Neuro 140/240 Biological and Artificial Intelligence 2021

List of Potential Projects

Project Title	Brief Description, Hypotheses, Questions	References	Difficulty level [0 = easy, 10 = hard]	Link to more information, data, code
Recognition of doodles 	(1) Evaluate recognition of handmade drawings (2) Compare model performance versus human performance (3) Evaluate transfer learning between doodles and real objects	Krizhevsky A, Sutskever I, Hinton G (2012) ImageNet Classification with Deep Convolutional Neural Networks. In: NIPS. Montreal. Kreiman G and Serre T. The role of feedback in vision (To appear)	5	https://quickdraw.withgoogle.com/data
Face recognition with arduino 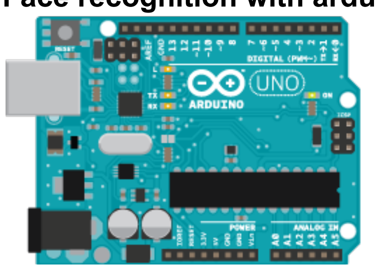	Build a face recognition system with arduino	Arduino tutorial	7	https://create.arduino.cc/
Landmark recognition with arduino	Build an arduino system that will recognize Harvard campus buildings. The system should be able to stand	Arduino tutorial	7	https://create.arduino.cc/

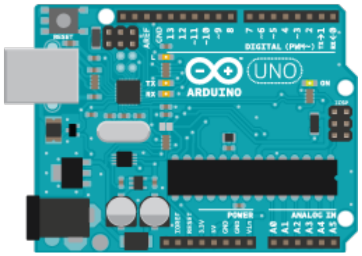	anywhere on campus and tell you where it is (using vision, not using GPS).			
Action recognition in controlled datasets 	Current action recognition datasets are easy due to confounding factors. Action recognition in controlled datasets can be challenging. Projects (1) Create controlled datasets and test them as in Jacquot et al 2020 (2) Evaluate state-of-the-art models of controlled datasets	[Soomro, Zamir, Shah, UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild] [Jacquot V, Ying J, Kreiman G. (2020) Can Deep Learning Recognize Subtle Human Activities? CVPR 2020.]	6	[https://www.crcv.ucf.edu/data/UCF101.php]
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 for underdetermined and overdetermined linear systems to assess robustness	[Poggio, Kur, Banburski. Double descent in the condition number.] [Casper S, Boix X, D'Amario V, Guo L, Schrimpf M, Vinken K, Kreiman G. (2021). Frivolous Units: Wider Networks are not really that Wide. AAAI Conference on Artificial Intelligence.]	5	

The role of color in visual recognition 	(1) Take a deep convolutional network (e.g. VGG16) pre-trained for object recognition (ImageNet) and test it after stripping away color from the images (2) Train and test network with grayscale images	[Krizhevsky A, Sutskever I, Hinton G (2012) ImageNet Classification with Deep Convolutional Neural Networks. In: NIPS. Montreal.]	3	[http://www.image-net.org/]
Size and position invariance in visual recognition COCO Common Objects in Context	(1) Take a deep convolutional network (e.g. VGG16) pre-trained for object recognition (e.g. MSCOCO) and test it after changing object positions (2) Same as (1), but changing object scales	[Singh, Davis. An analysis of scale invariance in object detection]	4	[http://cocodataset.org/]
Continual learning for video games 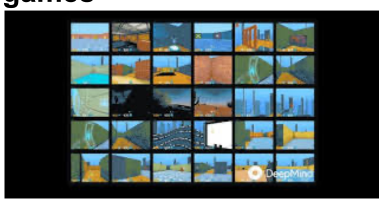	(a) Use reinforcement learning to teach a network to play a video game (b) Teach the network to teach a second video game without forgetting the previous one and without retraining on the previous one!	[Vinyals, Babuschkin, […] Silver Grandmaster level in StarCraft II using multi-agent reinforcement learning] [Kirkpatrick, …, Hadsell Overcoming catastrophic forgetting in neural networks]	7	
Create your own object recognition challenge	(1) Consider 2 or more object categories (2) Collect images (3) Train models (4) Test models e.g. (i) app that takes a picture and tells you where you are on campus, (ii) emotion recognition in faces, (iii) from short videos of favorite sport, predict winner (without looking at the score, purely from frames)		8	

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 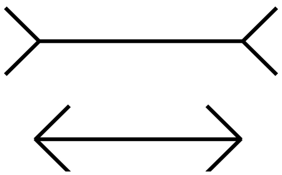	Are current computer vision systems susceptible to human visual illusions?	[Kreiman. The phenomenology of seeing.]()	8	
Working memory 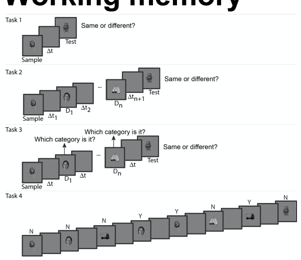	Create a model that can solve a variety of delay match to sample working memory tasks	[Miller. Working memory 2.0]()	9	
Pattern completion	Create an algorithm that can recognize occluded objects	[Tang et al. PNAS 2018]()	8	http://cocodataset.org/

				
Object Tracking	Create an algorithm that can continuously monitor the position (and orientation) of a single object or multiple objects in multiple consecutive frames of a video.		8	https://motchallenge.net/
Object Tracking: Object Files	Replicate a canonical visual cognition result from human visual psychology in a deep convolutional neural network -- training your network to create object representations that more closely match those of humans.	http://www.nyu.edu/gsas/dept/philo/courses/readings/2016.green.qd.pdf	9	
Play Pictionary with Recurrent Networks	Learn what constitutes a minimally recognizable visual concept using one or a combination of Sketch-RNN; Quick, Draw; or Generative Adversarial Network. Explore the learned representations to see whether the machine's minimal concept match naive human observers!	https://arxiv.org/pdf/1704.03477.pdf	7	https://magenta.tensorflow.org/assets/sketch_rnn_demo/index.html
Train Networks Using Evolutionary Algorithms	Assess the ability of an evolutionary algorithm to master standard computer vision tasks like object recognition.	https://openai.com/blog/evolution-strategies/	7	
Analysing Representations: Comparing Tasks	Use pretrained models and any one of various model-agnostic analyses (representational similarity analysis, canonical correlation analysis or principal components regression) to compare representations learned by networks trained on different tasks.	http://taskonomy.stanford.edu/	2.5	

Analysing Representations: Comparing Architectures	Use pretrained models and any one of various model-agnostic analyses (representational similarity analysis, canonical correlation analysis or principal components regression) to compare representations learned by different network architectures on the same task.		1	
Train Multiple Tasks without ‘Catastrophic’ Forgetting	Train the same model on multiple different tasks in such a way that allows the model to maintain high performance in each class. Use a dynamic memory bank or a modified training regimen.		3-5	
Train a Network to Recognize Compositional Objects 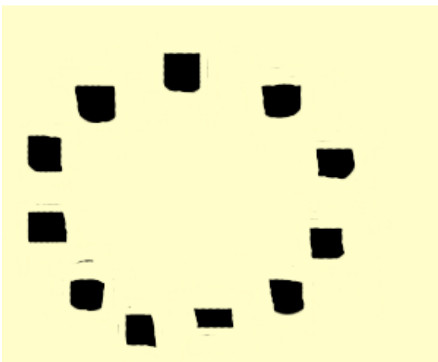	Most humans would call the image to the left a ‘circle’. Most computer vision algorithms would label this image a square. Can you train an algorithm to recognize this as a circle made of squares (while still preserving the individual labels of circle and square)?		5	
Hierarchical Categorization	A chihuahua is not just a chihuahua. A chihuahua is also a dog is also a mammal is also an animal is also an organism. What happens when you train a model to learn multiple hierarchical labels for the same input? What sorts of representations does it learn?		2	
Analyzing Representations: Models Trained to Match Brain	Much of the modeling that happens at the intersection of neuroscience and artificial intelligence involves		1	

Data	training a model on a certain task (like object recognition) and testing whether the representations it learns match those of the brain. What happens if you reverse this process? Does a model trained to match the brain do better at object recognition than an untrained baseline?			
Modeling Perceptual Bistability 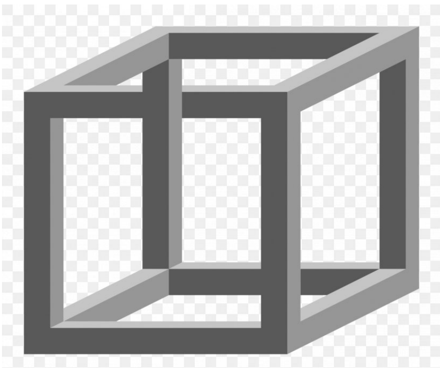	The image to the left is called a necker cube. It's an example of a 'bistable' stimulus -- a stimulus you can interpret one of two ways, but (importantly) not both at the same time. Perceptually bistable illusions suggest that the two halves of the brain (especially in the visual system) are sometimes engaged in an opponent process, sometimes competing, sometimes taking over for the other when one becomes fatigued. How could you model this? Adversarial relationships exist in neural network modeling, as well. Could there be a relationship?		4	
Perform 'Deep Neuroethology' on a Virtual Agent	Train a virtual agent in an environment of your choosing, then analyze the neural basis of its learned behaviors.	https://arxiv.org/pdf/1911.09451.pdf	9	
Explore Modularity in Neural Networks	The brain is often described as 'modular' -- different parts with different functional capacities working in tandem to produce sophisticated behavior. Classic modules include the 'fusiform face area' (for face processing) in visual cortex and broca's area (for the production of language) in auditory cortex. Under what sorts of constraints would you expect a computational system to produce modularity? How would you replicate those constraints in silicon?	https://plato.stanford.edu/entries/modularity-mind/	4	

Use TensorflowJS to Explore Gestural Communication through Games	Build a game (following a well-established example, linked on the right) that allows users to use their webcams and personalized gestures as controllers for a game. Collect a dataset of these personalized gestures to see if there exist commonalities across the gestures people use to enact certain types of control.		8	https://storage.googleapis.com/tfjs-examples/webcam-transfer-learning/dist/index.html
Analyzing Representations: Representations in Depth	How should we characterize the evolution of representations over the functional transformations they undergo across the layers of a neural network. Which operations push representational spaces closer together? Which operations pull them apart? If an operation is repeated in serial, how does that impact the representations? If an operation is drastically reduced (by ablating many or most of the weights to the operating layer) how is the final product of the model affected?		1	
Adventures in Artificial Ablation (The Neural Network Neuropsychologist)	Because artificial neural networks are artificial, you can open them up and play around with their nervous systems in ways that would be ethically questionable in live animals. Our ability to artificially ablate artificial neurons (effectively by setting all weights from that neuron to zero), we've observed that a significant amount of a neural network can be pruned while maintaining the model's overall functionality. What else might we learn?		3	An example in Tensorflow by a former student, Will Bryk: https://artificialneuroscientist.herokuapp.com/
Create Your Own Neural Network Module	Convolutions, rectified linearities, pooling… These are some of the most popular neural network 'modules' available in modern deep learning libraries. But they're far from the only ones possible. Using the	https://pytorch.org/docs/stable/nn.html#	3	

	brain as inspiration (or any other model system) create your own neural network module and plug it in. You never know what you might find.			
Compete in the AnimalAI Olympics!	Train an artificial agent to complete tasks many animals do with ease!		9	http://www.animalaiolympics.com/AAI/
Explore Tradeoffs Between Shape and Texture in Object Recognition Models	For all their success in object recognition, neural networks seem to be biased heavily towards texture. Explore why this bias emerges and whether you can shift the bias towards shape.	https://arxiv.org/abs/1811.12231	1	https://github.com/ColinConwell/ShapeCentral