Robot Learning Proseminar/Seminar SS 2021

Robot Learning Lab

Albert-Ludwigs-Universität Freiburg

Friday, 23rd April 2021

UNI FREIBURG

Evaluation

Evaluation	Bachelor's Due Date	Master's Due Date
Paper Abstract	18/06/2021	18/06/2021
Seminar Presentation	23/07/2021	16/07/2021
Paper Summary	30/07/2021	30/07/2021

- Abstract → At most 2 pages
- Presentation → At most 20 minutes
- Summary → At most 7 pages excluding bibliography and figures
- Final grade → Abstract, Presentation, Summary, Seminar participation

Bachelor's Seminar: https://rl.uni-freiburg.de/teaching/ss21/robotlearning-bachelor/

Master Seminar: https://rl.uni-freiburg.de/teaching/ss21/robotlearning-master/

Enrollment Procedure

Select three papers in decreasing order of preference → Register for the seminar in HISInOne → Students selected based on HISInOne Priority Suggestion → Students assigned a paper based on their paper preference

Bachelor's Form
Master's Form

On 29/04/2021

On 29/04/2021

Robot Learning

- Tremendous progress on complex, high dimensional data
  - Speech Recognition
  - Computer Vision
  - Natural Language Understanding
- Autonomous systems smart enough to operate in the real world

Perception

- Complex environments

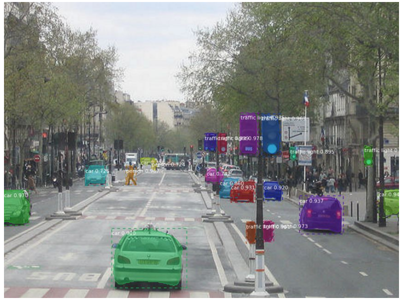

- Noisy observations and sensors

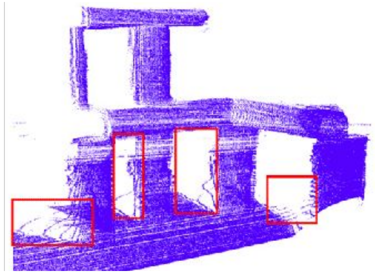

Mask R-CNN for object detection and instance segmentation on Keras and TensorFlow, Waleed et. al., 2017
Rohit Mohan and Abhinav Valada, "EfficientPS: Efficient Panoptic Segmentation", arXiv preprint arXiv:2004.02307, 2020.Patil, Ashok Kumar & Kumar, G Ajay & Kim, Tae-Hyoung & Chai, Young-Ho. (2018). Hybrid approach for alignment of a pre-processed three-dimensional point cloud, video, and CAD model using partial point cloud in retrofitting applications. International Journal of Distributed Sensor Networks. 14. 155014771876645. 10.1177/1550147718766452.

Unknown, Open World

- Unknown world → Many unlabelled samples
- Uncertainty estimation
- Adversarial attacks

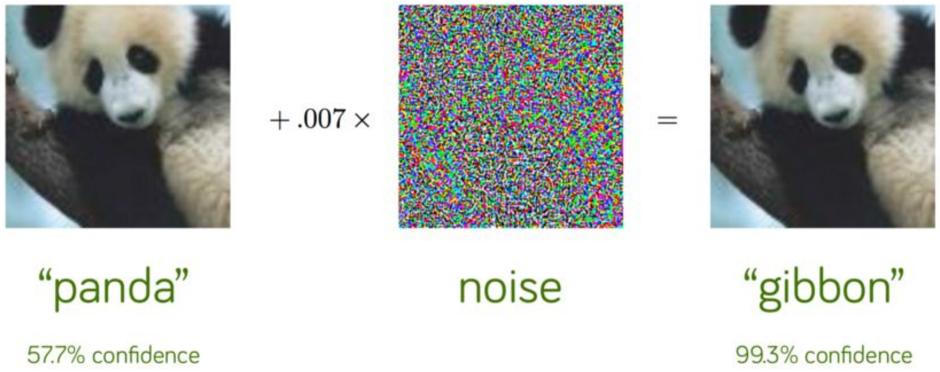

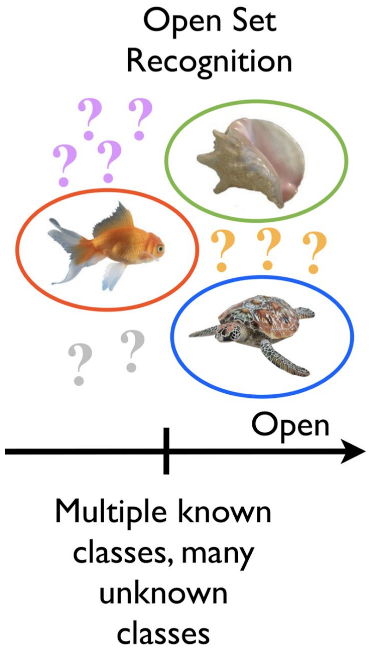

Towards Open Set Recognition, Scheirer et. al., 2012

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

Autonomous Decision Making

- Reinforcement learning for short- and long-term decision making

Reinforcement Learning

- Model free RL
  - Adapts to complex scenarios
  - Directly optimise policy
  - Data intensive

- Model-based RL
  - Learns a world model
  - Promise of better generalisation

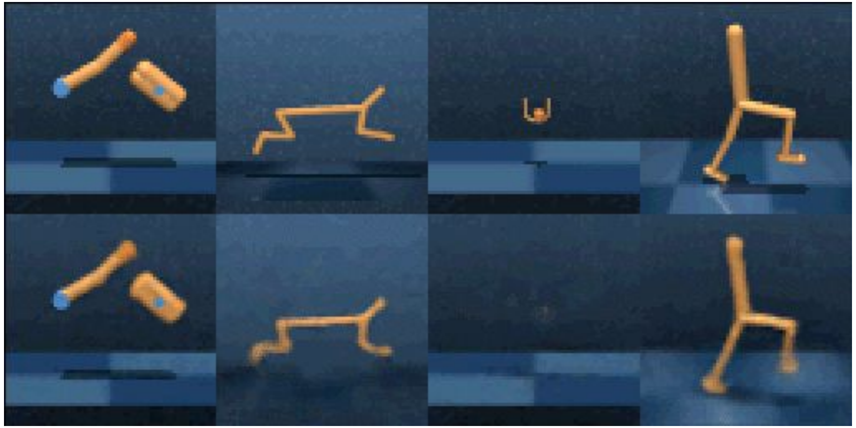

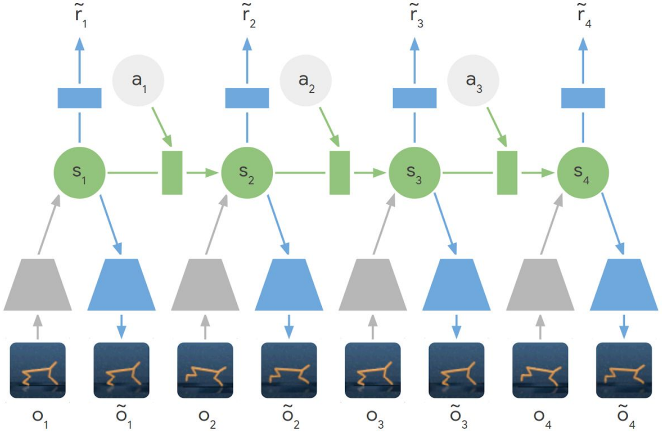

Learning Latent Dynamics for Planning from Pixels, Hafner et. al., 2019

Expensive Real World Data

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

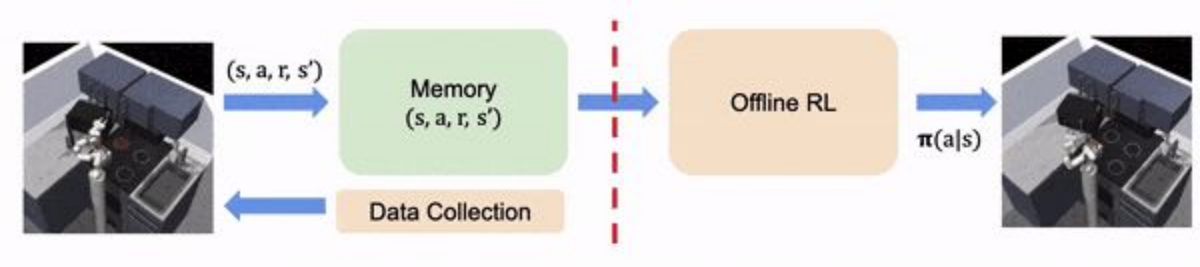

Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection. Sergey Levine, Peter Pastor, Alex Krizhevsky, Deirdre Quillen
D4RL: Datasets for Deep Data-Driven Reinforcement Learning, Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, Sergey Levine

Weak- and Self-Supervision

- Provide labels for simpler tasks
  - Object presence and absence
  - Consistency over time
  - Viewpoint invariance
- Reduce oversight
  - Automatic resets
  - Reward labelling

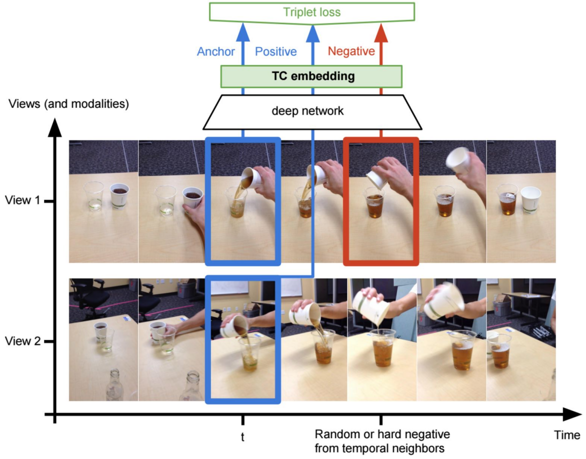

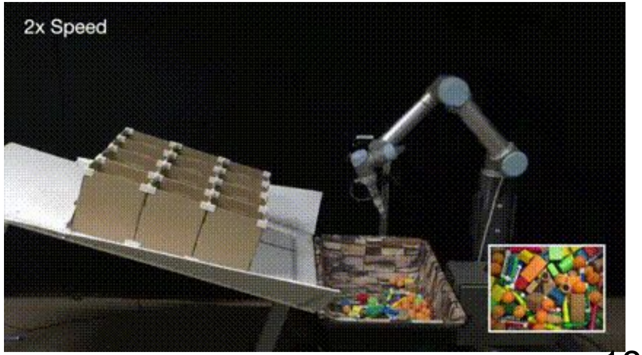

Time-Contrastive Networks: Self-Supervised Learning from Video, Sermanet et. al., 2018
TossingBot: Learning to Throw Arbitrary Objects, Zeng et. al., 2019.

Bachelor Seminar Topics

1. PolarNet: An Improved Grid Representation for Online LiDAR Point Cloud Semantic Segmentation

- Uses polar representation instead of spherical or BEV
- Proposes a novel ring-based CNN for use with the polar representation
- Beat the SOTA while using lesser parameters

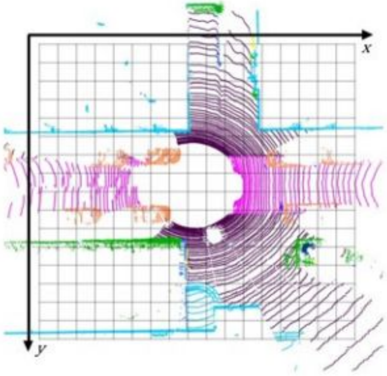

(a) Cartesian BEV

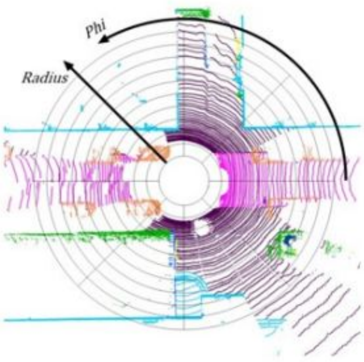

(b) Polar BEV

Source: Zhang et al. 2020

Supervisor: Nikhil Gosala - [Link to paper]

1. PolarNet: An Improved Grid Representation for Online LiDAR Point Cloud Semantic Segmentation

- Point clouds quantised and projected into a polar grid
- Ring CNN applied to polar grid to learn the semantic mapping.

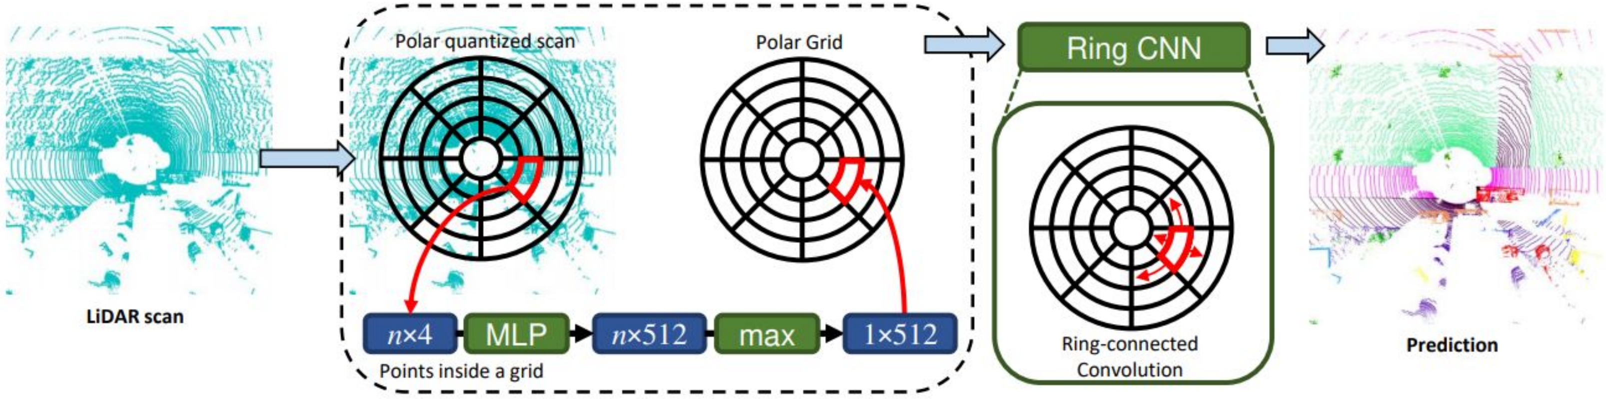

Source: Zhang et al. 2020

2. DroNet: Learning to Fly by Driving

- Addresses the challenge of drone navigation in urban environments
- Uses data from cars and bicycles to address lack of drone-relevant data
- Learns the steering angle and collision probability to control the drone

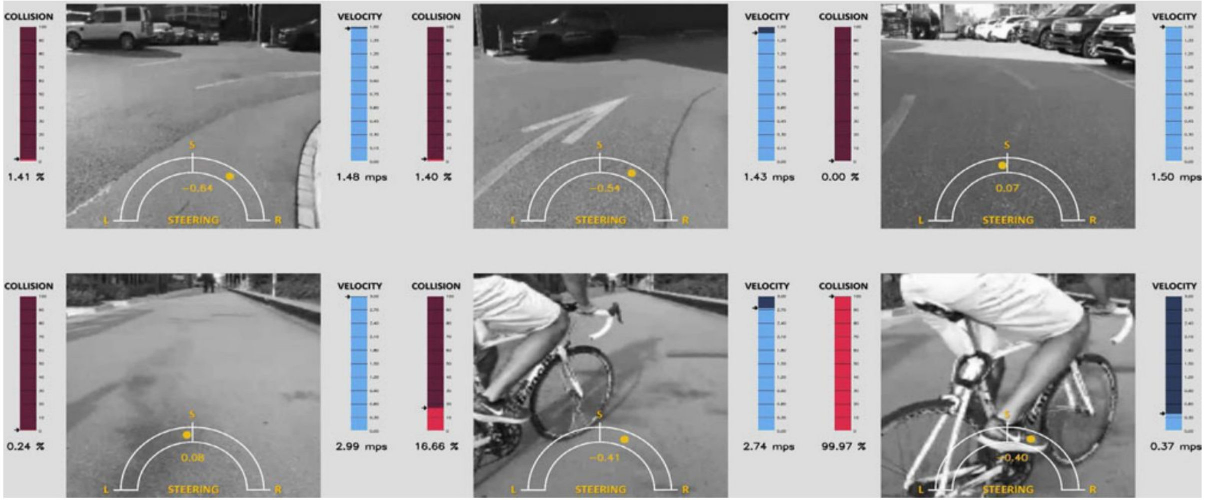

Source: Loquercio et al.

Supervisor: Nikhil Gosala - [Link to paper]

2. DroNet: Learning to Fly by Driving

Video Link

Supervisor: Nikhil Gosala - Link to paper

3. Let there be Color: Joint End to End Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification

- Approach to introduce colour into grayscale images
- Merges patch-level local features with image-level global features

(a) Cranberry Picking, Sep. 1911 (b) Burns Basement, May 1910 (c) Miner, Sep. 1937 (d) Scott's Run, Mar. 1937

Source: IIzuka et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

3. Let there be Color: Joint End to End Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification

- Local features capture texture or objects
- Global features capture time of day, location (indoors/outdoors)
- Changing the global features can help re-imagine the image

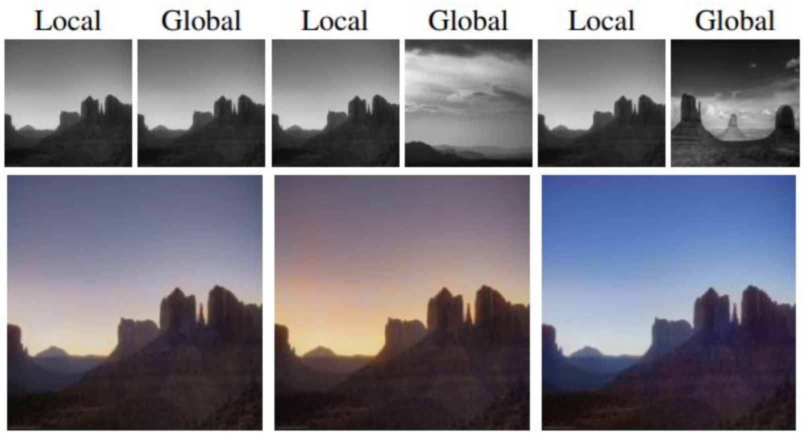

Source: *IIzuka et al.*

Supervisor: Juana Valeria Hurtado - [Link to paper]

4. Context Encoders - Feature Learning by Inpainting

- Approach to hallucinate missing / blacked out regions in image
- Validates the ability of a network to understand structure in scenes
- L2 loss → Captures overall structure
- Adversarial loss → Chooses one mode out of several possible ones

(a) Input context (b) Human artist (c) Context Encoder ($L2$ loss) (d) Context Encoder ($L2$ + Adversarial loss)

Source: Pathak et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

4. Context Encoders - Feature Learning by Inpainting

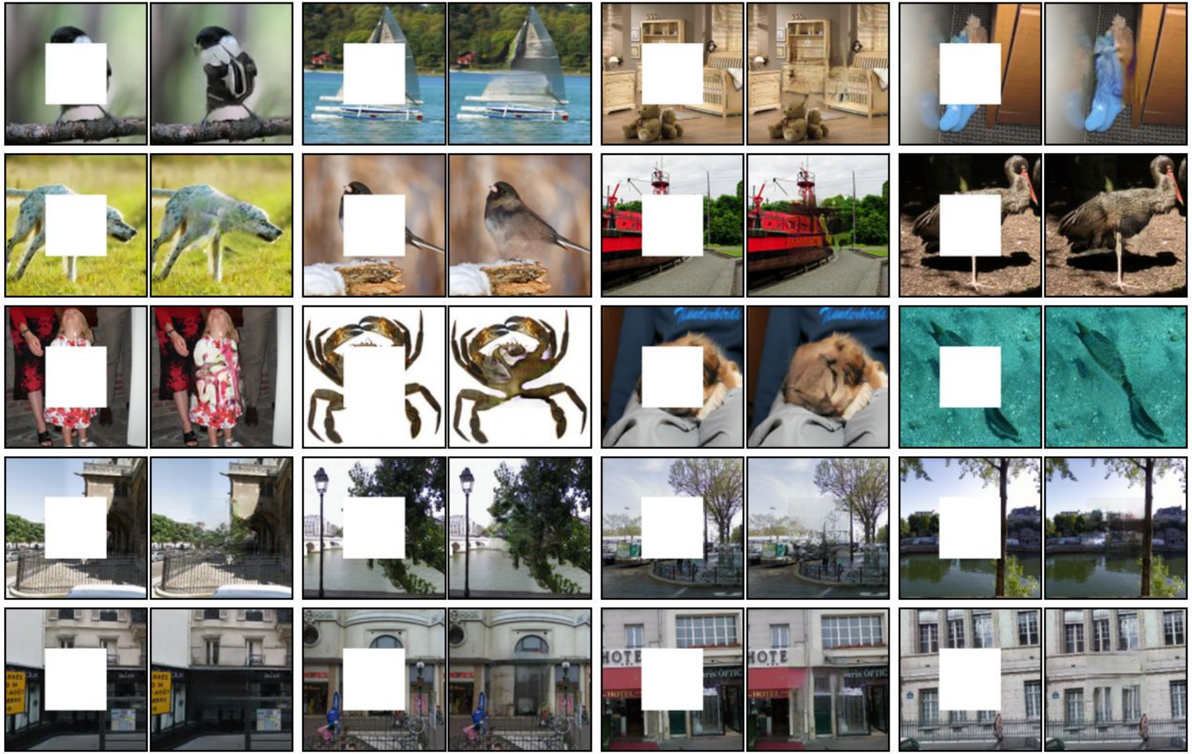

Source: Pathak et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

5. R2D2: Repeatable and Reliable Detector and Descriptor

- Approach to generate reliable and repeatable keypoints
- **Detect-and-Describe** instead of **Detect-then-Describe**
- Good keypoints:
  - Repeatable (generate strong cues, like corners)
  - Reliable (ability to uniquely match across images)
- Predict discriminativeness along with keypoints

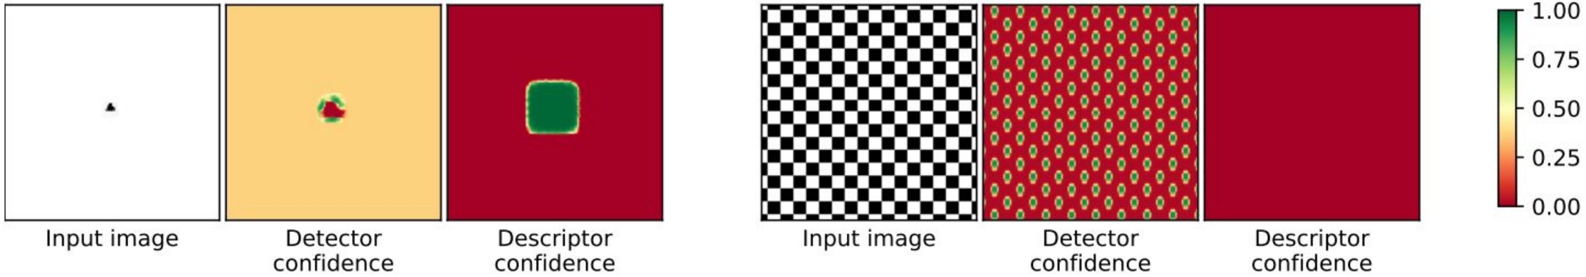

Source: Revaud et al.

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

5. R2D2: Repeatable and Reliable Detector and Descriptor

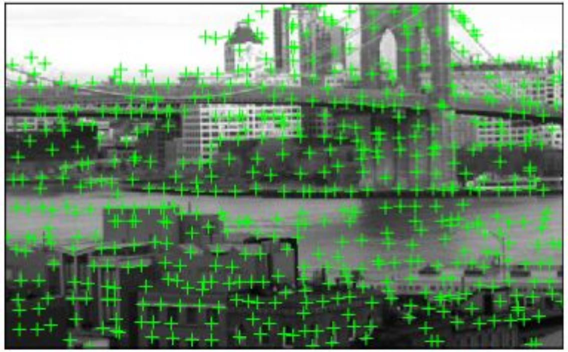

Ground Truth Keypoints

Keypoint Repeatability

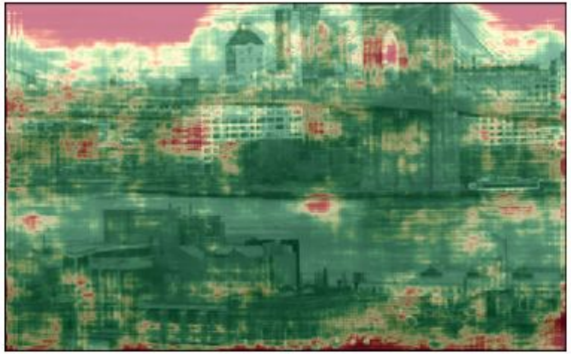

Keypoint Reliability

Source: Revaud et al.

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

6. SegMatch: Segment Based Place Recognition in 3D Point Clouds

- A place recognition approach for 3D point clouds for use in SLAM
- Generates and matches 3D segments:
  - Alleviates the need for accurate object detection or notion of “object”
  - More informative as opposed to keypoint descriptors

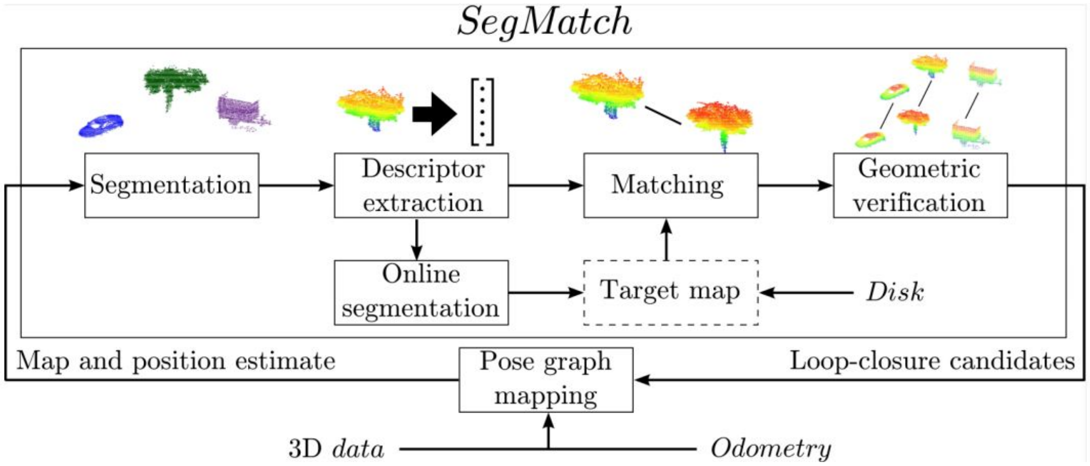

Source: Dube et al.

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

6. SegMatch: Segment Based Place Recognition in 3D Point Clouds

-

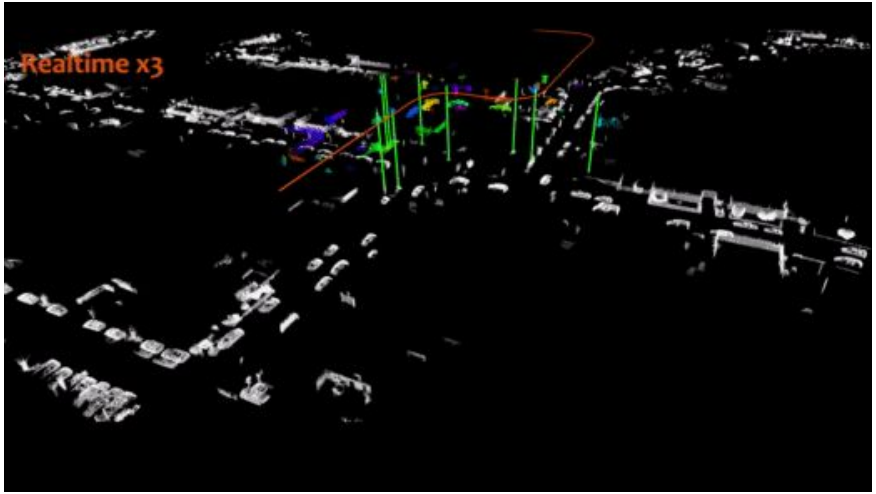

Segment Matching

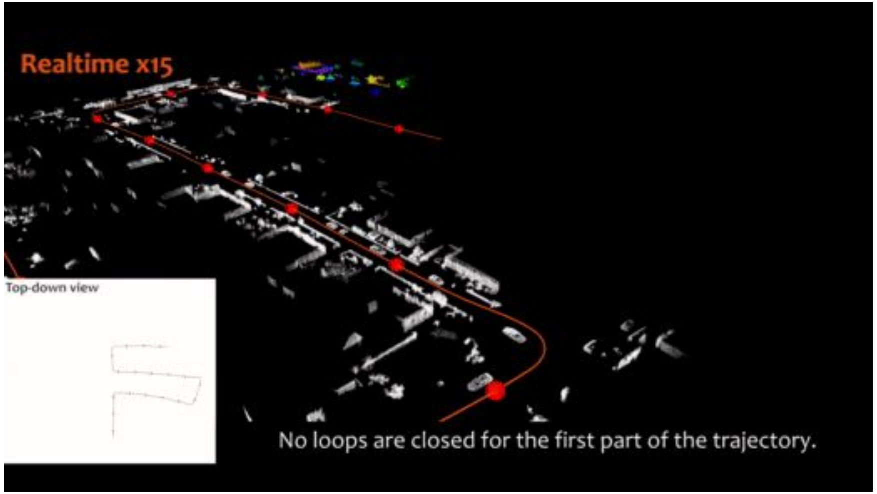

Loop Closure Detection

[Video Link]

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

7. Sim-to-Real Transfer of Robotic Control with Dynamics Randomization

- Bridge the “reality-gap” between simulation and real-world
- Randomizes the robot simulator dynamics for policy generalisation
- Policies trained in simulation can directly be used in the real world!
- LSTM used to model the memory nodes of the DeepRL network

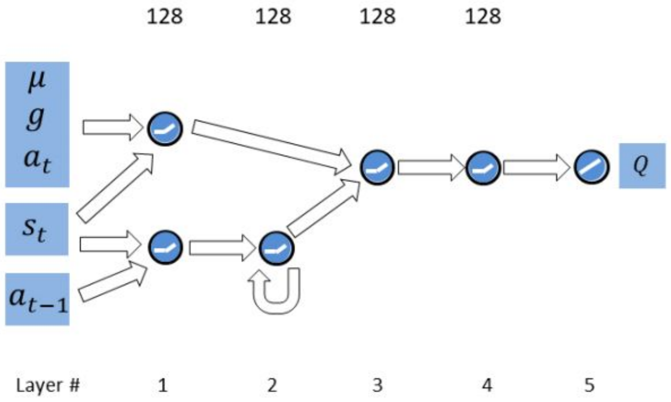

Policy Network

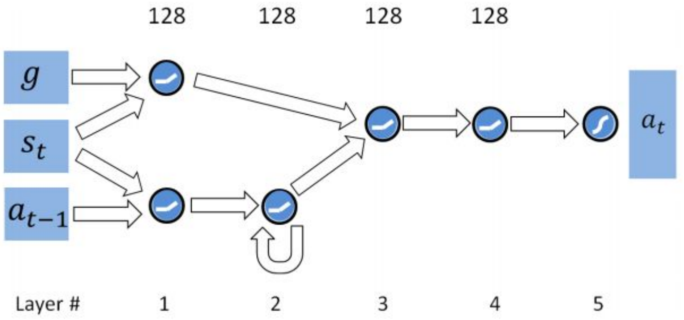

Value Network

Source: Peng et al.

Supervisor: Eugenio Chisari - [Link to paper]

7. Sim-to-Real Transfer of Robotic Control with Dynamics Randomization

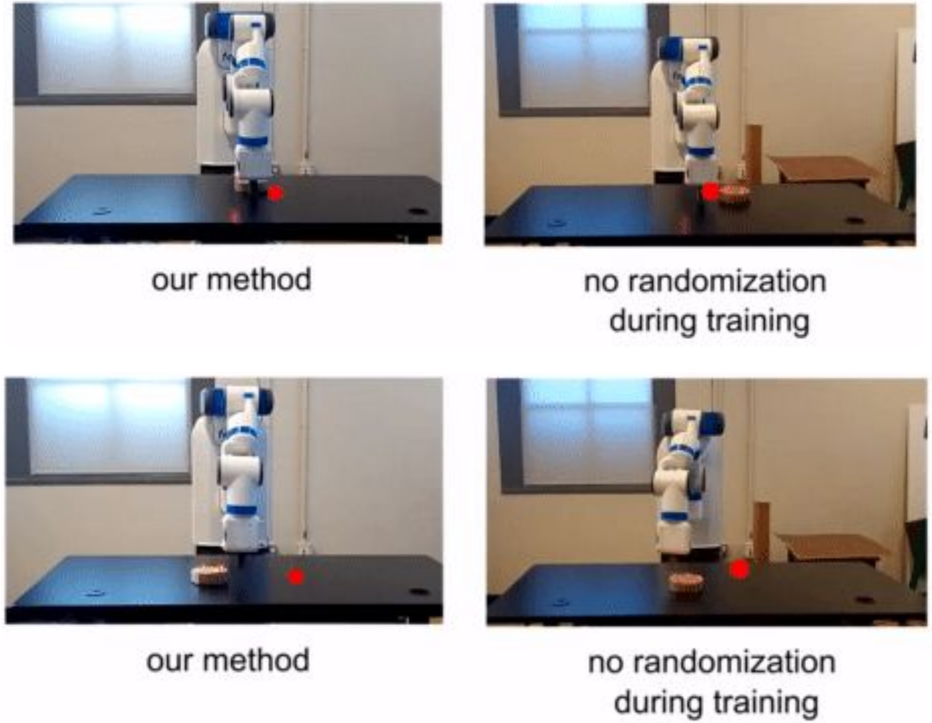

[Video Link]

Supervisor: Eugenio Chisari - [Link to paper]

8. Continuous Control for High-Dimensional State Spaces: An Interactive Learning Approach

- Introduces human corrective feedback into the policy optimisation phase
- No reward function needed
- Sample size greatly reduced because of expert intervention
- Human actively corrects the errors of the policy and physically interacts with the system during the learning process

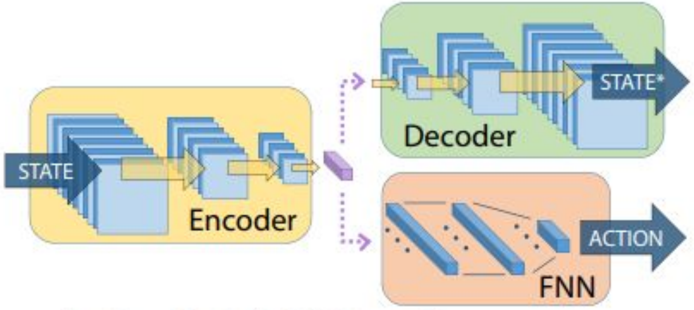

Source: Perez-Dattari et al.

8. Continuous Control for High-Dimensional State Spaces: An Interactive Learning Approach

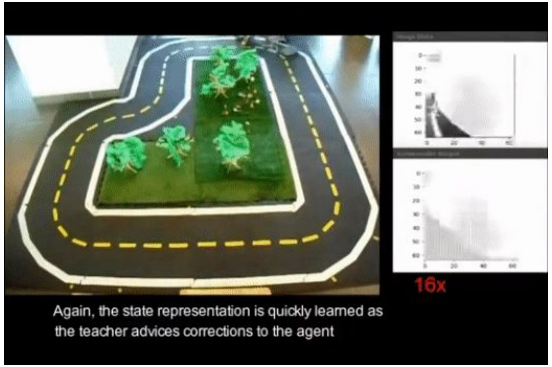

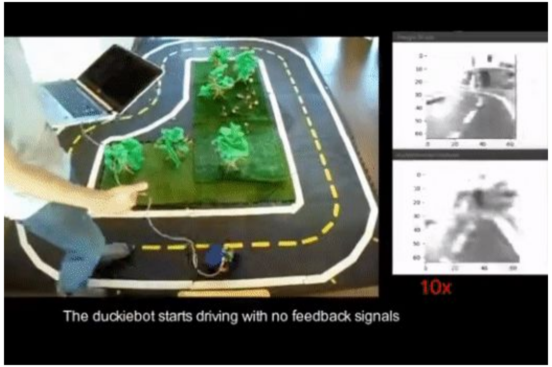

Supervisor: Eugenio Chisari - [Link to paper]

[Video Link]

9. Playing Atari with Deep Reinforcement Learning

- CNN to directly learn a control policy from high-dimensional input using RL
- Network is trained with a variant of Q-Learning
- Single model trained once plays multiple Atari games
- This paper led to a $500 million acquisition of DeepMind by Google

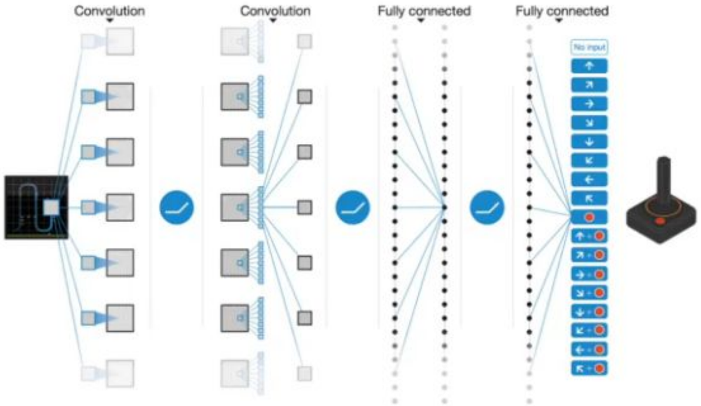

Source: Mnih et al.

Supervisor: Daniel Honerkamp - [Link to paper]

9. Playing Atari with Deep Reinforcement Learning

	B. Rider	Breakout	Enduro	Pong	Q*bert	Seaquest	S. Invaders
Random	354	1.2	0	-20.4	157	110	179
Sarsa [3]	996	5.2	129	-19	614	665	271
Contingency [4]	1743	6	159	-17	960	723	268
DQN	**4092**	**168**	**470**	**20**	**1952**	**1705**	**581**
Human	7456	31	368	-3	18900	28010	3690
HNeat Best [8]	3616	52	106	19	1800	920	**1720**
HNeat Pixel [8]	1332	4	91	-16	1325	800	1145
DQN Best	**5184**	**225**	**661**	**21**	**4500**	**1740**	1075

Source: Mnih et al.

10. Recurrent Models of Visual Attention

- CNNs have high computational complexity
- CNNs replaced with RNNs to control the computation complexity
  - RNN adaptively selects only the most relevant features for a given task
  - The next location to be attended to depends on the past regions and the task
- The computation time is independent of the input size
- Model is non-differentiable and is trained using RL for a specific task

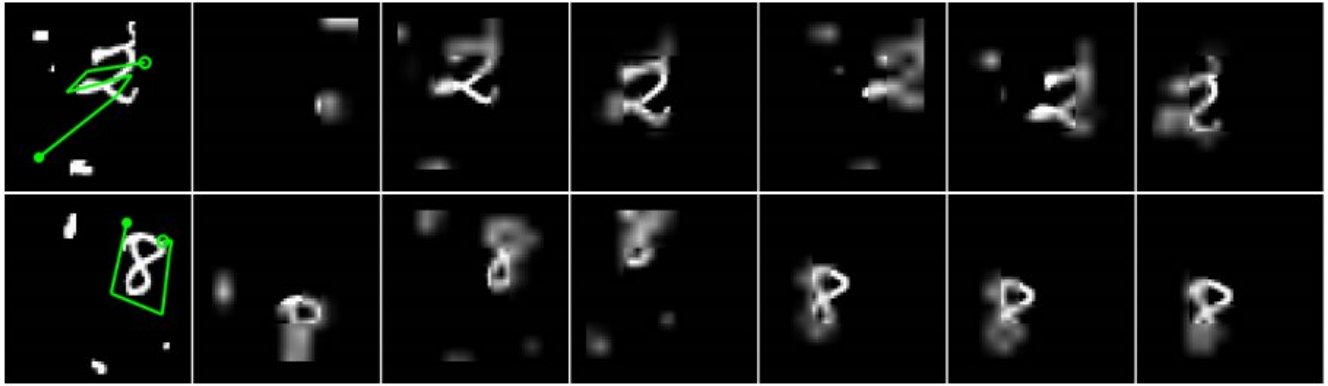

Source: Mnih et al.

Supervisor: Daniel Honerkamp - [Link to paper]

10. Recurrent Models of Visual Attention

- The network incrementally combines the different regions to learn a dynamic representation of the image

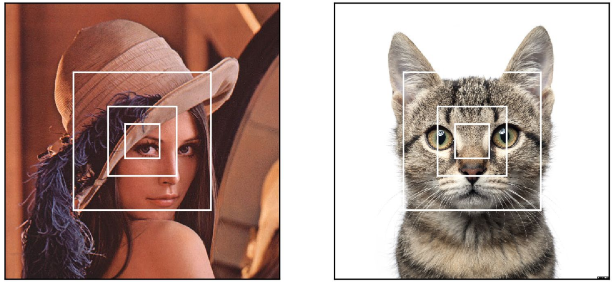

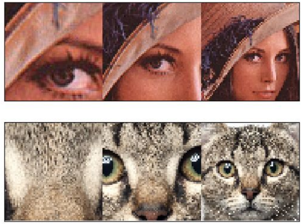

Source: *[kevinzakka GitHub]*

Supervisor: Daniel Honerkamp - [Link to paper]

11. Efficiency and Equity are Both Essential: A Generalized Traffic Signal Controller with Deep Reinforcement Learning

- An RL method to learn a traffic signal control policy to optimise traffic flow
- Reward function accounts for both Equity (Variance) and Efficiency (Mean)
- Adaptive discounting used to reduce the reward for transition phases

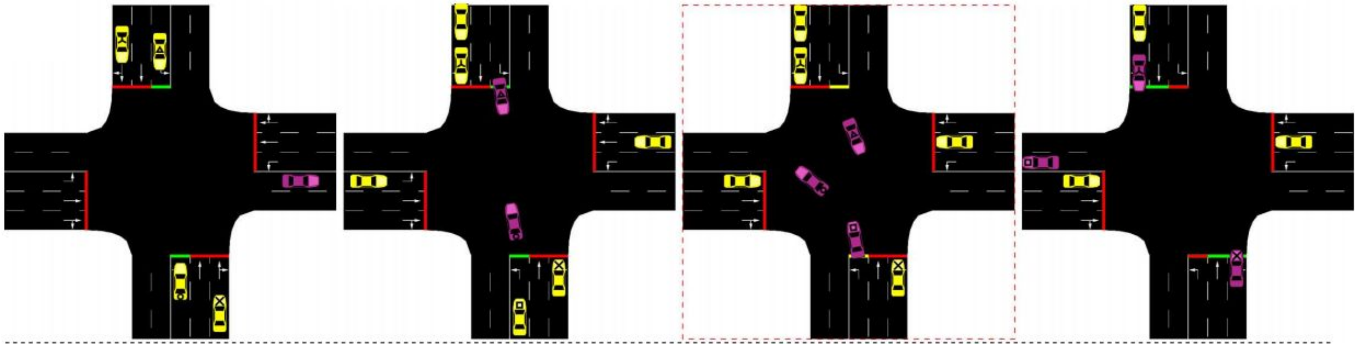

Source: Yan et al.

Supervisor: Dr. Tim Welschehold- [Link to paper]

12. Use the Force, Luke! Learning to Predict Physical Forces by Simulating Effects

- Approach to infer forces exerted on objects from a video
- Predicts physical information that can be used to recreate the interaction
- Force and contact point prediction used to deform object mesh to recreate the action

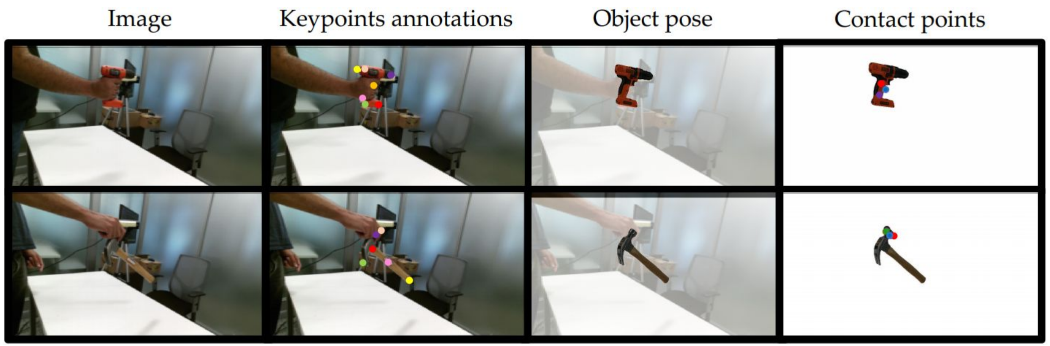

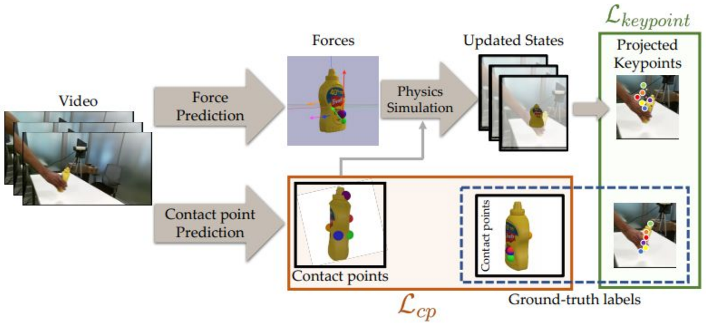

Source: Ehsani et al.

Supervisor: Dr. Tim Welschehold - [Link to paper]

12. Use the Force, Luke! Learning to Predict Physical Forces by Simulating Effects

Video Link

Supervisor: Dr. Tim Welschehold - Link to paper

Master Seminar Topics

1. Perceive, Attend, and Drive: Learning Spatial Attention for Safe Self-Driving

- Lot of computation in perception is wasted on regions that don't matter!
- An approach to automatically attend to important regions of the image for use with motion planning
- The attention mask increases safety by allowing focused computation

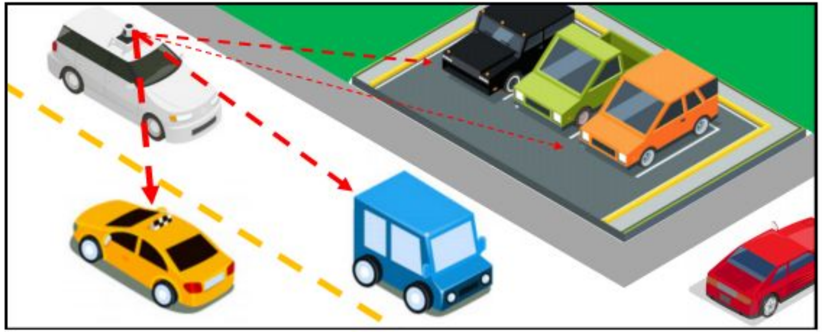

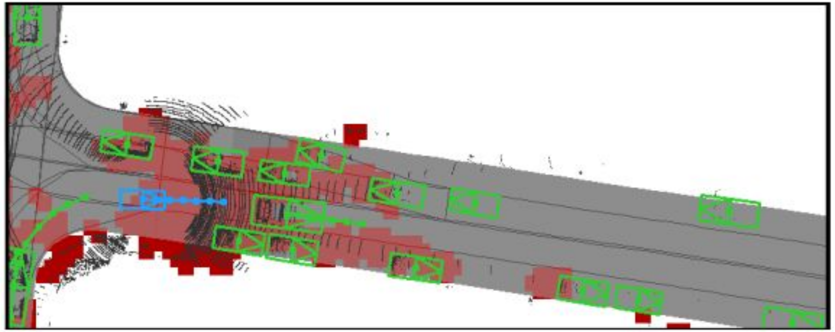

Source: Wei et al.

Supervisor: Nikhil Gosala - [Link to paper]

1. Perceive, Attend, and Drive: Learning Spatial Attention for Safe Self-Driving

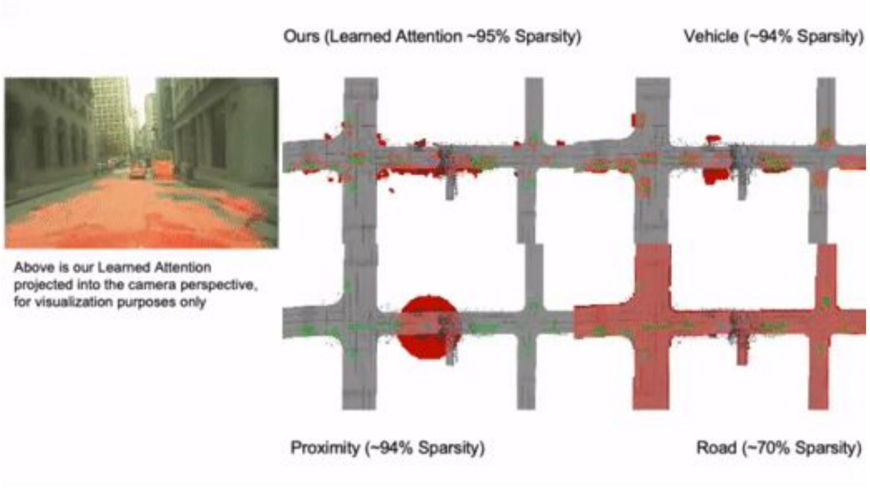

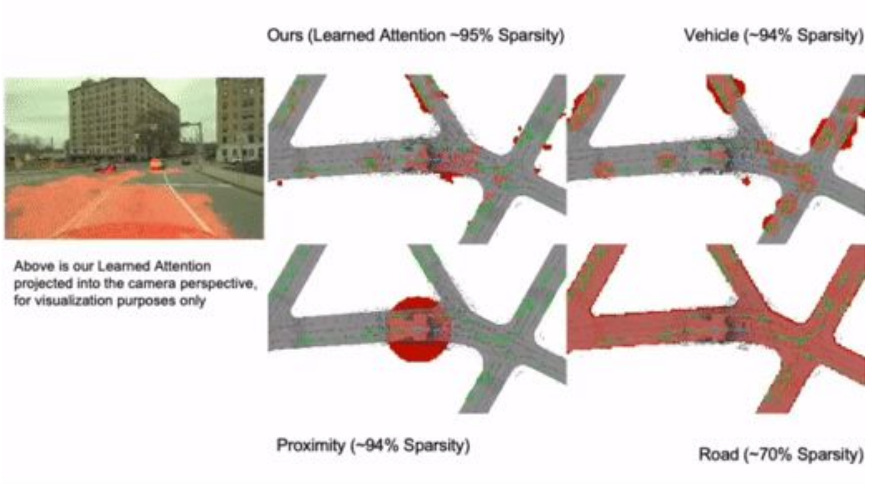

[Video Link]

Supervisor: Nikhil Gosala - [Link to paper]

2. Orthographic Feature Transform for Monocular 3D Object Detection

- An approach to convert a 2D frontal perspective view to a 3D orthographic view for 3D object detection in monocular images
- Orthographic Feature Transform maps features in frontal view to features in the Bird's eye view
- Generates a holistic view where object scale and distance are meaningful

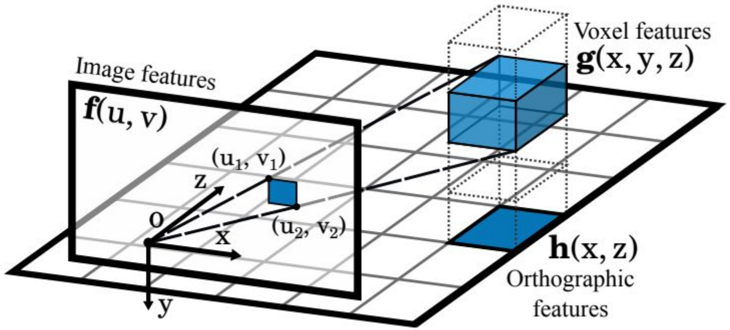

Source: Roddick et al.

Supervisor: Nikhil Gosala - [Link to paper]()

2. Orthographic Feature Transform for Monocular 3D Object Detection

Source: Roddick et al.

[Video Link]

3. Neural Scene Graphs for Dynamic Scenes

- A neural rendering approach for accurate view synthesis of dynamic scenes using scene graph representations
- The network decomposes the scene into static and dynamic objects and encodes their representations using scene graphs
- Novel scenes can be generated by manipulating the learnt scene graphs

(a) Neural scene graph in isometric view.

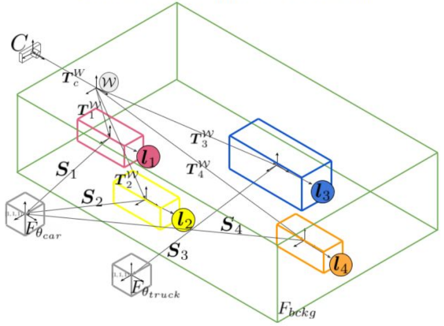

(b) Neural scene graph from the ego-vehicle view.

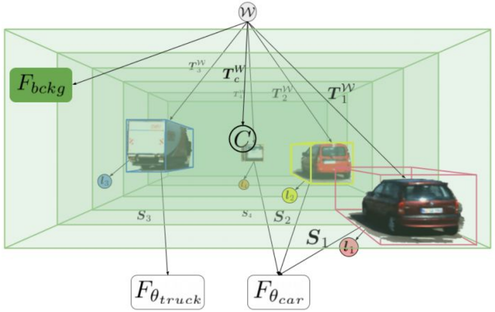

Source: Ost et. al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

3. Neural Scene Graphs for Dynamic Scenes

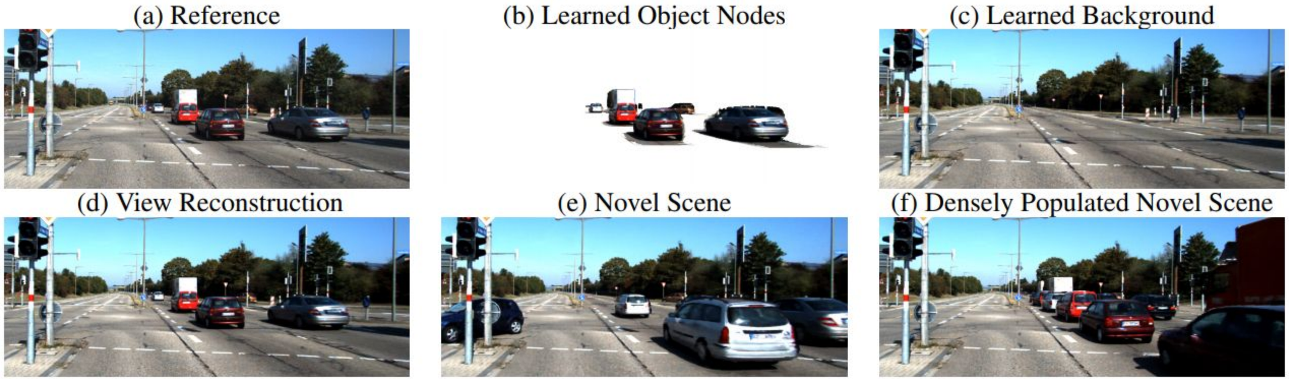

Source: Ost et. al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

4. Semantic Object Prediction and Spatial Sound Super-Resolution with Binaural Sounds

- Approach to perform semantic object prediction using binaural sound
- Network trained using a teacher-student approach
  - A vision network acts as a teacher
  - Audio network is the student and learns to generate the same results as the teacher
- Depth estimation and spatial sound super resolution used to boost results

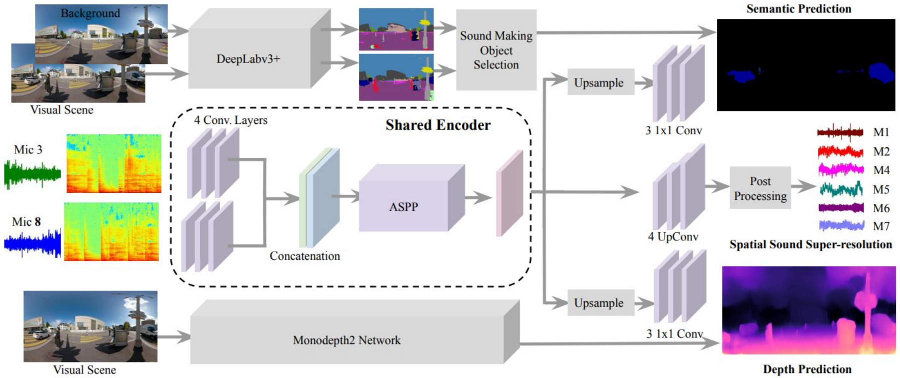

Source: Vasudevan et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

4. Semantic Object Prediction and Spatial Sound Super-Resolution with Binaural Sounds

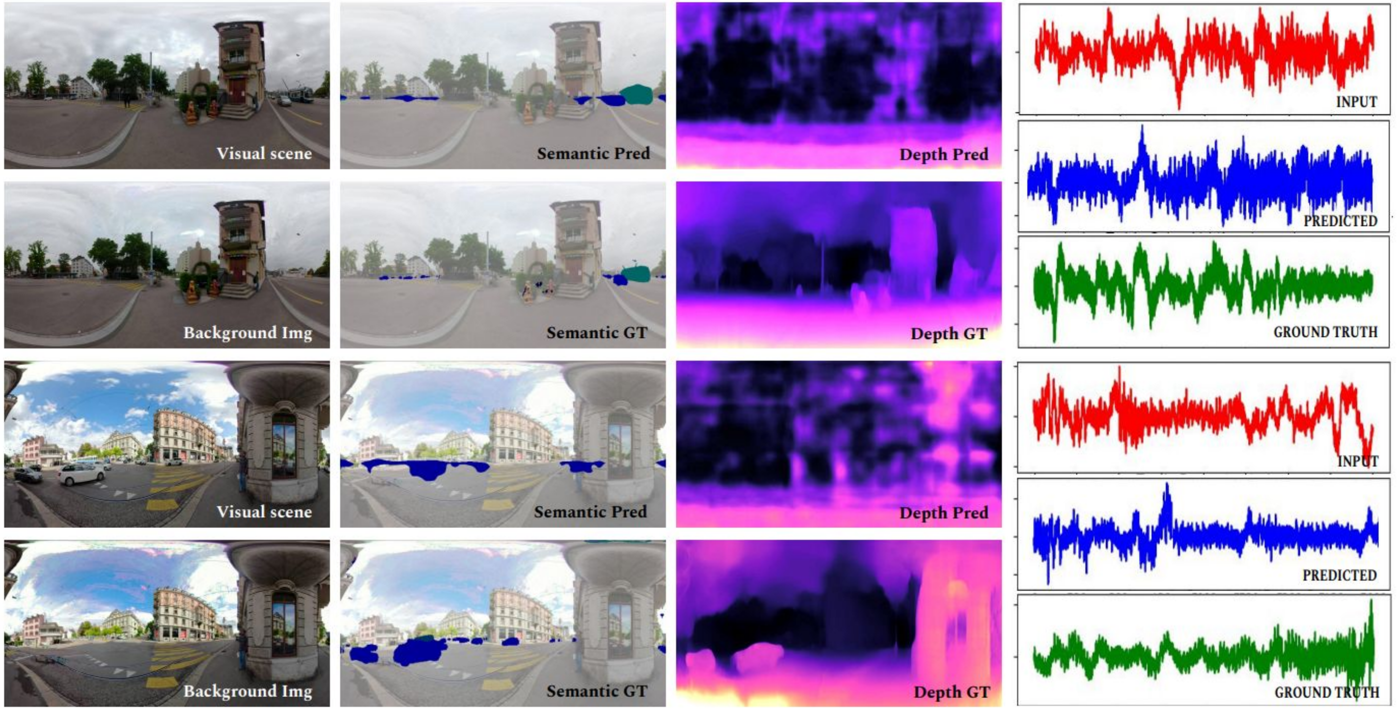

Source: Vasudevan et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

5. xMUDA: Cross-Modal Unsupervised Domain Adaptation for 3D Semantic Segmentation

- Transfer features learnt on one dataset to another dataset
- Exploits multiple modalities to improve the result of unsupervised domain adaptation between datasets
- Each modality learns to mimic the better qualities of the other, making the overall domain adaptation objective better

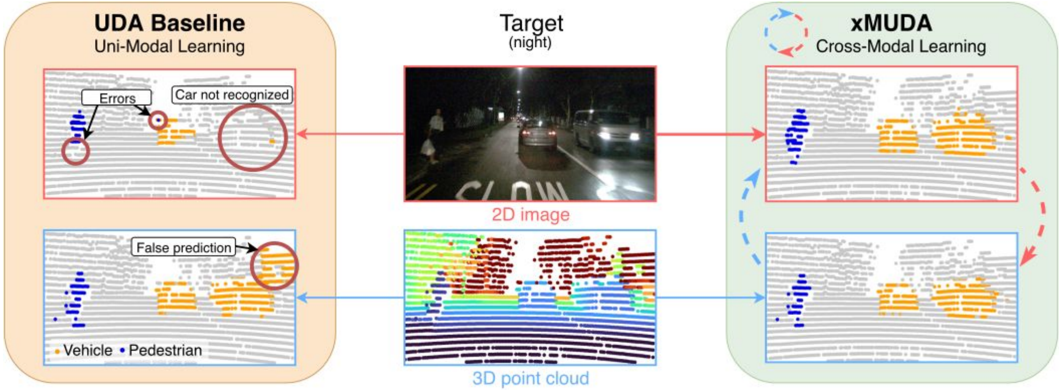

Source: Jaritz et al.

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

5. xMUDA: Cross-Modal Unsupervised Domain Adaptation for 3D Semantic Segmentation

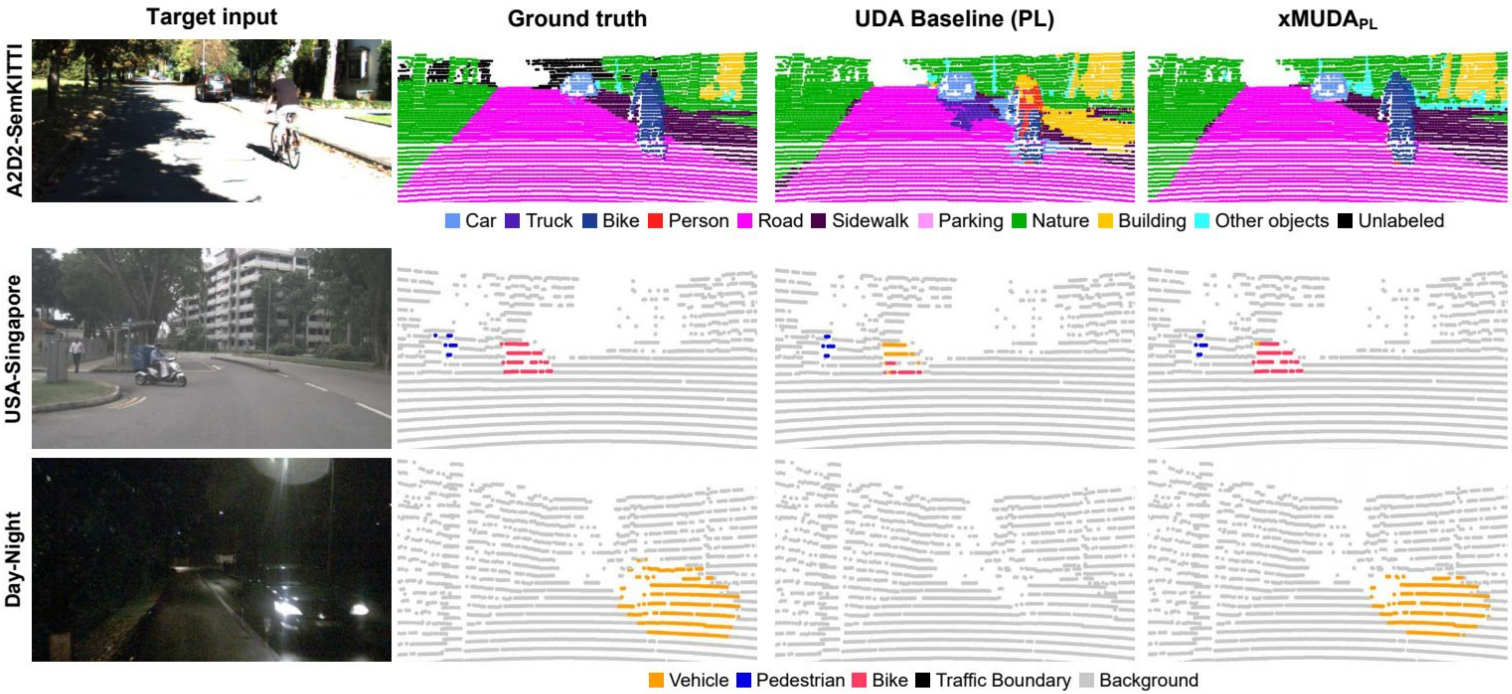

Source: Jaritz et al.

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

6. STaR: Self-Supervised Tracking and Reconstruction of Rigid Objects in Motion with Neural Rendering

- Self-supervised tracking and dynamic scene reconstruction approach from multi-view RGB videos
- Explicitly handles rigid motion, which is the main cause of failure in competing approaches
- NeRFs and rigid poses are jointly optimised to achieve the goal

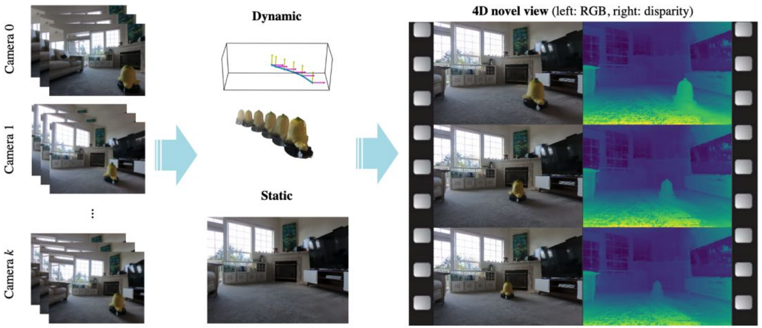

Source: Yuan et al.

Supervisor: Dr. Daniele Cattaneo - [Link to paper]

6. STaR: Self-Supervised Tracking and Reconstruction of Rigid Objects in Motion with Neural Rendering

[Video Link](#)

Supervisor: Dr. Daniele Cattaneo - [Link to paper](#)

7. Self-Supervised Correspondence in Visuomotor Policy Learning

- Approach to improve generalisation of visuomotor policy learning using self-supervised correspondences
- Self-supervised dense correspondence training generalises across different objects using very little data
- Also propose a technique for multi-camera time-synchronised dense spatial correspondence learning

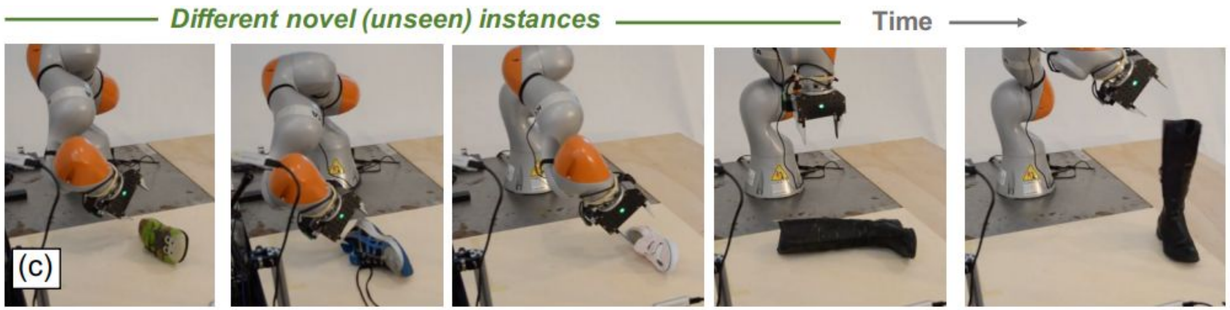

Source: Florence et al.

Supervisor: Eugenio Chisari - [Link to paper]

7. Self-Supervised Correspondence in Visuomotor Policy Learning

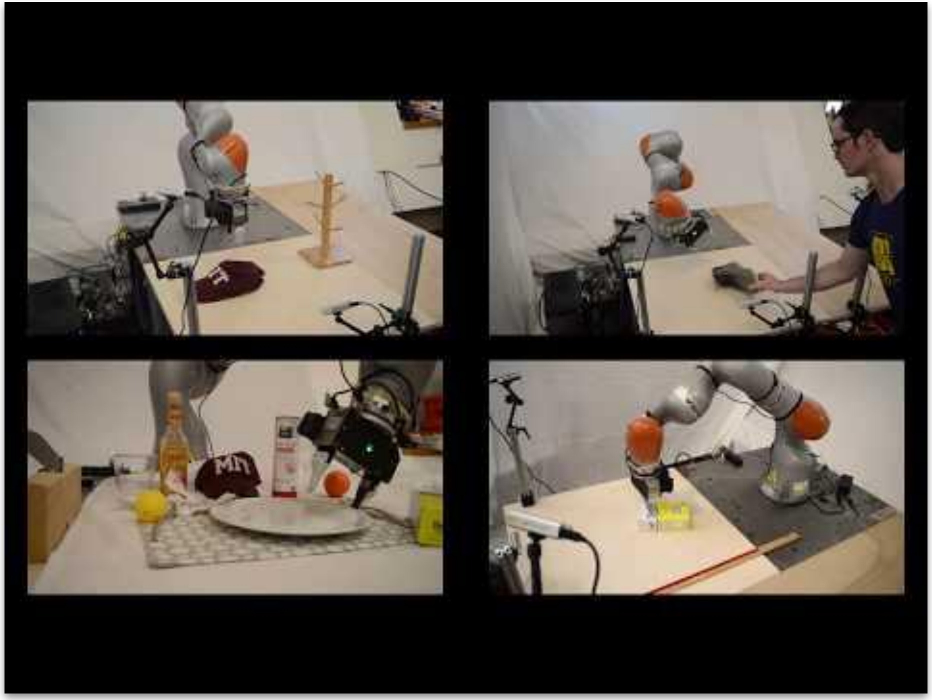

Video Link

Supervisor: Eugenio Chisari - Link to paper

8. Accelerating Reinforcement Learning with Learned Skill Priors

- Approach taps into previously learned skills to learn a *prior over skills*
- Deep Latent Variable model jointly learns embedding space of skills and skill priors
- Skill priors accelerate downstream policy learning on complex tasks

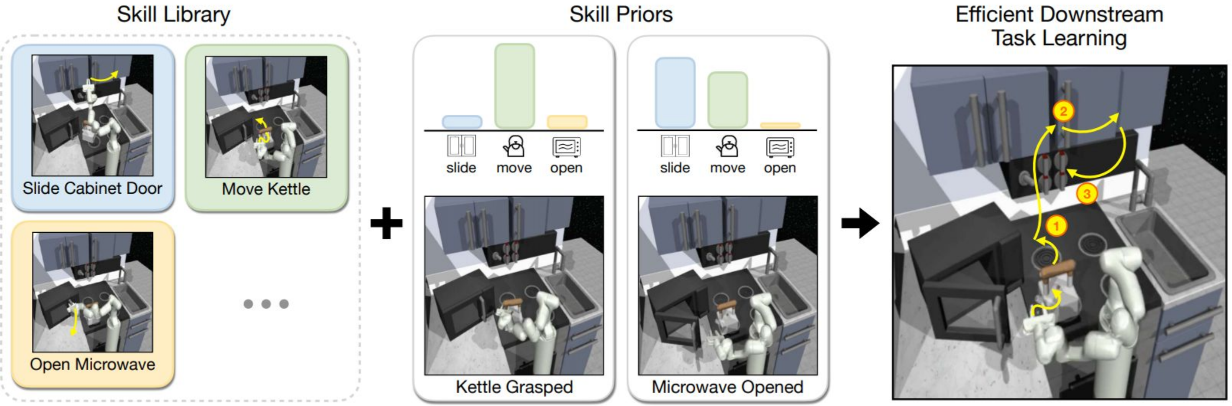

Source: Pertsh et al.

Supervisor: Eugenio Chisari - [Link to paper]

8. Accelerating Reinforcement Learning with Learned Skill Priors

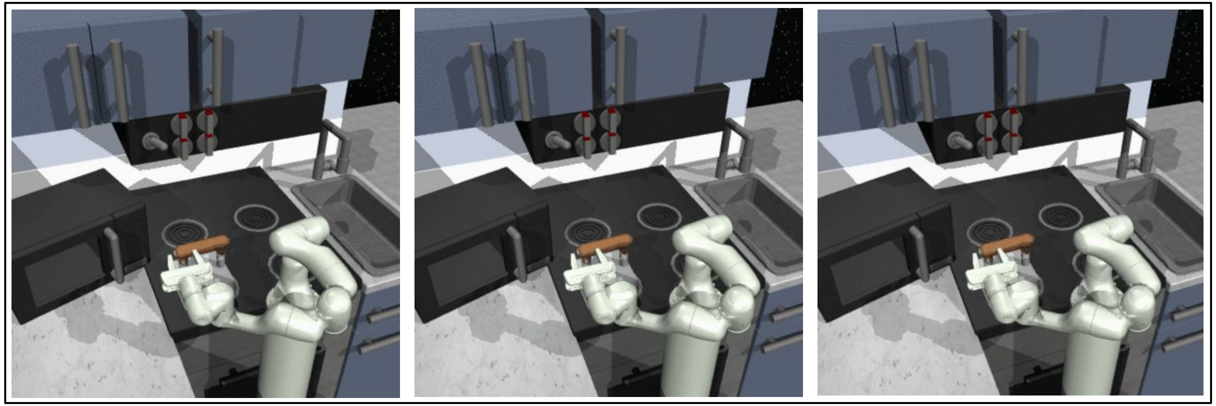

Competing approaches

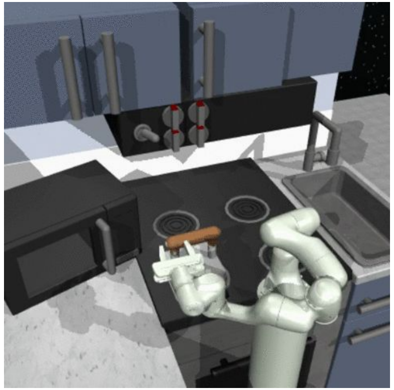

Their approach

Source: Pertsh et al.

Supervisor: Eugenio Chisari - [Link to paper]

9. Asymmetric Self-Play for Automatic Goal Discovery in Robotic Manipulation

- A zero-shot generalisation approach to solve multiple unseen tasks using asymmetric self-play
- Two agents play against one another - Alice and Bob
  - Alice proposes a challenge to Bob
  - Bob tries to solve that challenge
- Bob learns a policy by observing Alice’s trajectory and cloning its behaviour
- ASP always gives an achievable goal and comes up with many novel goals

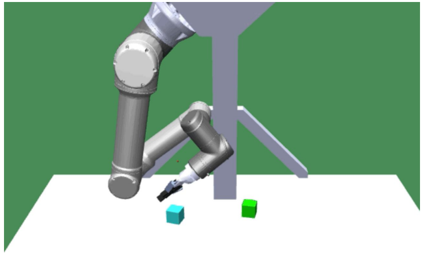

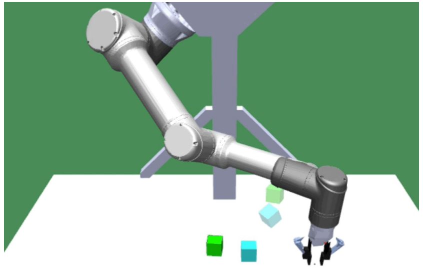

[Video Link]

Supervisor: Daniel Honerkamp - [Link to paper]

9. Asymmetric Self-Play for Automatic Goal Discovery in Robotic Manipulation

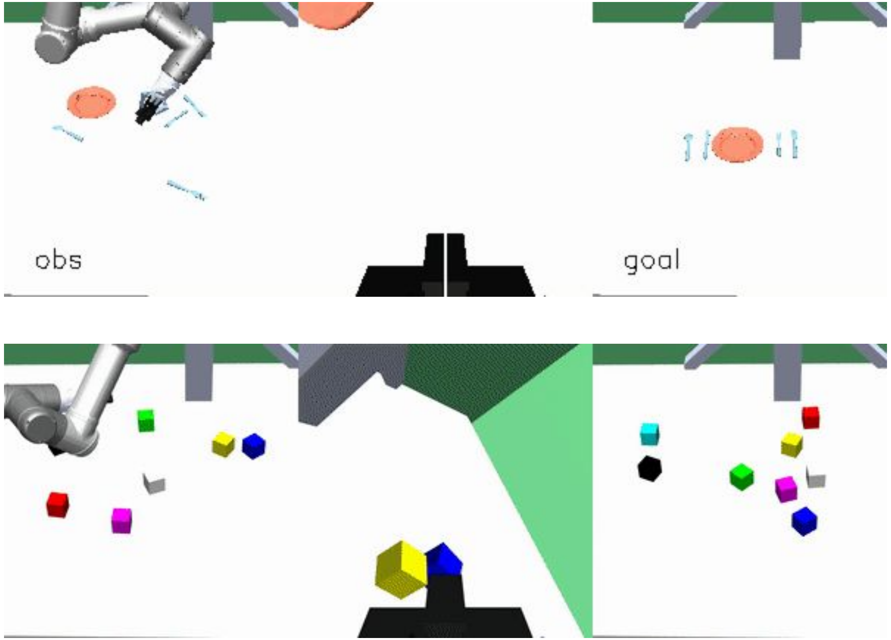

[Video Link]

Supervisor: Daniel Honerkamp - [Link to paper]

10. Language as a Cognitive Tool to Imagine Goals in Curiosity-Driven Exploration

- A DeepRL approach that automatically imagines out-of-distribution goals for goal discovery and open-ended learning
- The model interprets the OOD goal using separation of:
  - Learned goal-achievement reward function
  - Policy relying on deep-sets, gated attention and object-centered representation

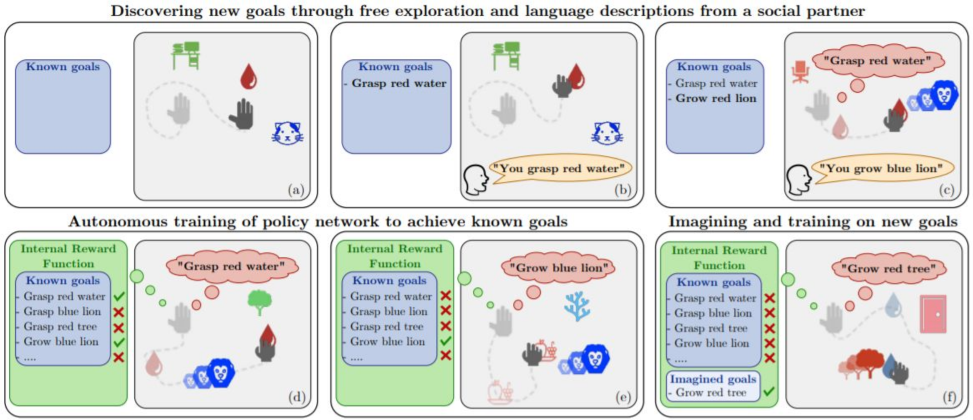

Source: Colas et al.

Supervisor: Daniel Honerkamp - [Link to paper]

10. Language as a Cognitive Tool to Imagine Goals in Curiosity-Driven Exploration

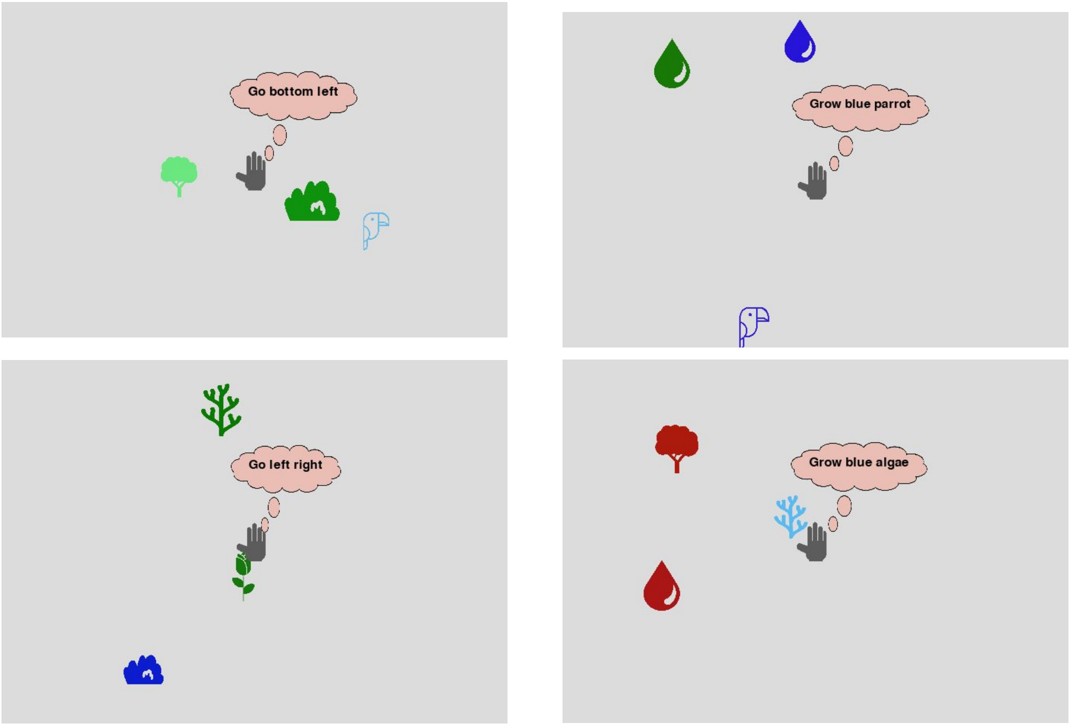

Goals provided by the social partner

Imagined goals
(Soma goals are meaningless)

Source: Colas et al.

Supervisor: Daniel Honerkamp - [Link to paper]

11. Reinforcement Learning with Videos: Combining Offline Observations with Interaction

- Model learns a policy using human demonstration videos along with online data collected by the robot
- Human experience videos don't have reward or motion cues
  - Model estimates a reward using domain-invariant representation of the images
- The model successfully learns challenging vision based skills with much less data

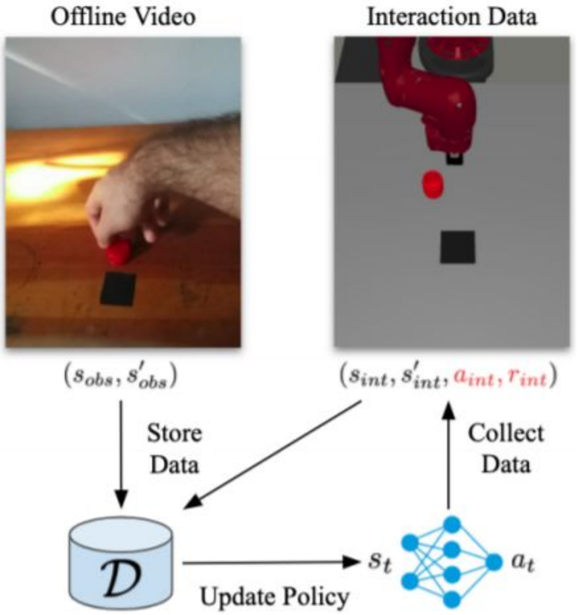

Source: Schmeckpeper et al.

Supervisor: Dr. Tim Welschehold - [Link to paper]

11. Reinforcement Learning with Videos: Combining Offline Observations with Interaction

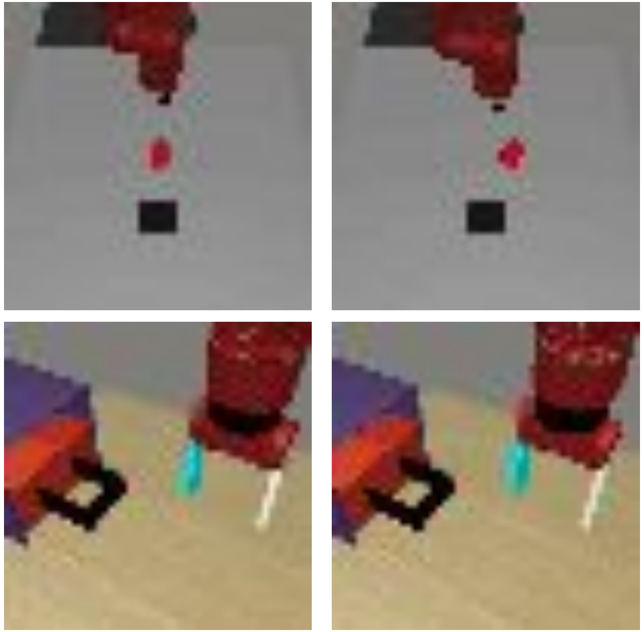

Source: Schmeckpeper et al.

Supervisor: Dr. Tim Welschehold - [Link to paper]

12. A GMM Layer Jointly Optimized with Discriminative Features within a Deep Neural Network Architecture

- Propose a NN with GMM as its final layer and jointly trained using Asynchronous SGD
- The best of both worlds:
  - Superior feature extraction of NNs
  - Well-studied and structured classification of GMMs
- Evaluate the improvement obtained by:
  - Joint training of the GMM and the other layers
  - Changing network depth in the GMM network and the vanilla one
  - Using a GMM model instead of a vanilla NN

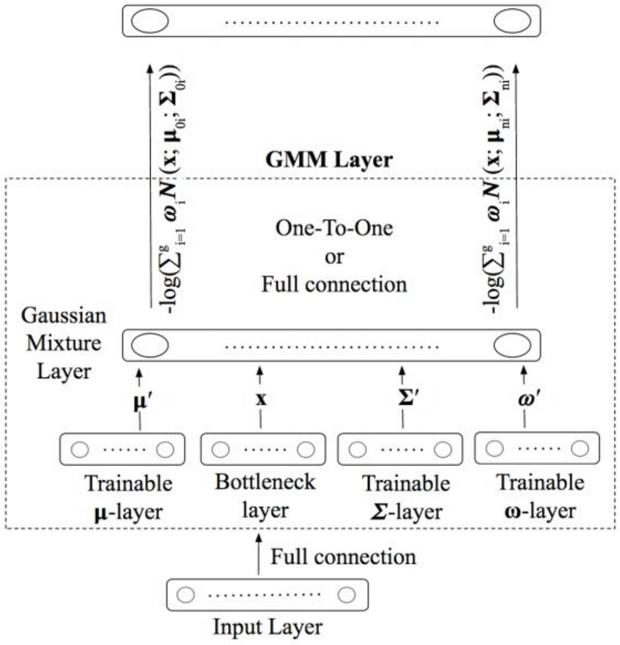

Source: Variani et al.

Supervisor: Dr. Tim Welschehold - [Link to paper]

Thank you!