Neuro 140. Biological and Artificial Intelligence 2022

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's 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	
Can GANs capture viewpoint and lighting variations?	Train a GAN on images of objects seen from different viewpoints and under different lighting conditions.		7	Contact TF for data
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	Poggio, Kur, Banburski. Double descent in the condition number.	5	

	underdetermined and overdetermined linear systems to assess robustness	https://www.pnas.org/content/117/44/27162		
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/
Training on deformations of meshes	3D objects are represented as meshes, and there are several methods to deform these meshes i.e. to modify their shape (http://www.open3d.org/docs/0.10.0/tutorial/Advanced/mesh_deformation.html) How do CNNs shape vs texture bias change when this is done?		7	
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=QilHGSYbjDQ	8	
Train your own NeRF	NeRF's are the new upcoming technique for		4	

	generating new, unseen views of objects. Collect some pictures of your favorite objects around the house/university and train a NeRF model on it. Some good questions to answer would be - how many images are needed? What kind of objects work well and what don't? What kind of viewpoints should be used in training for good performance?			
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.		4	
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	
Visual illusions	Are current computer vision systems susceptible to human visual illusions?	Kreiman. The phenomenology of seeing.	8	https://robertmaxwilliams.github.io

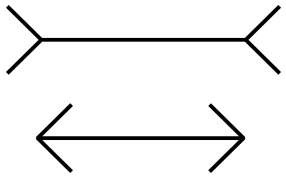	How do CNNs see these images? Can we create more such images automatically?	https://arxiv.org/pdf/1810.00415.pdf		
Working memory 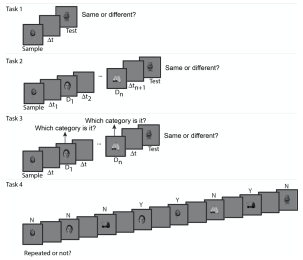	Create a model that can solve a variety of delay match to sample working memory tasks.	Miller. Working memory 2.0	9	