aList of class projects for Introduction to Computational Neuroscience 2018

- The idea is to work on a project that will allow you to study in some depth a phenomenon or question related to any aspect of neuroscience.

- A good option is to come with a project on your own on a question that interest you. However, you will need to explain the idea to any of the teachers who will have to check it before giving any approval and give you feedback to make the project feasible and interesting.

- Ideally, you should make teams of 3 or 4 people.

Here is a list of potential projects:

1. **How neural circuits internalize environmental rhythms? A reinforcement learning study** (given: Vladislav, Marharyta...)

2. **Exploring and visualizing an EEG + personality dataset** (given: Yevheniia,...)

3. **Training a VR dog using real reinforcement learning**
In this project you use Unity Machine Learning Agents to teach a virtual dog to sit, follow a ball, fetch a stick and so on. Training happens real-time inside VR, the trainer uses control button ("clicker") to give reward for desired behavior. Shaping principle from animal training is used - you always reward 50% of cases which are closer to the desired behavior.

4. **Synchronizing people using BCI.** In this project you will have to set up two portable EEG helmets and test whether it is possible for two people synchronize their EEG waves using the feedback from their own brains.

5. **Creating a game similar to Mindflex Duel** (http://store.neurosky.com/products/mindflex-duel) **with Emotiv Epoc device (EEG).** This involves processing the EEG signal from Emotiv's electrodes and building some game application that can be controlled with the processed signal. (2 ppl working on EEG, 2 person developing the game and communication between two devices)

6. **(PSYC LAB) Analyzing pupillary light reflex -** data about pupillary reactions (speed and amplitude of pupillary contraction, speed of return) has been collected in the Psychology Lab. This project will study how pupil reaction is related to other physiological processes.

7. **Reconstructing and visualizing a human brain made from word embedding**
   The goal of this project is to use word embeddings to represent the words of different parts of the brain in a distributed manner, and visualize the relative position of those vectors. This project involves training a large corpus of text using word2vec, selecting the vectors for words associated to different brain areas, and visualize their locations in 3D using dimensionality reduction techniques such as tSNE. If successful the output of the project should produce a brain where the distance between two different brain areas will be given by the similarity in their semantic roles in the text. Different corpora of neuroscience texts will be used and compared.

8. Learning to attend.

Central idea: attention is the key for learning (both in brains and in machines). But how does the agent know where to attend? Here we propose that it also has to learn that through reinforcement learning. In particular, we propose that instead of only being able to select different actions or different inputs, it can also select different features from its own system (e.g. the different features of the deep neural network) to enhance them.

Motivation: Neurobiologically this idea is motivated by the fact that the same area of the brain that selects actions *also selects which features to attend*. Furthermore, this area also assigns reward to actions and to attentional targets. Hence, it seems that there is an unitary mechanism for selecting actions and attention, both modulated by reinforcement learning.

Also, this idea is motivated by the fact that to process information efficiently we use attention: think about being at a loud party and trying to discuss deep learning with a friend. If there would be no attentional mechanism, the conversation would be continuously interrupted by different sounds that are stronger than the voice of your friend. The fact that we are able to focus in such situations shows the effect of attention: objectively nothing changes in the input (the party is still loud), but your system has enhanced only the relevant channel. There are many other facts supporting the general idea that learning in brains is dependent on attention and that attention can enhance the processing of single features of the network.

Implementation: The project is not very well defined yet, but we have ideas for starting. For example, we know that in the brain attention enhances the responsiveness of the neurons coding for this feature. In the artificial neural network the enhancement would mean lowering the threshold of activation of specific nodes, but how to learn which nodes and by how much should be affected? The appropriate enhancement of specific nodes should emerge during "learning to attend", but so far our group has no experience with how to train attention this way. The students should start by picking an Atari game with sufficient level of different elements (e.g. Seaquest) and try to use a standard DQN network where the agent could instead of only being able to select different actions also select different features from its own system (e.g. the different features of the deep neural network) to change their threshold. We envision that for effectively selecting features of its own layers to be enhanced, it might be necessary to add another separate network that tries to learn the importance of each node of the original network given the current state.

Suitable for: We hope to continue with these ideas in the long term. Hence, this project is suitable for students who like to explore new architectures, are interested in how our brains solve cognitive tasks and are potentially interested in joining our group later.

9. Challenge: predicting arousals during sleep

https://www.physionet.org/challenge/2018/

10. Emergent of grid cells via deep reinforcement learning

Grid cells: http://www.scholarpedia.org/article/Grid_cells

Project Settings: Grid world where food is spawn in random area. The agent or actor start from specific place (home) and see nothing but his last move as input. After the agent find the food he need to go home as fast as possible. Can the agent learn that ?

Implementation: most of the implementation for the project is done but the agent has vision.

Your task:

1. Remove the vision
2. Modify the network, and train it.
3. Open the network and look for grid-cells if they emerge.
4. If they didn't repeat (2-4).

Skills needed: Python, tensorflow.

11. Testing and / or Improving biologically realistic deep learning models.

Advances in DL have been breathtaking, but we believe that the brain still has some tricks up its sleeves. In this project the students would use available code for biologically realistic deep learning models. These models have not been applied to many tasks and they have not been optimized for standard DL problems. Hence, the students could either take the code and apply it to some other DL problems (e.g. a subset of Imagenet; or try to even build a deep reinforcement learning system based on this code and try it on Atari) or / and they could take the code appart and try to optimize its performance on MNIST (as done in the original paper - can one improve the network to beat the performance reported there?)

Original paper with background about biologically realistic deep learning: https://elifesciences.org/articles/22901

Code: https://github.com/elifesciences-publications/Segregated-Dendrite-Deep-Learning

12. Neural subimage search

Explore the use of deep learning models for subimage search: given images A and B, is the object on image B also on image A. Presumably image A is much bigger than image B. The idea is to pass both images through pre-trained deep neural network (i.e. ResNet) to produce convolutional features for both (dense classification layers are discarded). Then convolve image B features over image A features to calculate the match score. The scores could be used to predict the location of image B on image A. Global max pooling over all scores could be used to predict existence of object B on image A. Your task is to code the model, come up with evaluation metrics and experiment with different layers etc. MS COCO could be used as a dataset, comparing if segmented object from one class can be used to find objects of the same class from other images.

13. Higher order networks

Generalized linear models benefit from having higher-order features in the input, i.e. x_i * x_j. Could neural networks also take advantage of such features? One way to express this is through weighted pairwise correlations: h_i = sum_j(x_i * x_j * w_ij). If you transform the formula a bit - h_i = x_i * sum_j(x_j * w_ij) - then it is easy to see that this could be implemented similarly to residual connections, only instead of addition there is multiplication. Your task is to implement correlation connections and test if they provide any advantage compared to state-of-the-art neural networks. MNIST, CIFAR-10 or ImageNet could be used as a dataset.

14. Attention convolutional networks

We know from neuroscience that attention is an important feature of brain computations. Attention has been also successfully used in machine translation and image captioning. But is it also useful in standard convolutional networks? The idea is to split the network into two branches before the last convolutional layer: main branch and attention branch. Attention branch applies one-two additional convolutional layers and then spatial softmax to convert values into probability distribution over feature map. Main branch is multiplied with this probability distribution and summed spatially. This produces a final feature vector of the size of number of filters in the last convolutional layer. This provides at least two potential benefits: (1) the network can be applied to any size of images, (2) the attention layer outputs could assist in interpreting the classification failures, i.e. you can see where the network attended. Your task is to implement this attention architecture and experiment with different configurations. MNIST, CIFAR-10 or ImageNet could be used as a dataset.

15. Mental activity classification with OpenBCI EEG

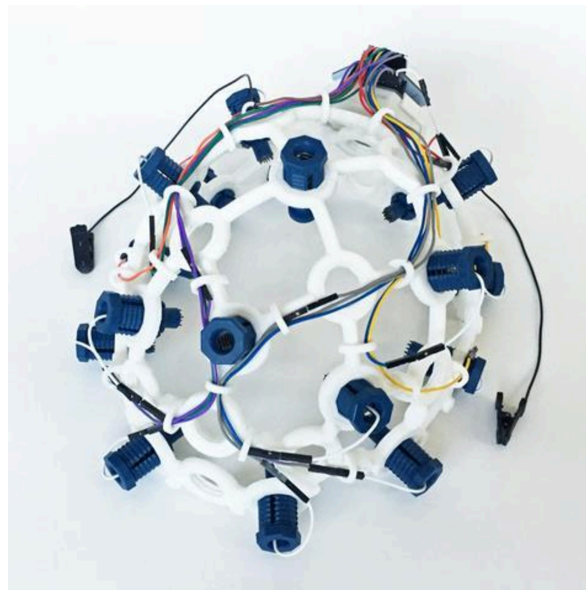

Am I thinking “left” or “right” at the moment?

Field of Brain-Computer Interfaces (BCI) uses neuroimaging and machine learning to answer such questions and use the emerging classifier to contral external devices.

Think “left” and computer cursor moves to the left, think “right” and it goes right.

Our lab worked on BCI before using Emotiv EPOC EEG device. That one was quite cheap and signal quality was poor. Now we have a OpenBCI 16-electrode EEG (4 times more expensive, much better signal) device in our lab. And we are looking for a person to take it for a spin.

The task is to read out the signal, record a dataset of person thinking “left” and “right” (or some other set of mental states) and train an ML model to measure how well we can tell them apart. Yes, it is not scifi and you can control an external device with sole power of your mind!

16. Surprising the predictive brain with teleportation in VR

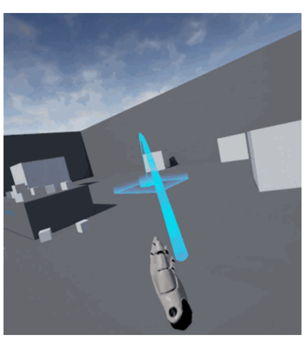

The brain has an amazing capability to adapt to new situations, and to do it fast! When thinking about new and never before seen experiences that we still manage to get used to with only a few trials, one in particular comes to mind - the act of teleporting in a virtual reality game. This is something no-one has ever done in real life, yet comes very natural in a virtual world. But is this always the case? Some pilot experiments have shown that when the person is moved to a new location via teleportation and also turned a few degrees in the process, the experience becomes much more disorienting and harder to predict. The same effect happens when the destination location is slightly changed from the predicted “landing spot”.

The task of the team is to come up and implement different modifications to the simple teleportation mechanic in VR to explore different ways the brain can or can not predict the results of teleportation. The performance of different interventions should be measured with a behavioural task (e.g. navigating a maze).

The work should be done with the help of a game engine (Unity or Unreal). Helpful starter templates exist. The computer graphics and virtual reality lab will provide needed workspaces and VR hardware (HTC Vive or Oculus Rift).

17. Challenge: can you find a DL architecture that can communicate like humans?
We humans are masters in communication and this is so easy for us that we tend to overlook the complexity of communication. In a recent paper it was claimed that the current progress is DL will not lead to agents that will and can communicate like us (see https://www.sciencedirect.com/science/article/pii/S1364661315002867 for the paper). These authors had developed a nice task that measures human communication skills without any language. See http://movie-usa.glencoesoftware.com/video/10.1073/pnas.1414886111/video-1 for the task. The task of the students would be to try to “solve” this task with currently available DL architectures. We anticipate that this will not be possible, as we would agree that current DL algorithms miss some components necessary for efficient communications. The students can discuss how to implement architectures that can solve this task and of course the students are more than welcome to try to create their own architectures (the grade for the project does not depend on actually being able to solve this task, as we do think that currently this is impossible)

> **Suitable for:** We hope to continue with these ideas in the long term. Hence, this project is suitable for students who like to explore new architectures, are interested in how our brains solve cognitive tasks and are potentially interested in joining our group later.

18. Genes and cognitive aging