The HLF Workshop on Deep Learning and Neuroscience

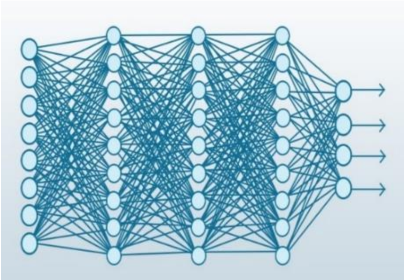

Understanding and modeling intelligence is one of the greatest problems in science and technology today. Making significant progress towards these challenges will require the interaction of several disciplines involving neuroscience and cognitive science in addition to computer science, robotics and machine learning. Deep Learning, a relatively new area of artificial intelligence and machine learning, has made huge advances in the last few years, and shown remarkable empirical success in many applications such as image categorization, face identification, action recognition, speech recognition, machine translation just to name a few.

Although the advances in performance during the last few years are exclusively an engineering feat, the cumulative effect of much increased computer power, the availability of manually labeled large data sets and a small number of incremental technical improvements, deep learning at its core draws inspiration from brain science, and is seen as the beginning of a breakthrough with the potential of opening new fields for science. Deep Learning has close connection with human brain, and at times, their learning has been compared with the learning in the human brain. Deep networks trained with Imagenet seem to mimic not only the recognition performance but also the tuning properties of neurons in cortical areas of the visual cortex of monkeys. Despite many such arguments by multiple researchers, these connections between deep learning and neuroscience are still poorly understood and remain an active topic of research.

The advancements in deep learning and their presumptive ability to mimic human brain has opened up several research questions. Do humans learn in a similar way as deep networks do i.e. back-propagating the error? Does brain also optimize some cost function as deep networks do? Can deep networks trained with millions of examples mimic an important subset of the basic building blocks for brain-like intelligence? Deep networks requires millions of examples to learn but children do not. Then how can deep networks mimic human learning? Are recent advancements in deep learning such as memory networks a step towards creating a prototype of human memory? In this workshop we plan to bring researchers from various fields such computer science, neuroscience, mathematics, behavior science, cognitive science together to discuss these questions.

Organizers

Arvind Agarwal, IBM Research, India
Felix Putze, University of Bremen, Germany
Subhrajit Roy, IBM Research, Australia

Mentor

Prof. Johannes Schöning, University of Bremen, Germany

Coordinates

Time and Date: Monday, September 24th, 14:30-16:00 (workshop session I)
Venue: Main Meeting Venue (Heidelberg University on University Square, 69117 Heidelberg)

Workshop Program

Introduction: Brief introduction to the workshop format and the three different subgroups for the subsequent discussion. Each participant of the workshop introduces himself or herself with one sentence.	14:30 – 14:45
Group Discussion: Split the group in smaller subgroups for more focused discussions on different topics. Each subgroup is moderated by one of the workshop organizers and stimulated by short impulse presentations, supported by posters.	14:45 – 15:30
Session-1: Knowledge Discovery & Learning **Moderator:** Arvind Agarwal **Impulse Presentation:** Ansif Arooj : Strengthening Learning by Prior Knowledge Lydia Braunack-Mayer: Man vs. Machine: Can machines really 'learn'? Arvind Agarwal : How Brains Represent Thousands of Objects	
Session-2: Computing in the Brain **Moderator:** Subhrajit Roy **Impulse Presentation:** Oliver Gäfvert: Differentiable Neural Computers Subhrajit Roy: How neuroscience can influence Deep Learning	
Session-3: Principles of Human Cognition for Next Generation Machine Learning **Moderator:** Felix Putze **Impulse Presentation:** Tapasya Patki: Modeling how humans process language Anwesha Das: Desh: Deep Learning for HPC System Health Resilience	
Final joint discussion: After the discussion, each moderator will give a short summary of the discussion. We will discuss next steps and opportunities to continue the workshop theme in other forms.	15:30 – 16:00

Detailed Program

Session-1: Knowledge Acquisition and Learning

We have known since antiquity that the seat of learning is the human brain. The process of knowledge acquisition and learning inside the brain has baffled researchers since the beginning of science, and it still remains an unsolved mystery. Learning is important because no one is born with the ability to function competently as an adult in society. It is especially important to understand the kinds of learning experiences that lead to transfer, defined as the ability to extend what has been learned in one context to new contexts. While understanding these processes has been a long-time dream of researchers, it has only been in the last decade that neuroscience researchers have been able to go inside the brain and observe how learning actually occurs at the molecular level. New technologies like diffusion imaging have opened up the brain's inner workings and allowed scientists to "see" what is going on inside the brain when people are engaged in learning. The last decade has also seen significant developments in deep learning which have enabled researchers to encode data characteristics in an artificial neural network and generate similar looking examples as provided in the training data. Using these networks, we are even able to control the characteristics that we want our generated examples to have. Given the advancements in both of these fields, there cannot be a better time than now to delve deeper and try to answer some basic questions that are central to the process of knowledge acquisition and learning.

1. How can brain organize information that it sees, hears, or feels? Does there exist some mapping of semantic concepts to the areas of the brain? How does the brain organize these concepts? Does it do it by size? Relevance to living things? Or is it just a random mess? Since there is no finite list of concepts, how does the brain map (and learn) from simple concepts to higher level complex concepts?
2. How is sensory input converted to information and knowledge, and how is it stored efficiently to be used later?
3. How do we build an ever increasing network to deal with the problem of learning ever increasing concepts, or the networks with the fixed architecture are sufficient enough?
4. How do we extend the existing learning to other categories with few examples? For example, given a model trained to generate images of human faces with brown hair, if we show it an image of an animal, can it generate animal with brown hair? How does it know “young book” does not make sense?

Title: Strengthening Learning by Prior Knowledge
Speaker: Ansif Arooj
Abstract: Forming new associations is a fundamental process of building our knowledge system. On recent cognitive neuroscience literature, we are able to identify, multiple components of memory are activated and having effect while having prior knowledge; It was noted that less familiar stimuli are more difficult to combine to create new knowledge and that this is because less familiar stimuli consume more working memory resources. These neuroscience’s researches have left many questions for the world of machine learning as: Having prior knowledge can strengthen the machine learning abilities?

Title: Man vs. Machine: Can machines really ‘learn’?
Speaker: Lydia Braunack-Mayer

Abstract: According to modern discourse, the development of machine learning is a critical step towards artificial intelligence. And, indeed, the names 'Machine Learning', 'Neural Network' and 'Deep Learning' suggest similarities between a computational method and the human ability to learn. But can a computational model acquire knowledge in the same way as a human? If humans learn by back propagation, by evaluating sensory input against feedback from a trainer, then achieving artificial intelligence may only be a matter of scale. Artificial intelligence could be constructed from a learner with access to a sufficiently complex dataset. However, testing input from the senses against feedback seems to be only one of the ways in which humans acquire knowledge. This casts doubt on the idea that a machine learning algorithm can truly 'learn'.

Title: How Brains Represent Thousands of Objects
Speaker: Arvind Agarwal
Abstract: For decades, neuroscientists have found parts of the brain that respond to specific objects - the fusiform face area (FFA) specializes in recognizing faces, the parahippocampal place area (PPA) becomes active when we see images of places. However there can't possibly be one such area for all the thousands of categories of objects that we recognize otherwise we'd soon run out of space. This work shows how different categories of objects are mapped in different parts of the brain. Experiments conducted on 5 human subjects by collecting their brain activity data indicate that the brain organizes concepts along a continuous map. We have seen similar advancements in the Deep Learning where we have been able to represent words as continuous representations distributed across several (but unknown) semantic dimensions. We have seen similar advancements in computer vision where we have been able to generate images by first providing some semantic meaning to these dimensions, and then controlling the value of those dimensions to generate meaningful and controllable images. However, in both of these areas (neuroscience or deep learning), it is not clear "how" the learning or the mapping happens. While there is a close connection between both fields, there are several open questions that are yet to be answered. For example, we know where information is represented in the brain, but not how it's represented. What the brain and network doing at the neuron level. They are able to map known concepts, but how do they map the higher level concepts e.g. man walking on the street.

Session-2: Computing in the Brain

Title: How neuroscience can influence Deep Learning
Speaker: Subhrajit Roy
Abstract: Machine learning allows machines to solve specific tasks without being explicitly programmed. In the recent past, it has provided the human society with means of practical image and speech recognition, self-driving cars, efficient web-search techniques, health predictions, etc. One school of Machine Learning researchers, taking inspiration from the intelligence and processing power of the human brain have proposed architectures termed as neural networks that mimic the structure and operation of the brain. The prime example of neural networks is Deep Neural Networks (DNN) that are being used by big companies such as IBM, Google, Facebook, etc., for solving various pattern recognition tasks.

The main goal of this discussion session is to talk about the impact that neuroscience might have on deep learning in future. Specifically, we would discuss whether including more bio-realistic models of neurons, synapses, and learning algorithms might be useful. For example, possible questions that can be discussed are:

1. Neurons in the brain communicate through spikes [1]. Do spiking neural networks hold potential for creating next generation low-power, high-latency and high performance deep networks? Discuss possible issues in training spiking neural networks and how to tackle them.

2. There is now plethora of evidence that neurons in the brain have nonlinear or active dendrites [2]–[4]. However, deep learning tends to ignore this. Will including models of neurons with non-linear dendrites increase the performance of deep networks?
3. It has been discovered that biological neural networks learn through both modification of the network architecture [4]–[6] and synaptic weights. However, deep learning tends to only implement the later. Will it be better to do both? If so, how?

For your reference, here are some articles that have tried to answer these questions [7]–[9].

Title: Differentiable Neural Computers
Speaker: Oliver Gäfvert: Differentiable Neural Computers
Abstract: Introduced in the paper "Hybrid computing using a neural network with dynamic external memory" by Graves et.al., Differentiable Neural Computers (DNCs) provide a way of incorporating more structure into artificial neural networks, something believed to be necessary due to their limited ability to represent variables and data structures and to store data over long timescales. Similar to the computational model of a Turing machine, a DNC consists of a controller, which is a recurrent neural network (usually an LSTM), and an external memory which the controller can learn how to read and write from. Training a DNC corresponds to learning an algorithm that manipulates the input data by reading and writing to the external memory. In this sense, a DNC is very similar to a conventional computer, where the controller can be thought of as the CPU and the memory as the RAM, but it is differentiable end-to-end. DNCs have proved to be able to learn much more complex tasks than what is possible with LSTMs, such as sorting lists, finding shortest paths in graphs and solving simple puzzles.

Session-3: Principles of Human Cognition for Next Generation Machine Learning

Title: Can the cognitive strength of deep learning enhance computing reliability?
Speaker: Anwesha Das
Abstract: Today's supercomputing systems and cloud computing systems have large scale events buried in text logs. With increasing scale and complexity of evolving systems fault tolerance is a challenge. Solutions based on machine learning, deep learning and natural language processing are being investigated in the context of log data mining. How well can we extract and leverage the information contained in the voluminous system logs? Can the cognitive strength of neuroscience enhance the reliability of computing systems? Do the fundamental principles of neuroscience always conform to the way technology evolves or computing can evolve in the future?

Some questions to be asked/discussed:

1. Human beings have an interesting quality. The memory is not just in the mind but also in the body, for example, say if I am near a mountain, where I have been as a child 15 years back, I may not remember the name or the exact location or when did I come here before, but my body says, "I have been here before, I feel I know this place well, I know how to go ahead".
   Machine Learning is like mind memory based on recent past known information. Can a machine learn to remember what it has forgotten or has not seen recently? The power is neither pattern recognition nor spatially/temporally/contextually related information but a more profound sensory perception coming from data seen any time in the past, the machine proposes the correlation. We need to understand the sensory cortex of the brain to enhance deep learning.
2. Neuroscience is brain-like. But is the evolution in technology really brain-like? Instead of augmenting to what has already evolved in the last few generations, can we rebuild the base better for a new outcome? For example, flying the plane (apart from kite flying) is different

from the evolution of the smart phone, (wired telecommunication systems → advancements in wireless networks → MAC protocols → cell phones → smart-phones)

3. There be implicit labeling as well as unconscious bias in learning. Can we enhance deep learning to incorporate both and see if they come handy in different computing environments, since requirements and assumptions may vary based on the characteristics of the systems?
4. Can there be cases where certain intrinsic properties of deep learning may not enhance the overall reliability of computing systems? Are there better ways to improve the cognitive strength of systems?

Title: Cognitive Processes in Neural Models

Speaker: Felix Putze

Abstract: Artificial neural networks do not only draw inspiration from fundamentals of neuroscience but also from several cognitive processes: Recurrent neural networks are equipped with special "memory cells". "Attention" mechanisms are used to learn relevant parts of a long feature sequence. What other cognitive processes (e.g. "episodic memory", "forgetting") provide potential inspirations to model beneficial aspects of human cognition? Are these concepts only a loose analogy or are they founded on plausible theories of cognition? Under which circumstances are we allowed to draw conclusions from the behavior of the machine learning models on human cognition?

Title: Modeling how humans process language

Speaker: Tapasya Patki

Abstract: Understanding human language better and enabling seamless conversation is a key goal for next-generation computing and machine translation algorithms. For example, Google's Neural Machine Translation system (GNMT) provides an advanced end-to-end language learning framework based on deep neural network techniques, which allows us to explore zero-shot translation for the first time --- a mechanism that can efficiently translate between arbitrary languages. Bridging the gap between how humans process language and teaching machines to do so is critical for this process. Current artificial neural network techniques, such as convolutional networks and reinforcement learning, cannot easily capture temporally irregular dependencies or semantic ambiguities. This poster will present details on how humans learn language, and how better computational models for predicting neural activity can be built using encoding/decoding as well as conceptual alignment. Such understanding of the human brain can help build better artificial neural network frameworks in the future.

References

[1] W. Gerstner and W. Kistler, "Spiking Neuron Models, " Cambridge University Press, 2002.

[2] L. L. Gollo, O. Kinouchi, and M. Copelli, "Active Dendrites Enhance Neuronal Dynamic Range," *PLOS Comput. Biol.*, vol. 5, no. 6, p. e1000402, Jun. 2009.

[3] D. Johnston, J. C. Magee, C. M. Colbert, and B. R. Cristie, "Active properties of neuronal dendrites," *Annu. Rev. Neurosci.*, vol. 19, pp. 165–186, 1996.

[4] P. Poirazi and B. W. Mel, "Impact of active dendrites and structural plasticity on the memory capacity of neural tissue," *Neuron*, vol. 29, no. 3, pp. 779–796, Mar. 2001.

[5] A. Knoblauch and F. T. Sommer, "Structural Plasticity, Effectual Connectivity, and Memory in Cortex," *Front. Neuroanat.*, vol. 10, 2016.

[6] R. Lamprecht and J. LeDoux, "Structural plasticity and memory," *Nat. Rev. Neurosci.*, vol. 5, no. 1, pp. 45–54, Jan. 2004.

[7] A. H. Marblestone, G. Wayne, and K. P. Kording, "Toward an Integration of Deep Learning and Neuroscience," *Front. Comput. Neurosci.*, vol. 10, 2016.

[8] P. O'Connor and M. Welling, "Deep Spiking Networks," *ArXiv160208323 Cs*, Feb. 2016.

[9] J. H. Lee, T. Delbruck, and M. Pfeiffer, “Training Deep Spiking Neural Networks using Backpropagation,” *ArXiv160808782 Cs*, Aug. 2016.

Reading Material

- The Where of What: How Brains Represent Thousands of Objects
- Machine Learning and the Language of the Brain
- Perception Science in the Age of Deep Neural Networks
- How Prior Knowledge Helps the Brain to Learn Something New
- An introduction to Generative Adversarial Networks (with code in TensorFlow)
- How the Brain Learns
- Deep learning with COTS HPC systems
- PANDA: Pose Aligned Networks for Deep Attribute Modeling
- Neuroscience meets Deep Learning
- Fault Tolerance in Distributed Neural Computing