Neuro 140. Biological and Artificial Intelligence Spring 2023

List of Potential Projects

Project Title	Brief Description, Hypotheses, Questions	References	Difficulty level [0 = easy, 10 = hard]	Link to more information, data, code
Building models that generalize well. There are three different options: **1. Weather modifications** **2. Day-night** **3. Real vs Cartoons/Sketches**	Neural networks are notoriously bad at generalizing to test data which is significantly different from train data. Recent efforts try to work across such shifts. Implement some recent works, and try to suggest modifications which might do well.	https://openaccess.thecvf.com/content/ICCV2021W/AVVision/papers/Musat_Multi-Weather_City_Adverse_Weather_Stacking_for_Autonomous_Driving_ICCVW_2021_paper.pdf https://team.inria.fr/rits/computer-vision/weather-augment/	7	
The problem of parameters in linear systems	Current deep convolutional neural networks are typically underdetermined. Why is it that they do not overfit? Compute condition numbers, rademacher averages for underdetermined and overdetermined linear systems to assess robustness	Poggio, Kur, Banburski. Double descent in the condition number. https://www.pnas.org/content/117/44/27162	6	

Sharpened and faded object boundaries	It is widely known that CNNs are biased to textures rather than shapes. Taking a dataset with segmentation maps, sharpen or blur the edges in the training data. How does this impact the texture loving nature of CNNs?		7	http://www.image-net.org/
Impact of changing Transition Function in Deep RL 	(a) Use reinforcement learning to teach a network to play a video game like PACMAN. (b) Transition function defines the probabilities with which PACMAN ghosts move. How does the RL agent perform when the probabilities of ghost movements are different in testing than training?	https://github.com/tychovdo/PacmanDQN https://www.youtube.com/watch?v=QiIHGSYbjDQ	8	
Enforcing brain like activations	Recent works have trained linear models which take as input a CNN layer's activations and map them to neuronal activations collected from brain measurements. Reproduce these results and build on it.		5	
Graphical humor	Write an algorithm that will predict human judgments on whether an image is funny or not (or quantitative values on how funny an image is).		10	https://docs.google.com/document/d/1P3GvuU-YNAJI5qWF0HAJaZ8zVJ5-BVk1UZW-z9wKFDM/edit

				
Visual illusions 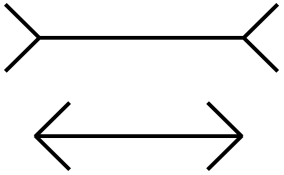	Are current computer vision systems susceptible to human visual illusions? How do CNNs see these images? Can we create more such images automatically?	Kreiman. The phenomenology of seeing. https://arxiv.org/pdf/1810.00415.pdf	8	https://robertmaxwilliams.github.io
Working memory 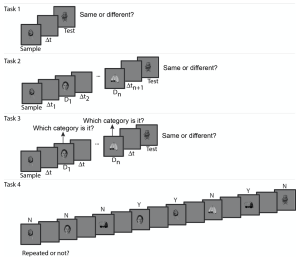	Create a model that can solve a variety of delay match to sample working memory tasks.	Miller. Working memory 2.0	9	https://docs.google.com/document/d/1oV3Asf-8qEqA96_QWvBH-qOZuYJhHp-u32Rpe_GpE5E/edit
Turing project	Test state-of-the-art algorithms as human imitators	Zhang et al. Human or machine? Turing tests for vision and language	4	https://drive.google.com/drive/folders/1vKplQUud271et4MPcqs8VXosGl_atOhP