassified	UU		19b. TELEPHONE NUMBER (Include area code) 703-941-1100

Sticky Reasoning within Learning Representations.

Final report.

2019-09-13

**Yann LeCun
New York University**

In this project, we addressed two questions: (1) how to learn abstract representations of data such as text and video; (2) how design trainable architectures that have persistent memory and can perform long chains of reasoning.

1. Learning abstract representations of text and images

One goal of learning representations of data is to produce representations in which the explanatory factors of variations are disentangled. We first proposed an architecture and learning procedure that automatically learns representations of images in which the content is separated from the pose. The method is completely unsupervised, not requiring that the pose or the identity of the object in the image be provided [Mathieu NIPS 2016].

While this work used generative adversarial networks (GAN), we found the concept wanting and proposed a new formulation of GANs called Energy-Based GAN. Our method was the first to allow the generation of realistic images at 128x128 resolution. Generative models, such as GANs are a key component of unsupervised procedures that can generate complex patterns from latent variables drawn from a simple distribution [Zhao ICLR 2017] (over 550 citations to date).

Next, we developed an unsupervised representation learning method based on the concept of regularized auto-encoder. An auto-encoder is a deep learning architecture that is trained to reproduce its input on its output, while funneling the representation through an information bottleneck. In previous methods, the information bottleneck is implemented by making the dimension of the representation small (e.g. PCA, hourglass auto-encoders), by making the representation sparse (e.g. sparse auto-encoder with L1 regularization of the code), or by normalizing the code vector and adding noise to it (e.g. Variational Auto-Encoders). The method we proposed was to bound the information content (the entropy) of the code by forcing it to be

indistinguishable from a distribution with a known entropy. The method uses an adversarial method to compare the reference distribution with the code distribution. This system can generate text whose content can be manipulated by playing with the latent representation [Zhao ICML 2018] (collaboration with Sasha Rush, Harvard)

Lastly, we proposed an architecture for unsupervised learning of dependency graphs that can be transferred from one task to another in natural language processing [Zhao NIPS 2018] (collaboration with Russ Salakhutdinov, CMU).

2. Reasoning and Memory in Deep Learning Systems

A major challenge in modern machine learning is to get learning systems to perform long chains of reasoning. Recurrent neural nets have been the topic of extensive research along those lines. But traditional recurrent nets don't seem to be able to remember a state for very long. In fact, the ability to memorize (to possess stable states) was long thought to be incompatible with gradient-based learning procedures such as back-propagation through time (BPTT), as the gradients would vanish through time steps where the state is stable. However, one can escape this conundrum by using ideas from quantum mechanics: it is possible to store memories, not in stable states, but in stable orbits. One simply needs to constrain the dynamics of the network to be reversible to solve the vanishing gradient problem. It is possible to read out a stable memory from a continuously-evolving orbit through a non-linear pooling operation on the state vector. This is somewhat similar to a quantum-mechanical system whose dynamics is unitary (and therefore invertible with no stable state), and where a "measurement" removes the information relative to the rotating phase thereby producing a stable state. The method was shown to perform tasks that require long-term storage of information with higher accuracy than previous methods [Henaff ICML 2016].

Next, we proposed a deep learning architecture that is capable of maintaining a current state of the world in its memory, and to update it in the face of new information (like an event). For example, if a system reads a sentence like "john walks to the kitchen", it should maintain a memory slot for "john" storing a vector that encodes John's attributes, such as "location=kitchen". Similarly, the system should possess a memory slot for "kitchen" whose state vector represents the fact that John is present in it. The architecture is composed of a number of recurrent sub-networks, each of which is meant to

represent the attributes of an entity (such as John or the kitchen). When the system is fed a sentence describing an event, it produces a key vector that is compared each of the entities' key vectors and allows the entity's state vector to be updated of the keys match. The system is trained by being fed a sequence of events and a question about the final state of the world. It is then given the answer to the question and, through backpropagation, to update its parameters to produce the given answer. This was the first method to solve all the 20 of the so-called bAbI question-answering tasks [Henaff ICLR 2017].

This led us to pursue work on learning predictive world models. A world model is a trainable function (a neural net) that predicts the next state of the world, given the previous state and an action taken by the agent. One of the main issues in learning such models is to deal with the intrinsic unpredictability of the real world (or of simulated stochastic environments). Once a model of the world is trained, it can be used to plan a sequence of actions so as to optimize a particular objective. For example, if a learning agent is to land a spaceship on the moon or rendez-vous with another spacechip in orbit, it must possess a model of its own dynamics. Planning a trajectory under these conditions is called Model Predictive Control (MPC). But the problem is somewhat complicated when the actions are discrete. We propose a method to learn a predictive model and to learn a policy with discrete and continuous action spaces. Planning by minimizing an objective can be seen as a special case of a general form of reasoning. Many forms of reasoning and inference can be reduced to optimization problems [Henaff 2017a] [Henaff 2017b].

[Zhao NIPS 2018] GLoMo: unsupervised learning of transferable relational graphs

Z Yang, J Zhao, B Dhingra, K He, WW Cohen, RR Salakhutdinov, ...
Advances in Neural Information Processing Systems, 8950-8961

[Zhao ICML 2018] Adversarially Regularized Autoencoders for Generating Discrete Structures

J Zhao, Y Kim, K Zhang, AM Rush, Y LeCun
International Conference on Machine Learning (ICML 2018)

[Henaff 2017b] Prediction under uncertainty with error-encoding networks

M Henaff, J Zhao, Y LeCun
arXiv preprint arXiv:1711.04994

[Henaff 2017a] Model-based planning with discrete and continuous actions

M Henaff, WF Whitney, Y LeCun
arXiv preprint arXiv:1705.07177

[Henaff ICLR 2017] Tracking the World State with Recurrent Entity Networks

M Henaff, J Weston, A Szlam, A Bordes, Y LeCun
international Conference on Learning Representations (ICLR 2017)

[Zhao ICLR 2017] Energy-based generative adversarial network

J Zhao, M Mathieu, Y LeCun
international Conference on Learning Representations (ICLR 2017)

[Henaff ICML 2016] Orthogonal RNNs and Long-Memory Tasks

M Henaff, A Szlam, Y LeCun
international Conference on Machine Learning (ICML 2016)

[Mathieu NIPS 2016] Disentangling factors of variation in deep representation using adversarial training

MF Mathieu, JJ Zhao, J Zhao, A Ramesh, P Sprechmann, Y LeCun
Advances in Neural Information Processing Systems, 5040-5048
(2016).