Assignment #1 CPSC 533V Jan 2020
due in class, Tue Jan 21

1) Read the paper:
"*A critique of pure learning and what artificial neural networks can learn from animal brains*"
https://www.nature.com/articles/s41467-019-11786-6.pdf

   a) Give two implications (of the many possible implications) of the proposed point-of-view for reinforcement learning.
   b) Give an argument that goes at least partly against the given point of view.

   Keep your answers short, i.e., one paragraph for each implication and each argument. Be original!

2) Learning for Dynamics and Control (L4DC) is a workshop/conference that looks at topics lying at the intersection of control and reinforcement learning. https://l4dc.mit.edu/

   Recommended videos to watch from the first (2019) version of L4DC include:

   "John Tsitsiklis (MIT): "The Shades of Reinforcement Learning"
   "Emma Brunskill (Stanford University): "Efficient RL When Data is Costly"
   "Russ Tedrake (MIT): "Learning manipulation — and why I (still) like F=ma"
   "Sergey Levine (UC Berkeley): "Robots That Learn By Doing"
   "Angela Schoellig (University of Toronto): "Machine Learning for Robotics"

   Watch one of the talks, and answer the following questions:
   a) Briefly summarize 3 things that you learned with regard to RL and/or control. These could be related to any of history / open problems / theory / practice.
   b) List three concepts/assumptions/issues from the talk that you do not fully understand, i.e., that you would have difficulty explaining to someone else.