Deep Learning for Autonomous Systems Seminar WS 2020/21

Robot Learning Lab

Albert-Ludwigs-Universität Freiburg

Friday, November 6th 2020

Procedure

http://rl.uni-freiburg.de/teaching/ws20/deeplearningforautonomoussystems

- Students should select three papers out of the list in preference order (highest first).
- Places will be assigned based on priority suggestions of HisInOne and motivation of student by Nov 25, 2020.
- Students are required to prepare a 20 minutes talk, write an abstract and a summary.
- The Seminar will be held as a virtual ”Blockseminar” on Feb 5, 2021.

Procedure

- The details of presentation and slides should be discussed with the supervisor two weeks before the presentation.
- Abstract should be 2 pages long and is due on Jan 08, 2021.
- Summary is due on Feb 19, 2021 and should be max. 7 pages long (latex, a4wide, 11pt) not including the bibliography and figures. Significantly longer summaries will not be accepted.
- The final grade is based on the oral presentation, the written abstract, the summary, and participation in the blockseminar.

Deep Learning For Autonomous Systems

- Deep learning has led to impressive progress on complex, high dimensional data
  - Speech Recognition
  - Computer Vision
  - Natural Language Understanding
- Now enable autonomous systems and robots to operate in the real world

Perception

- Complex environments

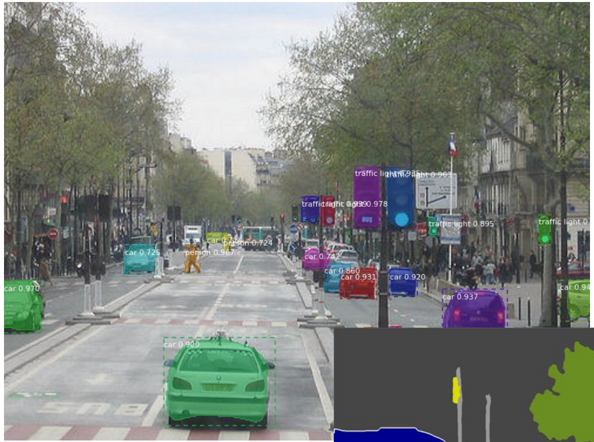

- Noisy observations and sensors

Mask R-CNN for object detection and instance segmentation on Keras and TensorFlow, Waleed et. al., 2017
Rohit Mohan and Abhinav Valada, "EfficientPS: Efficient Panoptic Segmentation", arXiv preprint arXiv:2004.02307, 2020.

Unknown, Open World

- Open Set Recognition: recognise unknowns
- Uncertainty estimation
- Adversarial attacks

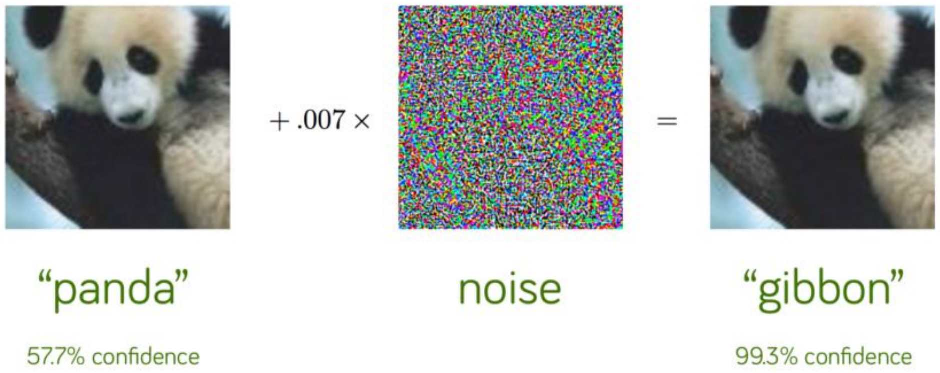

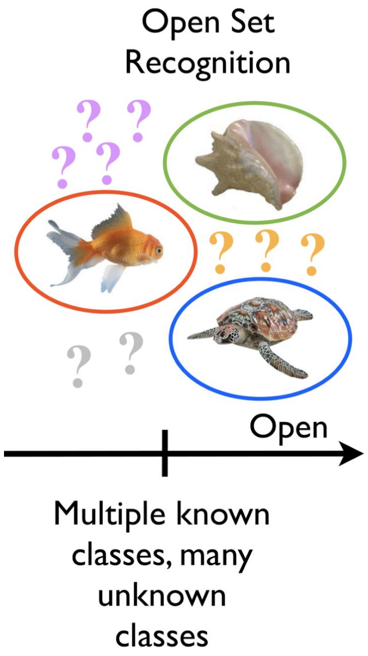

Towards Open Set Recognition, Scheirer et. al., 2012

Explaining and Harnessing Adversarial Examples, Goodfellow et. al., 2014

Autonomous Decision Making

Reinforcement learning for short- and long-term decision making

Slide from Raia Hadsell

Continuous Control

- Model free RL successful on difficult continuous control domains
  - Directly optimise policy
  - Comparably data efficient on stationary tasks
- Model based RL catching up
  - Learn a world model
  - Promise of better generalisation

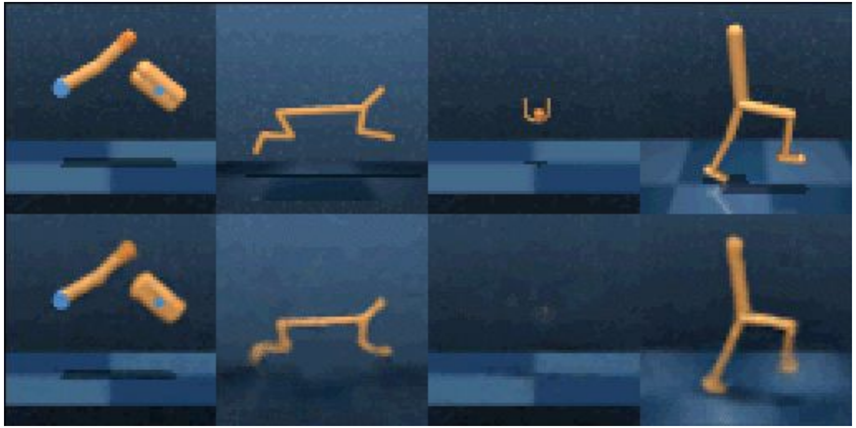

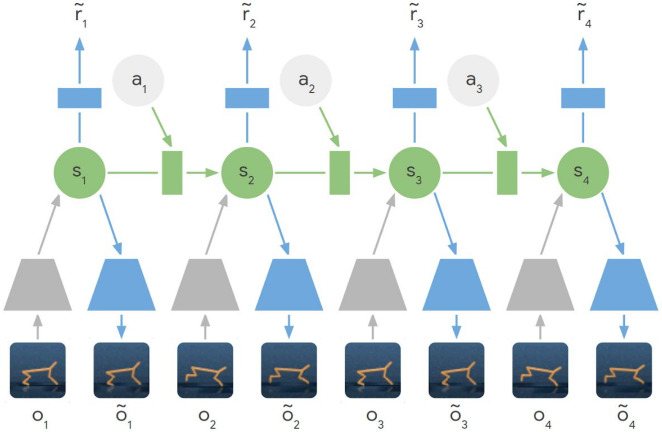

Learning Latent Dynamics for Planning from Pixels, Hafner et. al., 2019

Expensive Real World Data

- Sim2Real
  - Domain Adaptation
  - Action and dynamics noise
- Offline RL
  - Large amounts of unstructured data
  - Little annotated / expert data

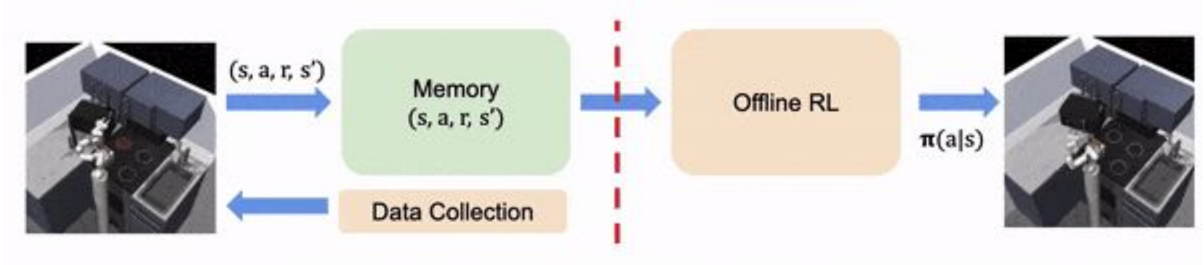

Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection. Sergey Levine, Peter Pastor, Alex Krizhevsky, Deirdre Quillen
D4RL: Datasets for Deep Data-Driven Reinforcement Learning, Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, Sergey Levine

Self- and Weak-Supervision

- Pretext tasks
  - Object presence and absence
  - Consistency over time
  - Viewpoint invariance
- Reduce oversight
  - Automatic resets
  - Reward labelling

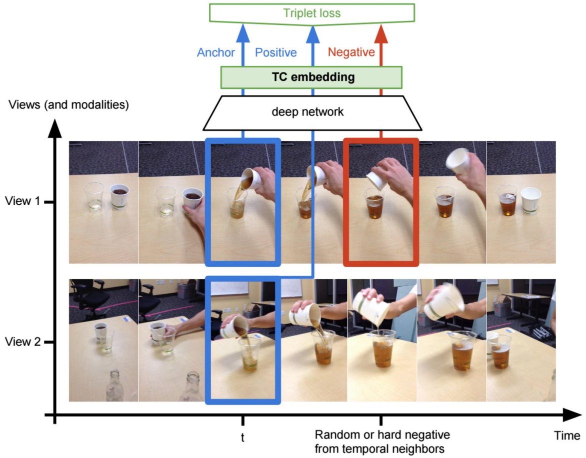

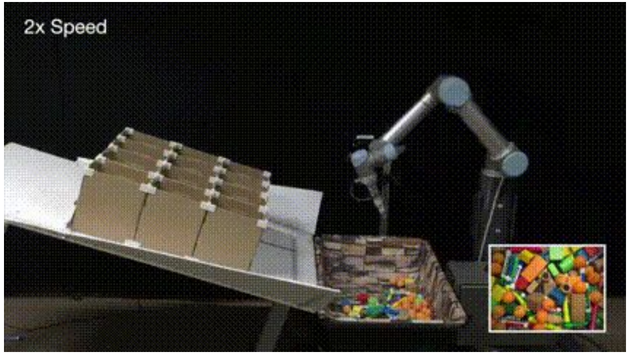

Time-Contrastive Networks: Self-Supervised Learning from Video, Sermanet et. al., 2018
TossingBot: Learning to Throw Arbitrary Objects, Zeng et. al., 2019.

Seminar Topics

Topic 1: End-to-End Robotic Reinforcement Learning without Reward Engineering

- Real-world applications of reinforcement learning must specify the goal of the task by means of a manually programmed reward function
- This work enables a robot to learn from a modest number of examples of successful outcomes, followed by active queries, where the robot shows the user a state and asks whether that state represents success.

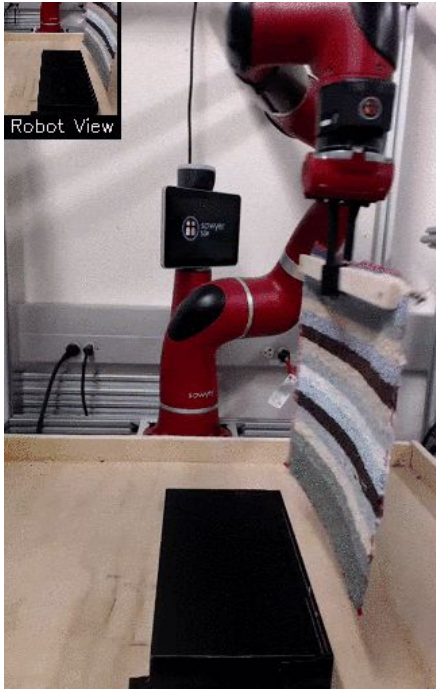

Source: Singh et al. 2019

Supervisor: Eugenio Chisari - Paper link: https://arxiv.org/abs/1904.07854

Topic 1: End-to-End Robotic Reinforcement Learning without Reward Engineering

- First, learn reward function based on goal images and labels specified by users.
- Then train an RL agent on a task based on this reward function. To avoid undesired behaviours the robot periodically queries the user to provide labels for images.

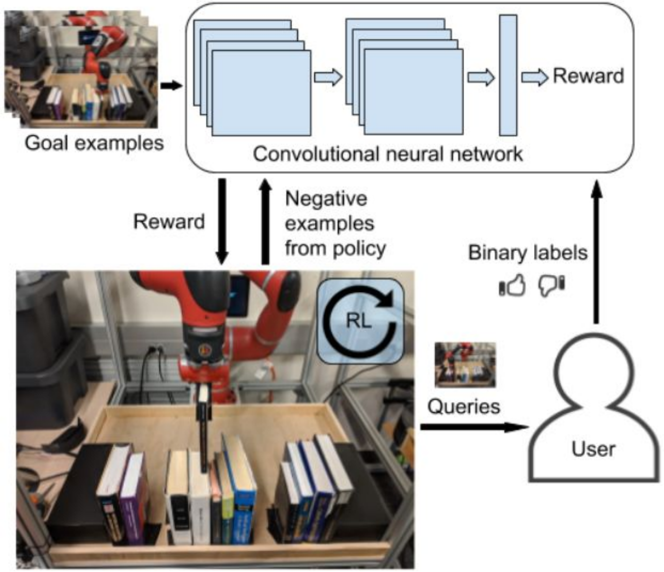

Source: Singh et al. 2019

Supervisor: Eugenio Chisari - Paper link: https://arxiv.org/abs/1904.07854

Topic 2: Learning quadrupedal locomotion over challenging terrain

- Quadrupedal locomotion is in princip very powerful. But achieving the generality and robustness of animal locomotion across diverse environments is very challenging.
- Goal: learn robust locomotion across very challenging terrains.

Source: Lee et al. 2020

Supervisor: Eugenio Chisari - Paper link: https://arxiv.org/abs/2010.11251

Topic 2: Learning quadrupedal locomotion over challenging terrain

- Rely only on proprioceptive measurements from joint encoders and an inertial measurement unit (IMU), the most durable and reliable sensors
- Distill a priviledged teacher policy
- Synthesise terrains that follow an adaptive difficulty schedule

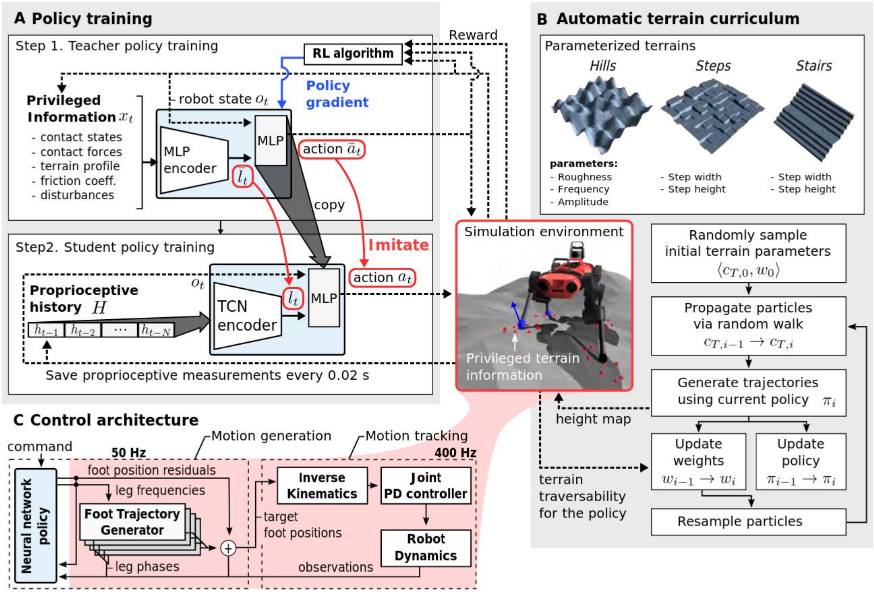

Source: Lee et al. 2020

Supervisor: Eugenio Chisari - Paper link: https://arxiv.org/abs/2010.11251

Topic 3: One Policy to Control Them All: Shared Modular Policies for Agent-Agnostic Control

- Most RL approaches learn policies specific to a particular agent. Can we instead learn a single global policy that generalises to a wide variety of agents?
- Assume each actuator as its own agent, sharing the same network and pass messages to propagate information

Supervisor: Daniel Honerkamp - Paper link: https://arxiv.org/pdf/2007.04976.pdf

Topic 3: One Policy to Control Them All: Shared Modular Policies for Agent-Agnostic Control

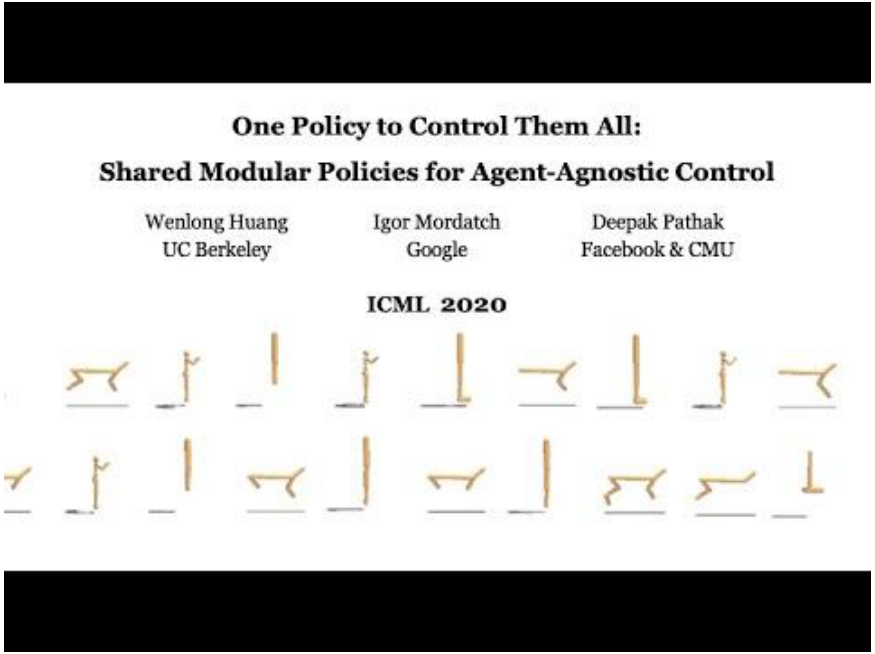

Source: Huang et al. 2020

Supervisor: Daniel Honerkamp - Paper link: https://arxiv.org/pdf/2007.04976.pdf

Topic 4: Opal: Offline Primitive Discovery for accelerating offline Reinforcement Learning

- Offline RL: an agent may have access to large amounts of undirected offline experience data, while access to the online environment is severely limited.
- Idea: extract a continuous space of recurring and temporally extended primitive behaviors before using these primitives for downstream task learning.

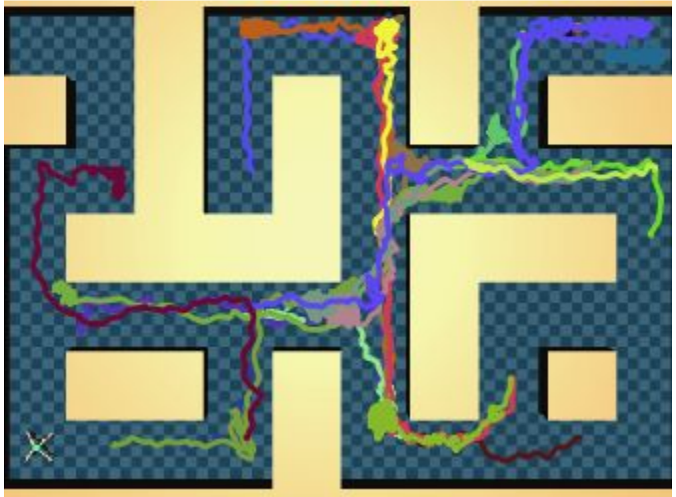

Source: Ajay et al. 2020

Supervisor: Daniel Honerkamp - Paper link: https://arxiv.org/pdf/2010.13611.pdf

Topic 4: Opal: Offline Primitive Discovery for accelerating offline Reinforcement Learning

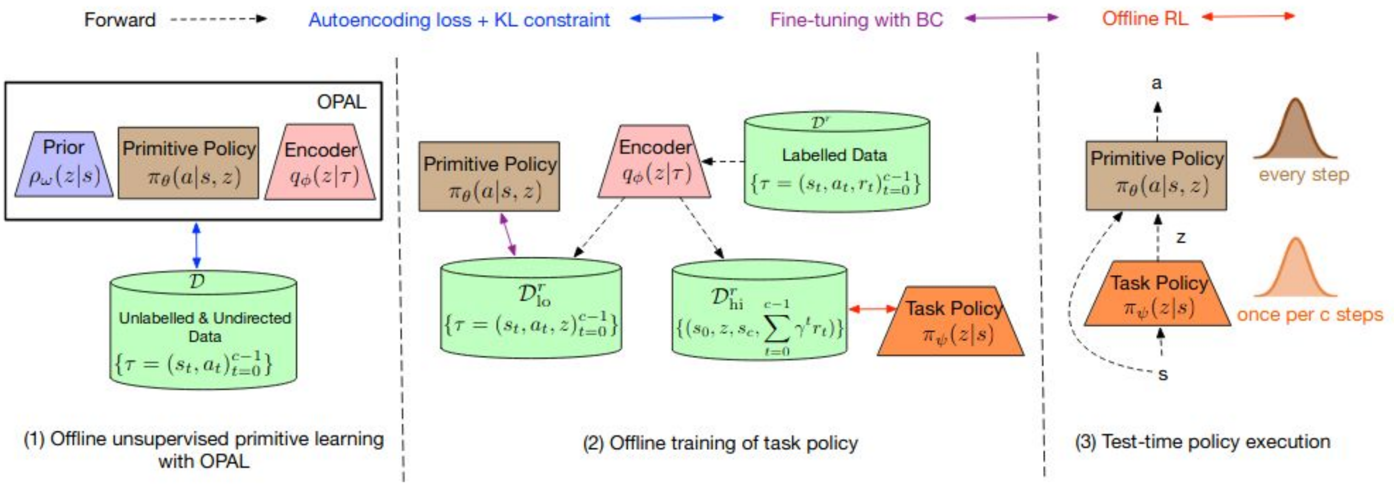

Source: Chen et al. 2019

Supervisor: Daniel Honerkamp - Paper link: https://arxiv.org/pdf/2010.13611.pdf

Topic 5: Euclideanizing Flows: Diffeomorphic Reduction for Learning Stable Dynamical Systems

- Robotic tasks often require motions with complex geometric structures.
- Goal: Learn such motions from a limited number of human demonstrations by exploiting the regularity properties of human motions e.g. stability, smoothness, and boundedness

Supervisor: Dr. Tim Welschehold - Paper link: https://arxiv.org/pdf/2005.13143.pdf

Topic 5: Euclideanizing Flows: Diffeomorphic Reduction for Learning Stable Dynamical Systems

- Instead of explicitly learning a stable dynamical system, view demonstrations as motions on a manifold which is linked, under a smooth bijective map, to a latent Euclidean space
- This results in an expressive class of diffeomorphisms suitable for learning stable and smooth dynamical systems

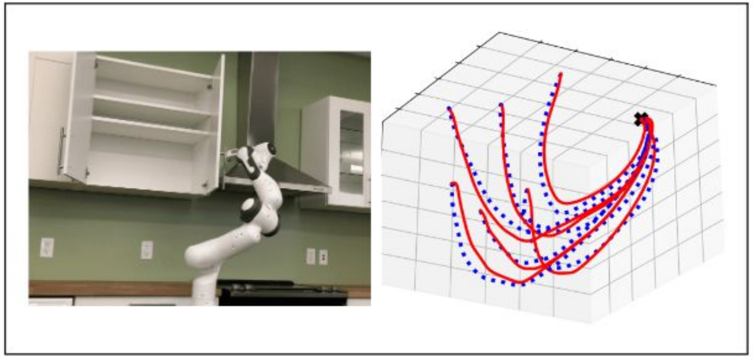

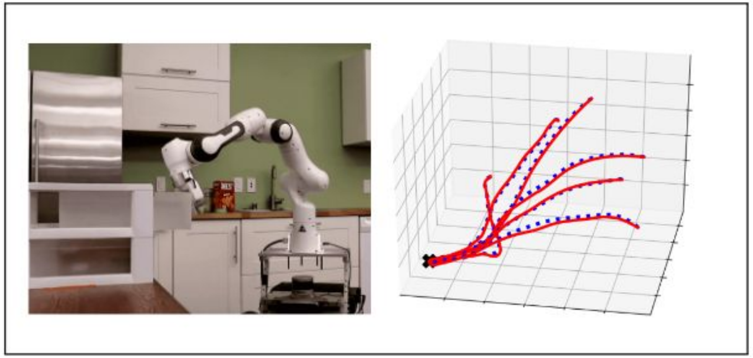

Supervisor: Dr. Tim Welschehold - Paper link: https://arxiv.org/pdf/2005.13143.pdf

Source: Rana et al. 2020

Topic 6: Neural Dynamic Policies for End-to-End Sensorimotor Learning

- Training policies directly in raw action spaces forces the agent to make decisions at each point in training, limiting its scalability to complex tasks
- Dynamical systems used in classical robotics lack the flexibility of deep learning
- Idea: reparameterize action spaces with differential equations to embed dynamics structure into NNs

Supervisor: Dr. Tim Welschehold - Paper link: https://biases-invariances-generalization.github.io/pdf/big_15.pdf

Topic 6: Neural Dynamic Policies for End-to-End Sensorimotor Learning

- Reparameterize the action space with non-linear differential equations corresponding to a dynamical system, train it end-to-end.
- ‘Deep’ part of the policy only needs to reason in the lower-dimensional space of building a dynamical system, so overall policy can easily reason in the space of trajectories.

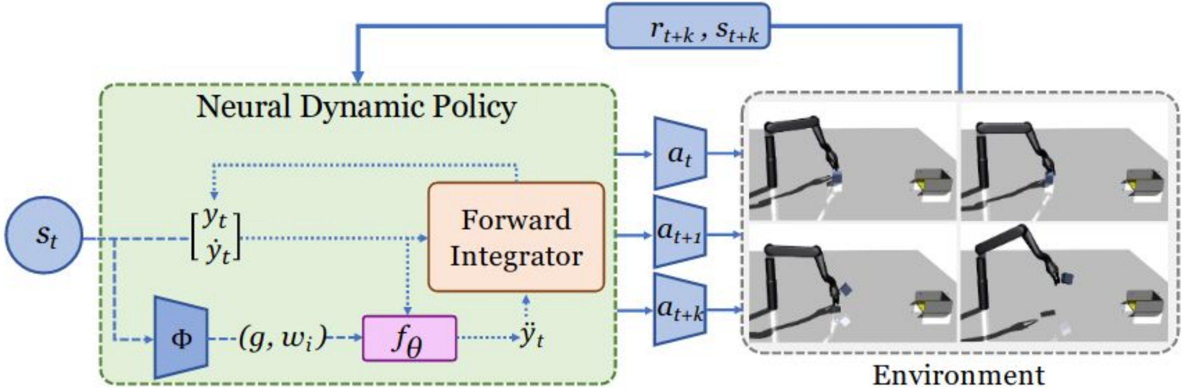

Supervisor: Tim Welschehold - Paper link: https://biases-invariances-generalization.github.io/pdf/big_15.pdf

Topic 7: Consistent Video Depth Estimation

- Reconstruct dense, geometrically consistent depth for all pixels in a monocular video
- Leverage a conventional structure-from-motion reconstruction to establish geometric constraints on pixels in the video and a learning-based prior, i.e., a convolutional neural network trained for single-image depth estimation
- At test time, we fine-tune this network to satisfy the geometric constraints of a particular input video

Supervisor: Nikhil Gosala - Paper link: https://arxiv.org/pdf/2004.15021.pdf

Topic 7: Consistent Video Depth Estimation

- Conventional approaches: Incomplete depth on moving objects
- Learning based: flickering and geometrically inconsistent

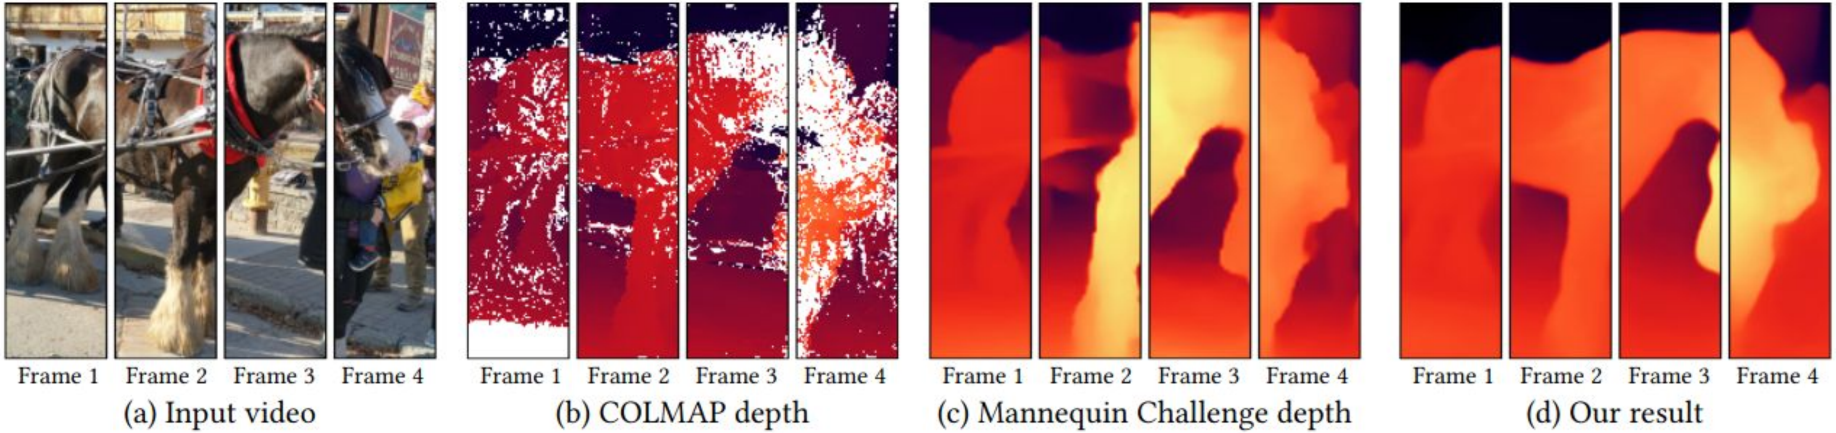

Source: Luo et al. 2020

Supervisor: Nikhil Gosala - Paper link: https://arxiv.org/pdf/2004.15021.pdf

Topic 8: Learning Depth with Very Sparse Supervision

- Can a three dimensional perception system be trained with the data that a robot would observe interacting with the environment?
- Novel global-local network architecture that takes images and extremely sparse depth measurements, down to even a single pixel per image

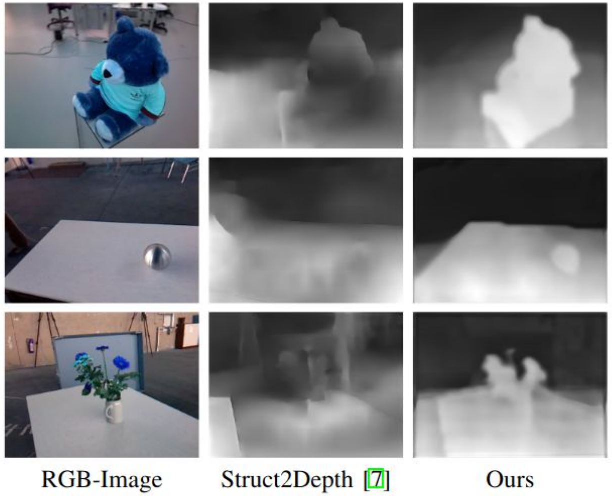

Source: Loquerico et al. 2020

Supervisor: Nikhil Gosala - Paper link: http://rpg.ifi.uzh.ch/docs/IROS20_Loquercio.pdf

Topic 8: Learning Depth with Very Sparse Supervision

- From flow and images estimate global parameters *g* representing camera motion
- Local module applies them to the optical flow field to generate final depth estimates

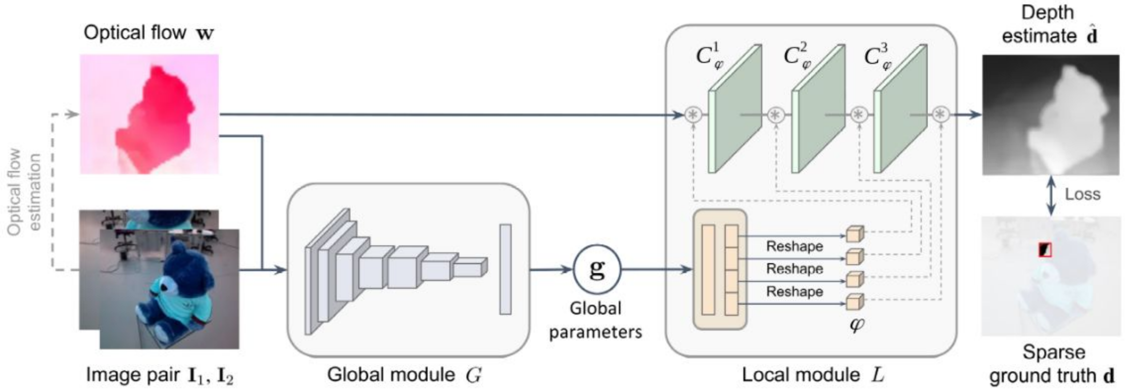

Source: Loquerico et al. 2020

- Strong results given as little as a single depth ground truth pixel

Supervisor: Nikhil Gosala - Paper link: http://rpg.ifi.uzh.ch/docs/IROS20_Loquercio.pdf

Topic 9: ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision and Language Tasks

- Visual understanding: corresponding understanding or *grounding* between vision and language
- Dominant strategy: start with separate pretrained language and vision models pretrained for other large-scale tasks – often resulting in myopic groundings
- New: joint model for learning task-agnostic visual grounding from paired visiolinguistic data

Supervisor: Juana Valeria Hurtado - Paper link: https://papers.nips.cc/paper/8297-vilbert-pretraining-task-agnostic-visiolinguistic-representations-for-vision-and-language-tasks.pdf

Topic 9: ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision and Language Tasks

- ViLBERT: Vision-and-Language BERT
- Separate streams for vision and language processing that communicate through co-attentional transformer layers
- Training on proxy tasks: predicting semantics of masked words and image regions given the unmasked inputs, and predicting whether an image and text segment correspond

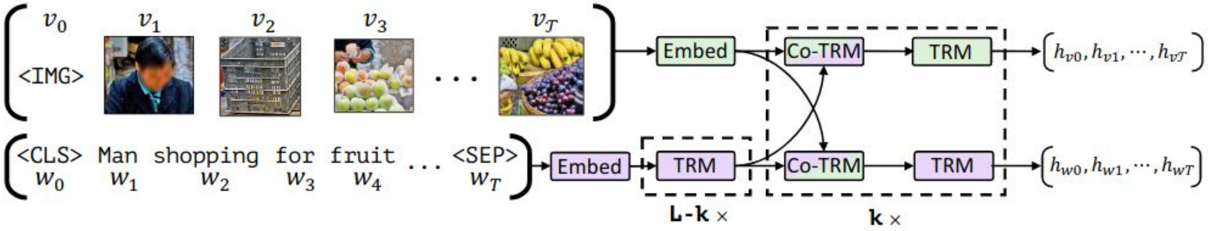

Source: Lu et al. 2020

Supervisor: Juana Valeria Hurtado - Paper link: https://papers.nips.cc/paper/8297-vilbert-pretraining-task-agnostic-visiolinguistic-representations-for-vision-and-language-tasks.pdf

Topic 10: Embodied Visual Recognition

- Passive visual systems typically fail to recognize objects in the amodal setting where they are heavily occluded
- Embodied visual recognition: agent is free to move in the environment to perform object classification, amodal object localization, and amodal object segmentation.

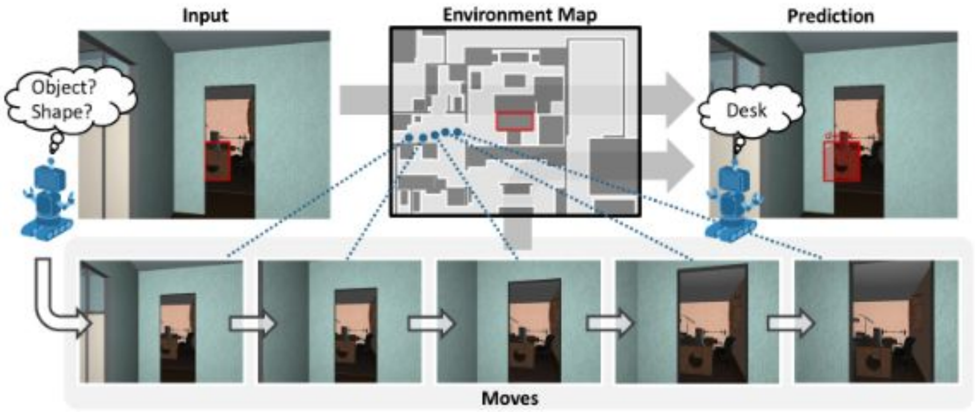

Source: Yang et al. 2019

Supervisor: Juana Valeria Hurtado - Paper link: https://arxiv.org/pdf/1904.04404.pdf

Topic 10: Embodied Visual Recognition

- Embodied Mask R-CNN to learn to move for visual recognition.
- Make predictions at each step with the aim to improve visual recognition performance on the target object in the first frame.
- Learned moves are different from shortest-path moves and generalize well to unseen environments

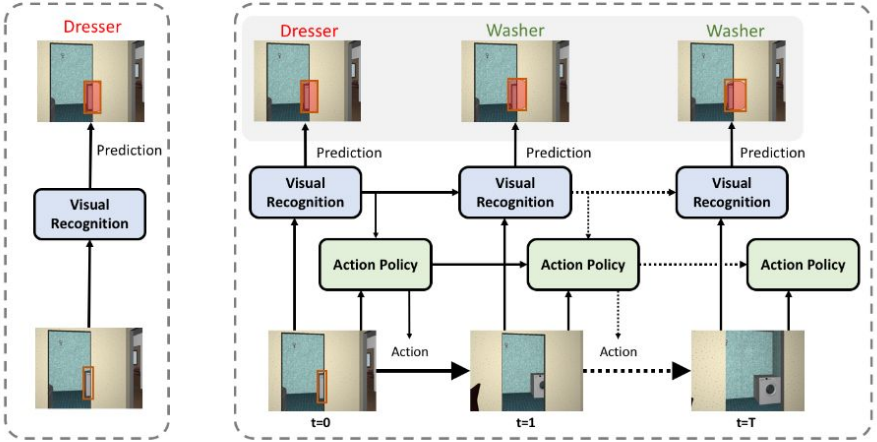

(a) Passive Visual Recognition (b) Embodied Visual Recognition

Source: Yang et al. 2019

Supervisor: Juana Valeria Hurtado - Paper link: https://arxiv.org/pdf/1904.04404.pdf

Topic 11: All about Structure: Adapting Structural Information across Domains for Boosting Semantic Segmentation

- Unsupervised domain adaptation for semantic segmentation: transfer knowledge learned upon synthetic datasets with ground-truth labels to real-world images without any annotation
- Disentangle images into domain-invariant structure and domain-specific texture representations

Supervisor: Dr. Daniele Cattaneo - Paper link: https://arxiv.org/abs/1903.12212

Topic 11: All about Structure: Adapting Structural Information across Domains for Boosting Semantic Segmentation

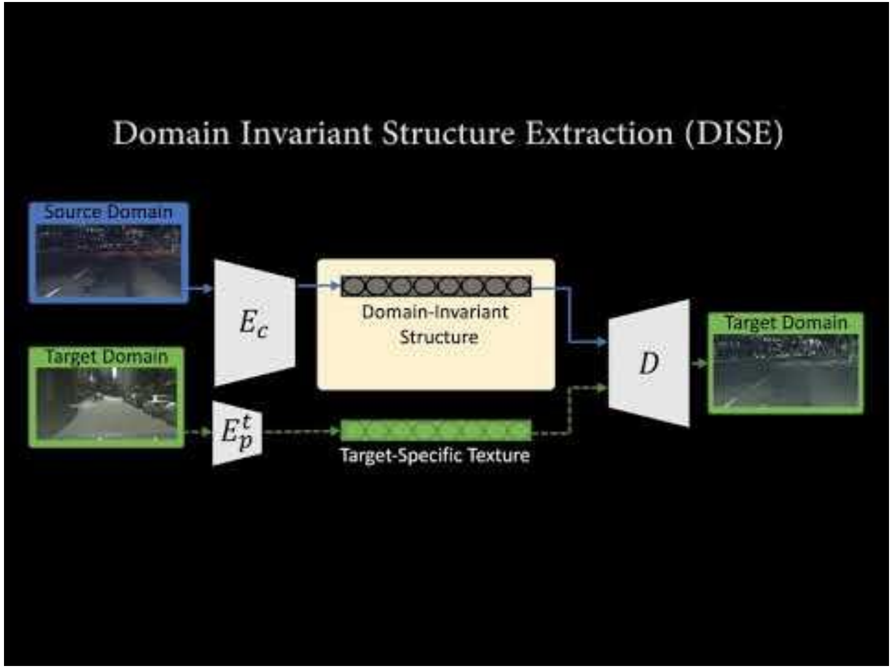

Supervisor: Dr. Daniele Cattaneo - Paper link: https://arxiv.org/abs/1903.12212

Topic 12: PnPNet: End-to-End Perception and Prediction with Tracking in the Loop

- Joint perception and motion forecasting in the context of self-driving vehicles
- PnPNet: end-to-end model that takes as input sequential sensor data, and outputs at each time step object tracks and their future trajectories
- Tracking module generates object tracks online from detections and exploits trajectory level features for motion forecasting

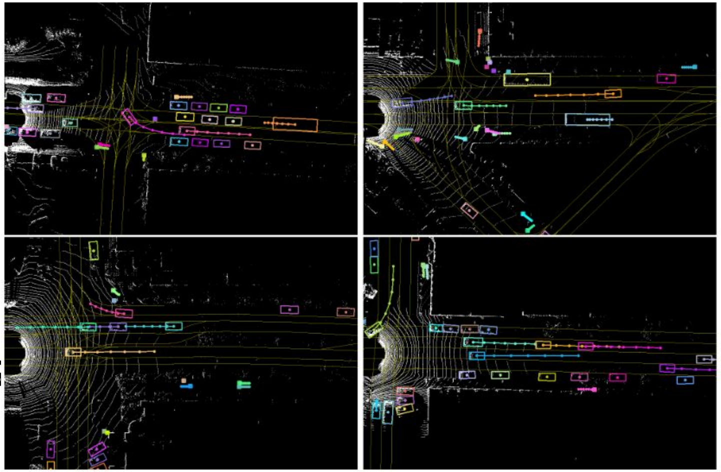

Source: Liang et al. 2020

Supervisor: Dr. Daniele Cattaneo - Paper link: https://arxiv.org/abs/2005.14711

Assessing Interest

- Fill out the form: https://forms.gle/RdUgJfM5XKYQVGzW9
- Places will be assigned based on priority suggestions of HisInOne and motivation of the student by 25/11/2020.