Robot Learning Seminar
WS 2021/22

Robot Learning Lab

Albert-Ludwigs-Universität Freiburg

Friday, 22nd October 2021

UNI FREIBURG

Evaluation

Evaluation	Due Date
Paper Abstract	14/01/2022
Seminar Presentation	11/02/2022
Paper Summary	25/02/2022

- 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

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

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
Form	Until 25/10/2021	On 26/10/2021	On 26/10/2021

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

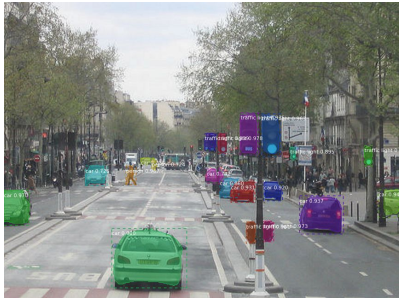

- Noisy observations and sensors

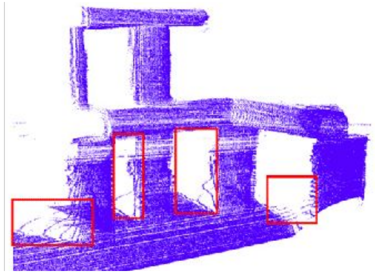

Mask R-CNN for object detection and instance segmentation on Keras and TensorFlow, Waleed et. al., 2017
Rohit Mohan and Abhinav Valada, "EfficientPS: Efficient Panoptic Segmentation", arXiv preprint arXiv:2004.02307, 2020.Patil, Ashok Kumar & Kumar, G Ajay & Kim, Tae-Hyoung & Chai, Young-Ho. (2018). 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

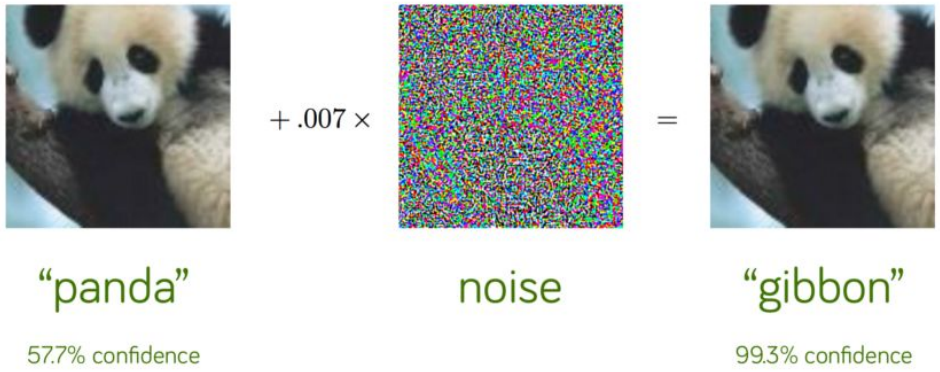

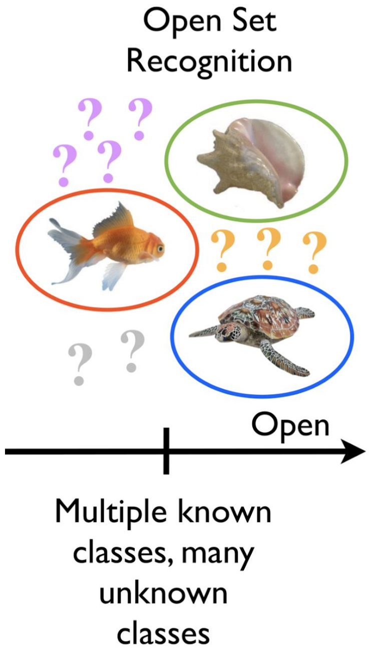

Towards Open Set Recognition, Scheirer et. al., 2012

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

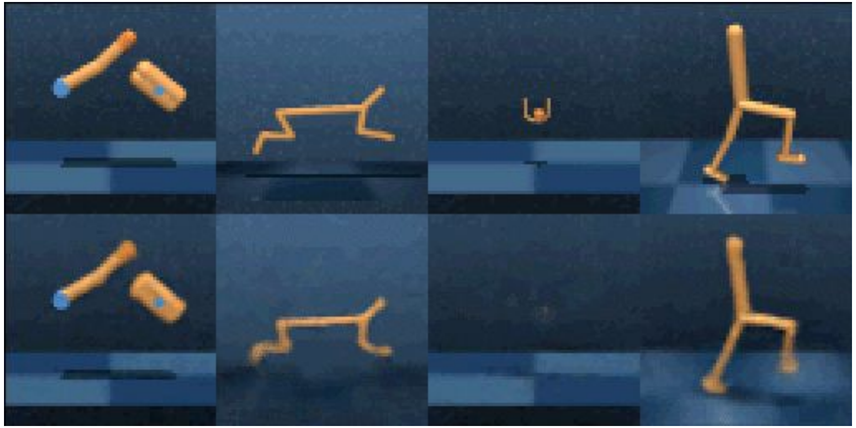

- Model-based RL
  - Learns a world model
  - Promise of better generalisation

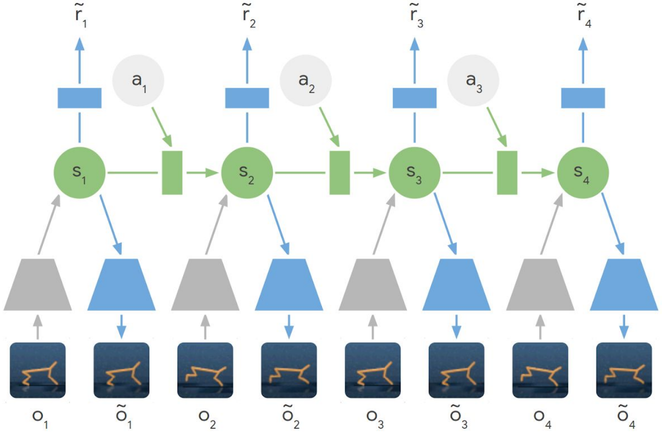

Learning Latent Dynamics for Planning from Pixels, Hafner et. al., 2019

Expensive Real World Data

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

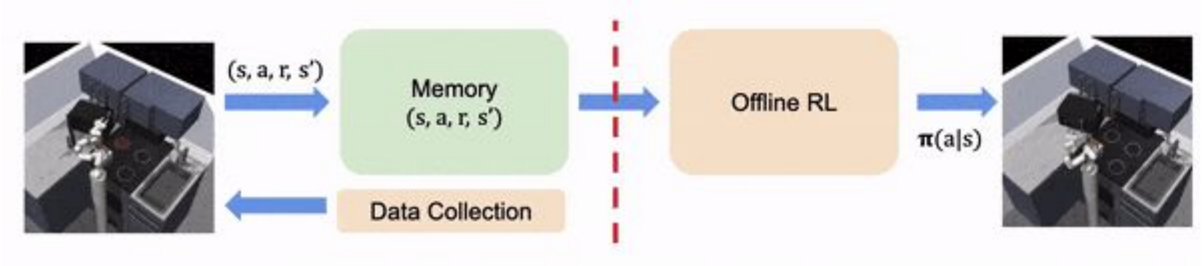

Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection. Sergey Levine, Peter Pastor, Alex Krizhevsky, Deirdre Quillen
D4RL: Datasets for Deep Data-Driven Reinforcement Learning, Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, Sergey Levine

Weak- and Self-Supervision

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

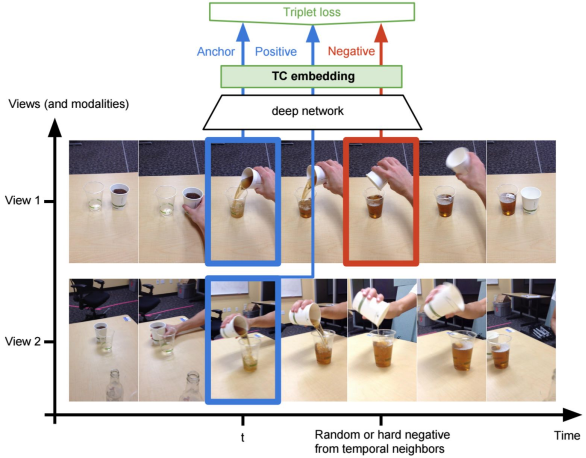

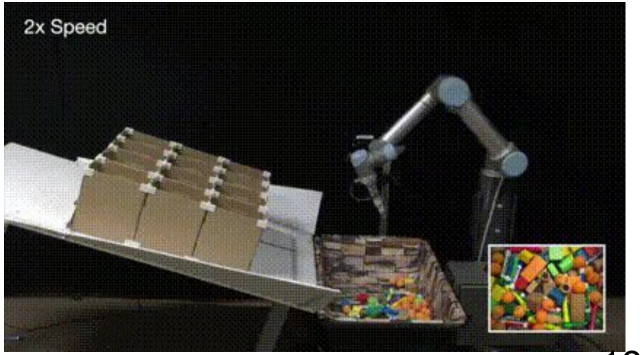

Time-Contrastive Networks: Self-Supervised Learning from Video, Sermanet et. al., 2018
TossingBot: Learning to Throw Arbitrary Objects, Zeng et. al., 2019.

Seminar Topics

1. Implicit Behavioral Cloning

- Behavioral Cloning is one of the easiest methods to learn real-world robotic skills
- Propose policy to be modeled implicitly, not explicitly
- Demonstrate great improvement in contact-rich and discontinuous tasks through experimentation

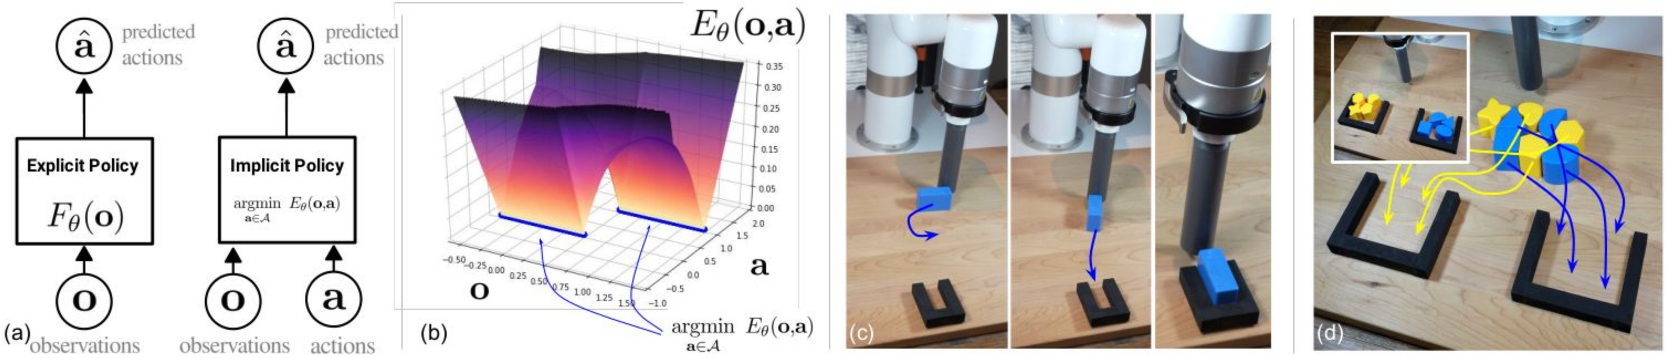

Source: Florence et al.

Supervisor: Eugenio Chisari - [Link to paper]

1. Implicit Behavioral Cloning

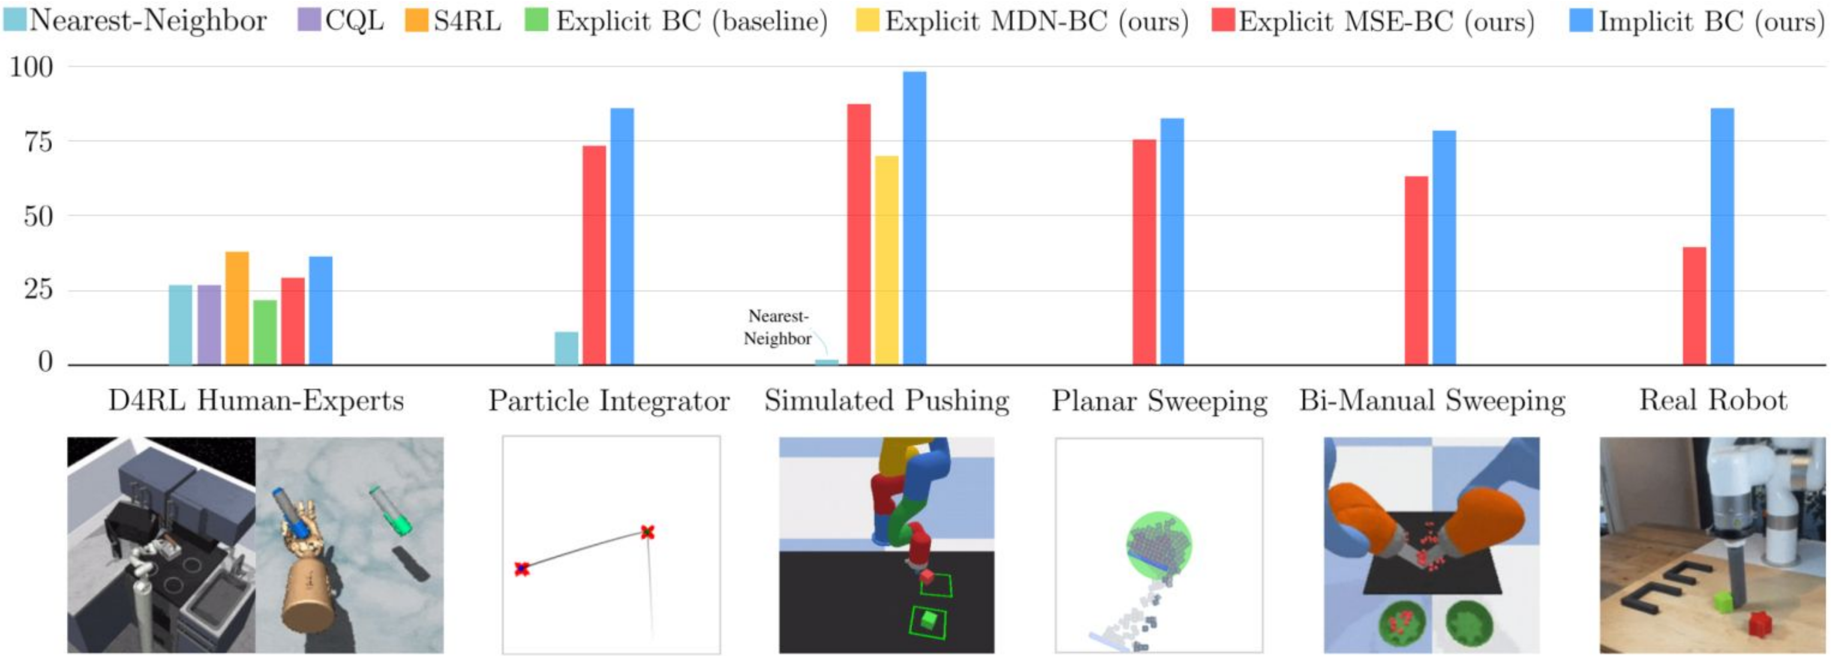

Source: Florence et al.

Supervisor: Eugenio Chisari - [Link to paper]

2. Error-Aware Imitation Learning from Teleoperation Data for Mobile Manipulation

- An approach to efficiently collect human demonstrations for mobile manipulation and learn long-horizon tasks from them
- Learned detection of out-of-distribution states
- Evaluation in a realistic simulated kitchen environment

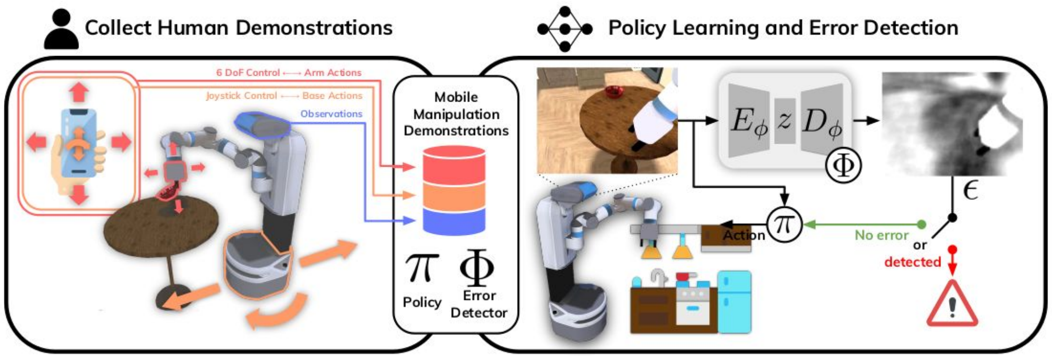

Source: Wong et al.

Supervisor: Eugenio Chisari - [Link to paper]

2. Error-Aware Imitation Learning from Teleoperation Data for Mobile Manipulation

Source: Wong et al.

Supervisor: Eugenio Chisari - [Link to paper]

3. PnPNet: End-to-End Perception and Prediction with Tracking in the Loop

- End-to-end approach for object detection, tracking, and prediction in autonomous driving
- Utilizes joint trajectory representation in tracking and prediction
- Shows greater robustness against long-term occlusions

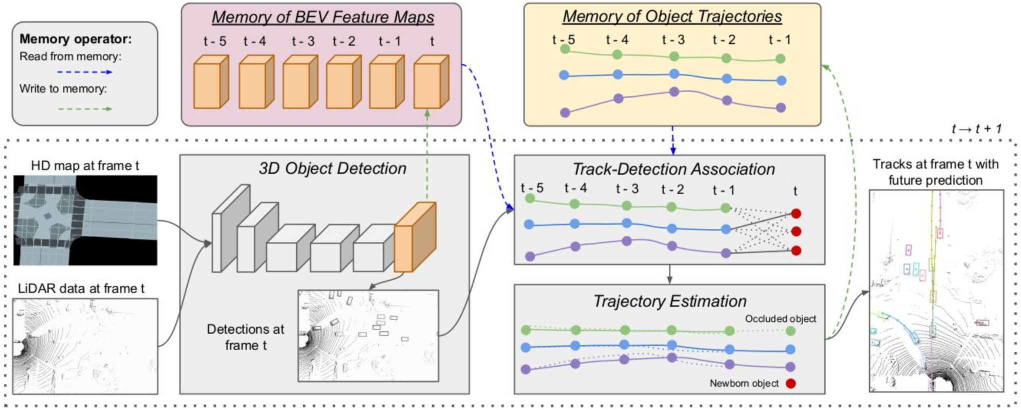

Source: Liang et. al.

Supervisor: Nikhil Gosala - [Link to paper]

3. PnPNet: End-to-End Perception and Prediction with Tracking in the Loop

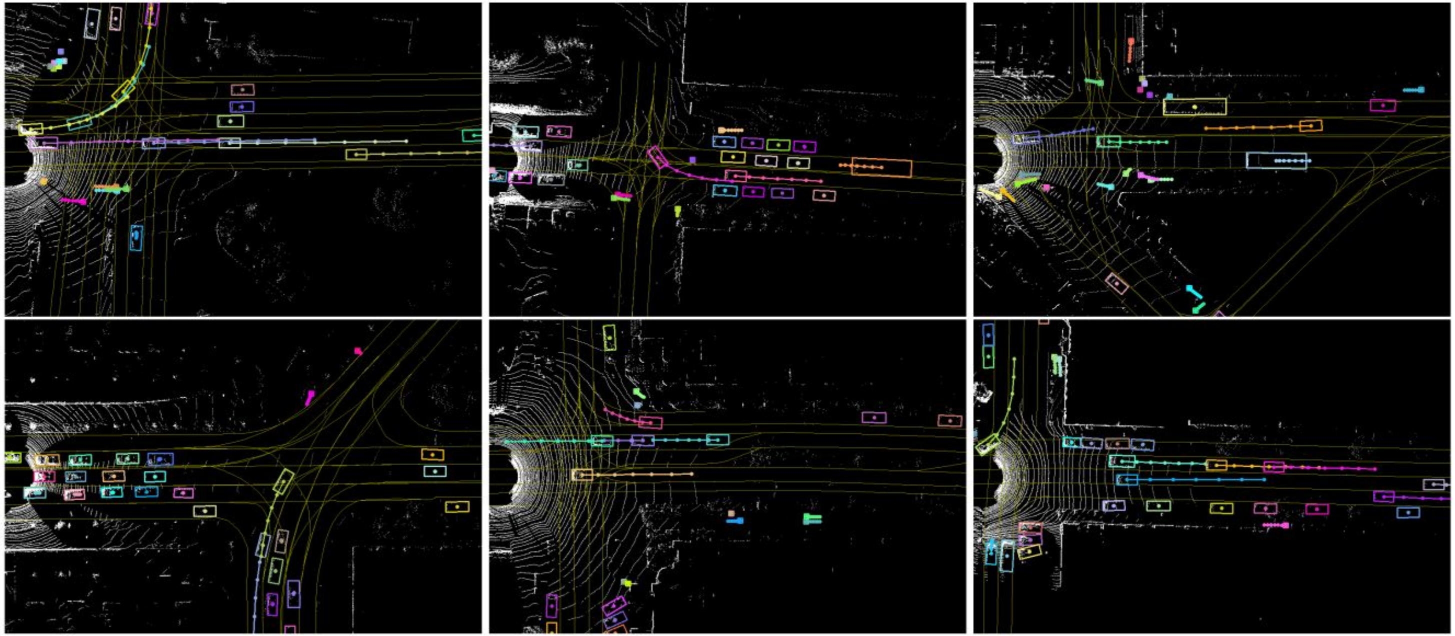

Source: Liang et. al.

Supervisor: Nikhil Gosala - [Link to paper]

4. CertainNet: Sampling-free Uncertainty Estimation for Object Detection

- Approach for assessing uncertainty in each output of a detection network
- Mapping of images to a hyper space, which is used to determine uncertainty in objectness, location, size, and class
- Transfers better to out-of-domain data than competing approaches

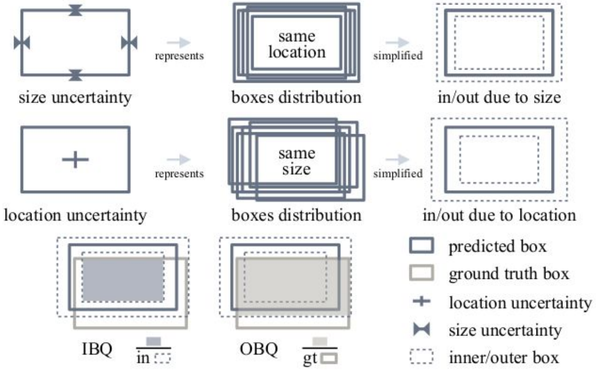

Source: Liang et. al.

Supervisor: Nikhil Gosala - [Link to paper]

4. CertainNet: Sampling-free Uncertainty Estimation for Object Detection

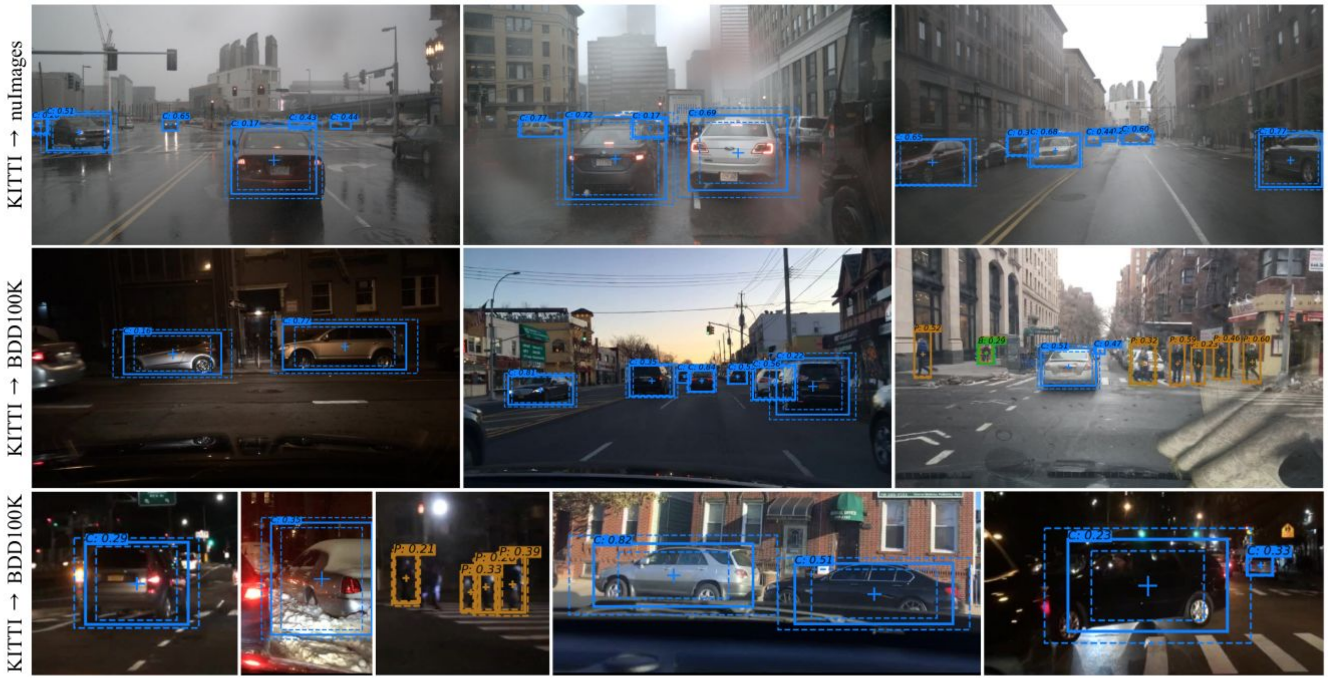

Source: Gasperini et. al.

Supervisor: Nikhil Gosala - [Link to paper]

5. XIRL: Cross-embodiment Inverse Reinforcement Learning

- Addresses the problem of learning a task from demonstrations of differently embodied agents
- A task progress encoder is used to assess task completion
- Delta in completion provides dense reward for RL task under a given embodiment

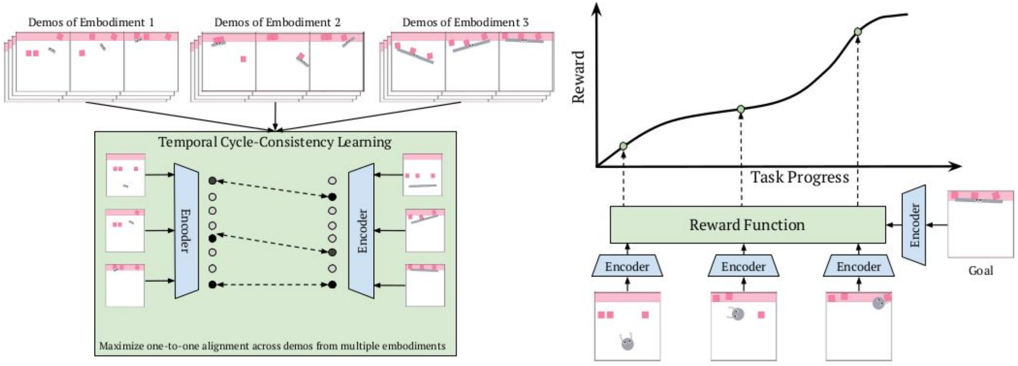

Source: Zakka et al.

Supervisor: Daniel Honerkamp - [Link to paper]

5. XIRL: Cross-embodiment Inverse Reinforcement Learning

Source: Zakka et al.

Supervisor: Daniel Honerkamp - [Link to paper]

6. Explore and Control with Adversarial Surprise

- Approach to avoid reward engineering by providing intrinsic motivation for exploration based on an adversarial game
- Two policies compete, attempting to either minimize or maximize *surprise*
- The approach learns more quickly and explores more thoroughly than competing baselines

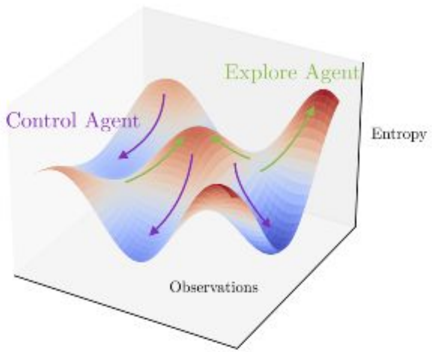

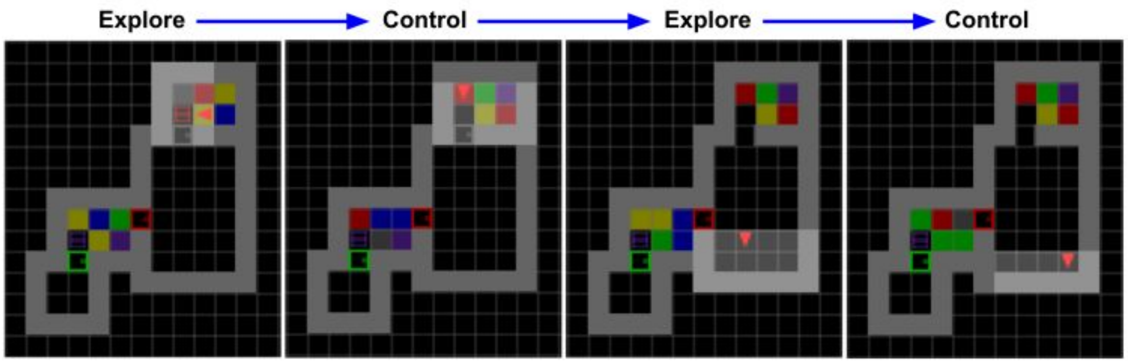

Source: Fickinger et al.

Supervisor: Daniel Honerkamp - [Link to paper]

6. Explore and Control with Adversarial Surprise

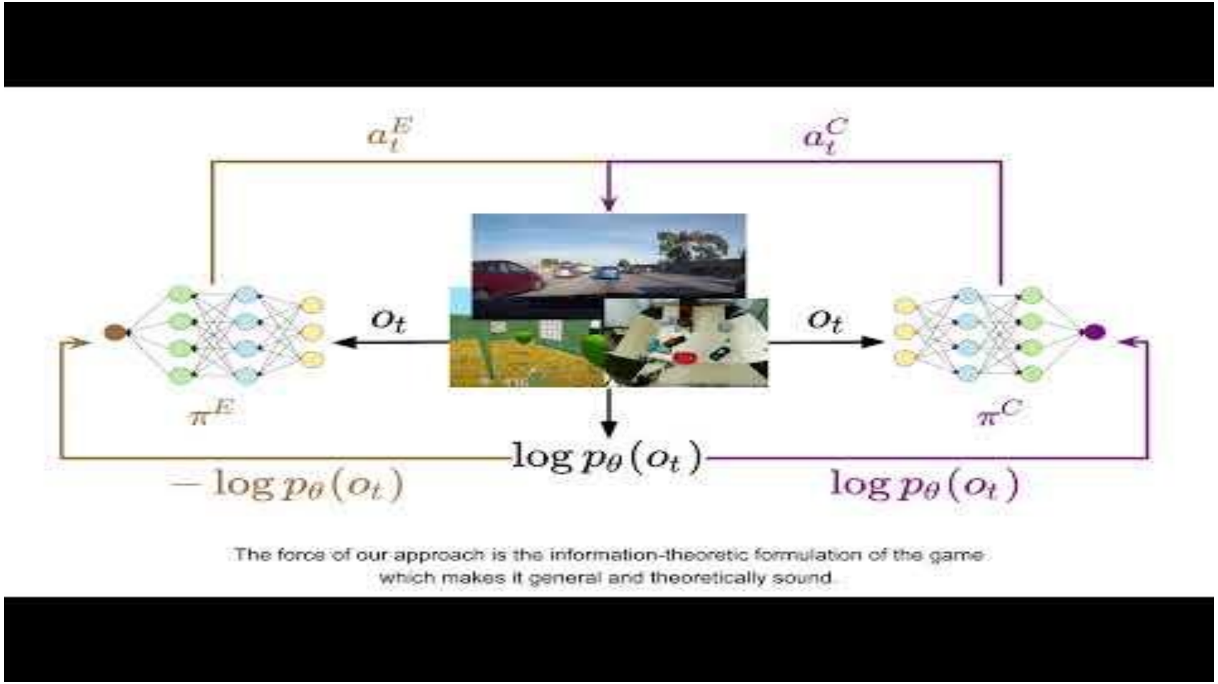

Source: Fickinger et al.

Supervisor: Daniel Honerkamp - [Link to paper]

7. Weakly Supervised Multi-Object Tracking and Segmentation

- Addresses high annotation cost for instance segmentation data
- Extracts heatmaps from classification head of network to use as foreground cues
- Uses a supervised learning task to support an unsupervised one

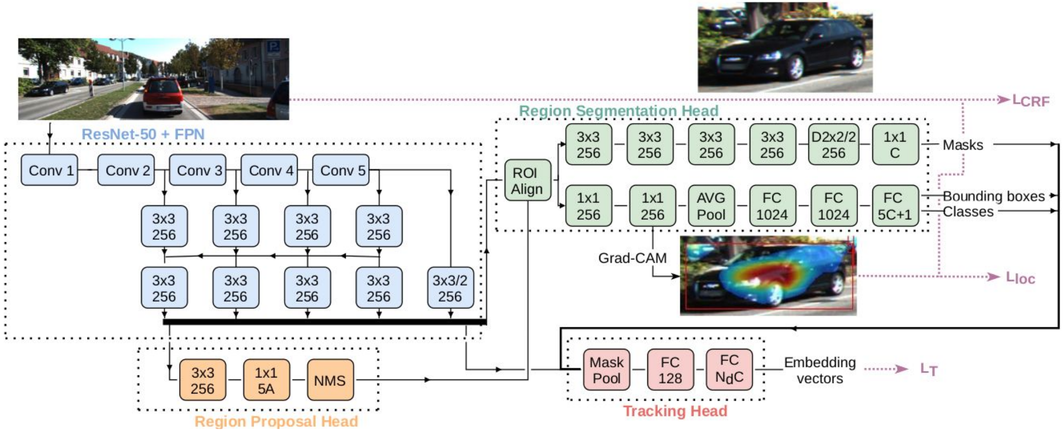

Source: Ruiz et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

7. Weakly Supervised Multi-Object Tracking and Segmentation

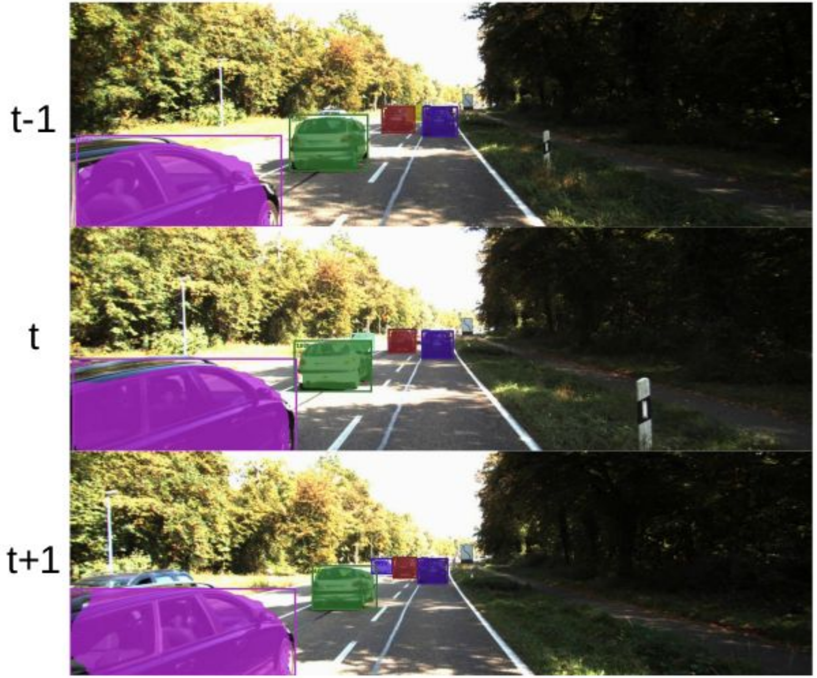

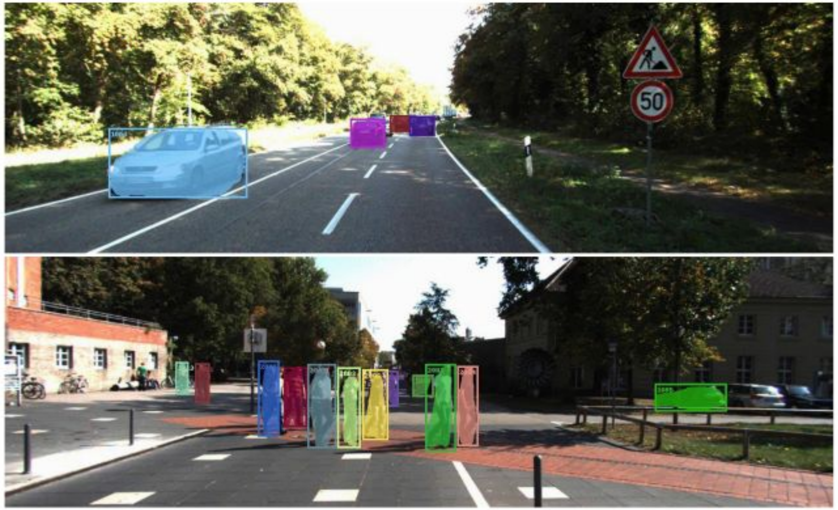

Source: Ruiz et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

8. Instance-Aware Predictive Navigation in Multi-Agent Environments

- Approach for predicting motions in multi agent scenarios and using this prediction to select safe control actions
- Learn to predict multiple hypothetical scenes including agent interactions
- Two-Stage action selection strategy based on prediction capability

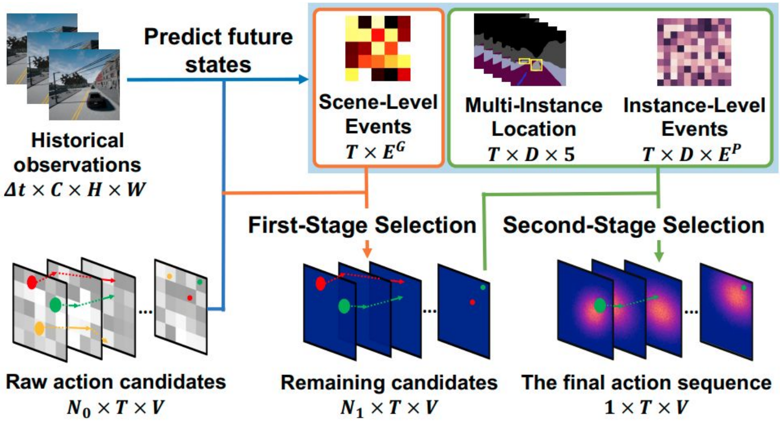

Source: Cao et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

8. Instance-Aware Predictive Navigation in Multi-Agent Environments

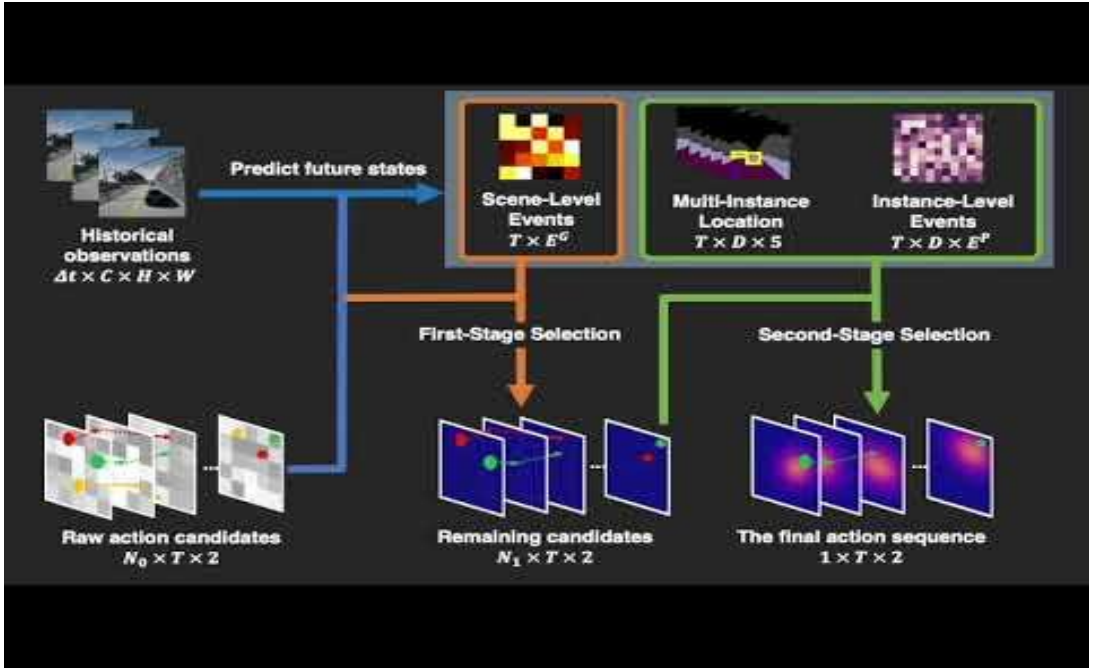

Source: Cao et al.

Supervisor: Juana Valeria Hurtado - [Link to paper]

9. Quasi-Dense Similarity Learning for Multiple Object Tracking

- Approach to integrating visual similarity as strong prior in tracking objects
- Learn dense similarity metric between RoIs generated by proposal network
- Nearest neighbor search in embedding space of similarity metric is used to associate regions and match them

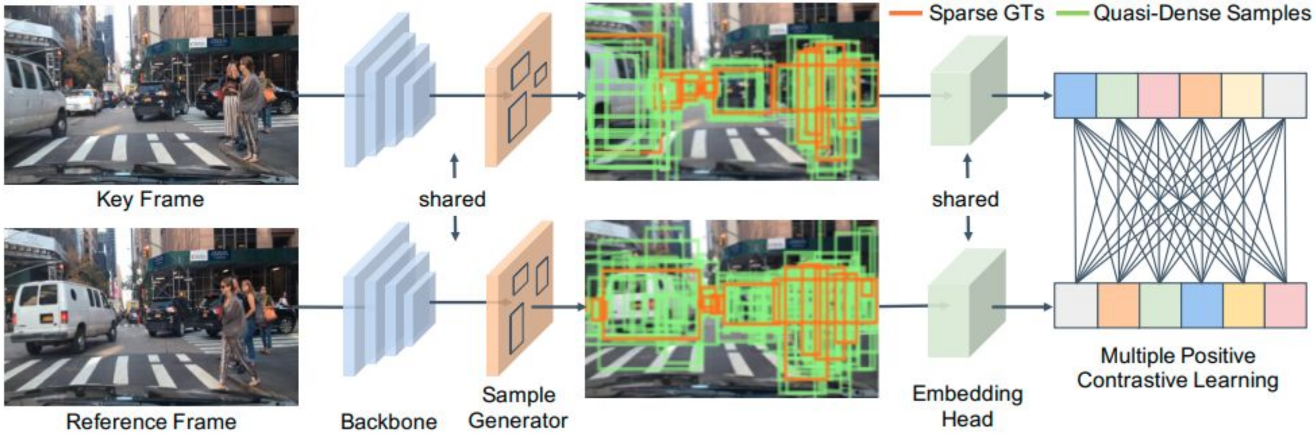

Source: Pang et al.

Supervisor: Rohit Mohan - [Link to paper]

9. Quasi-Dense Similarity Learning for Multiple Object Tracking

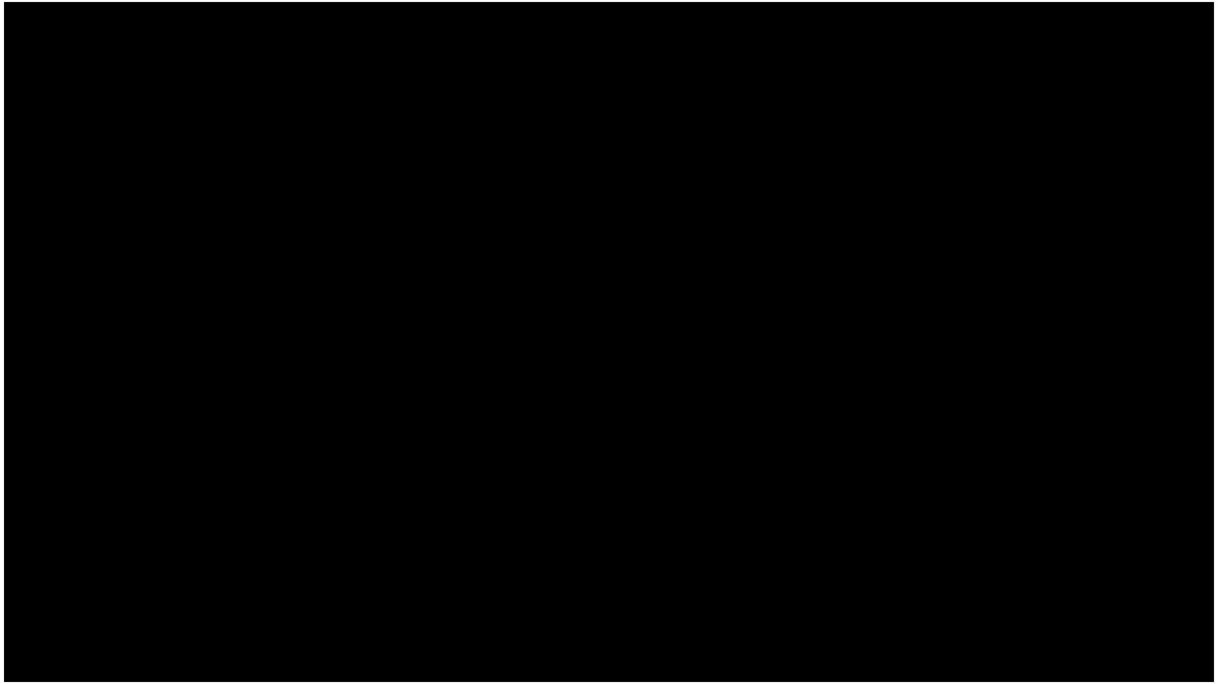

Source: Pang et al.

Supervisor: Rohit Mohan - [Link to paper]

10. Rethinking Space-Time Networks with Improved Memory Coverage for Efficient Video Object Segmentation

- Approach for video object segmentation
- Learn general RGB correspondence as opposed to object-wise correspondence
- Use and analyze uncommon similarity function

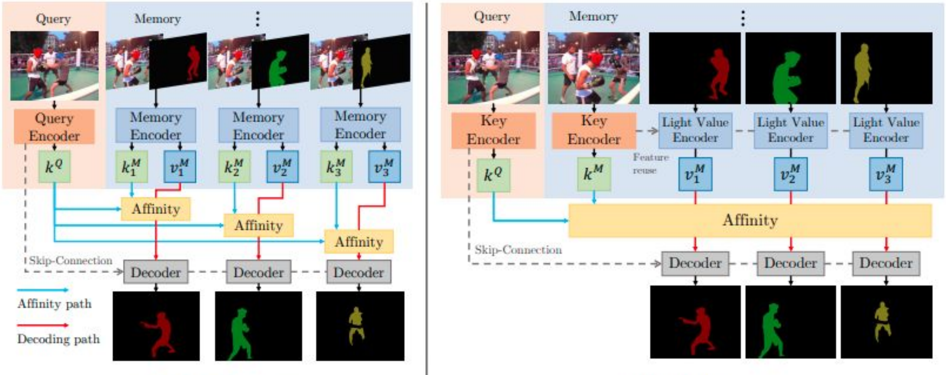

(a) STM and variants (b) STCN (Ours)

Source: Cheng et al.

Supervisor: Rohit Mohan - [Link to paper]

10. Rethinking Space-Time Networks with Improved Memory Coverage for Efficient Video Object Segmentation

Source: Cheng et al.

Supervisor: Rohit Mohan - [Link to paper]

11. GTI: Learning to Generalize Across Long-Horizon Tasks from Human Demonstrations

- Approach to generating policies for novel tasks from separate demonstrations in imitation learning
- Exploit overlap of tasks in latent space to explore task space
- Learned policies are reactive and robust

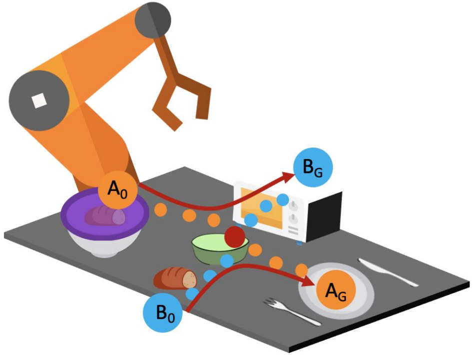

Source: Mandlekar et al.

Supervisor: Adrian Röfer - [Link to paper]

11. GTI: Learning to Generalize Across Long-Horizon Tasks from Human Demonstrations

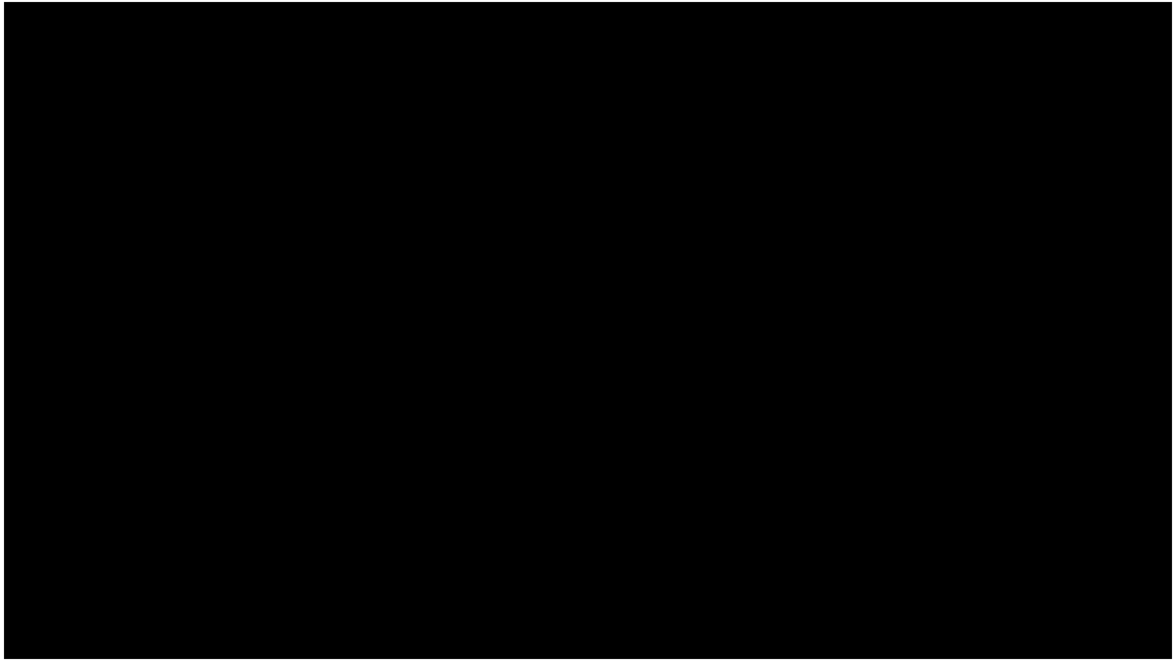

Source: Mandlekar et al.

Supervisor: Adrian Röfer - [Link to paper]

12. Hierarchical Planning for Long-Horizon Manipulation with Geometric and Symbolic Scene Graphs

- Visually grounded long-horizon planning of rearrangement tasks
- Use visually grounded scenegraph representation to structure task state
- Train two policies: One for high-level planning, one for immediate action selection

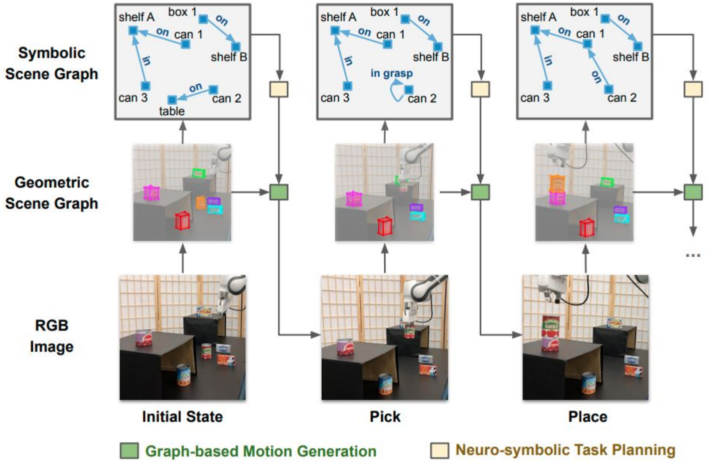

Source: Zhu et al.

Supervisor: Adrian Röfer - [Link to paper]

12. Hierarchical Planning for Long-Horizon Manipulation with Geometric and Symbolic Scene Graphs

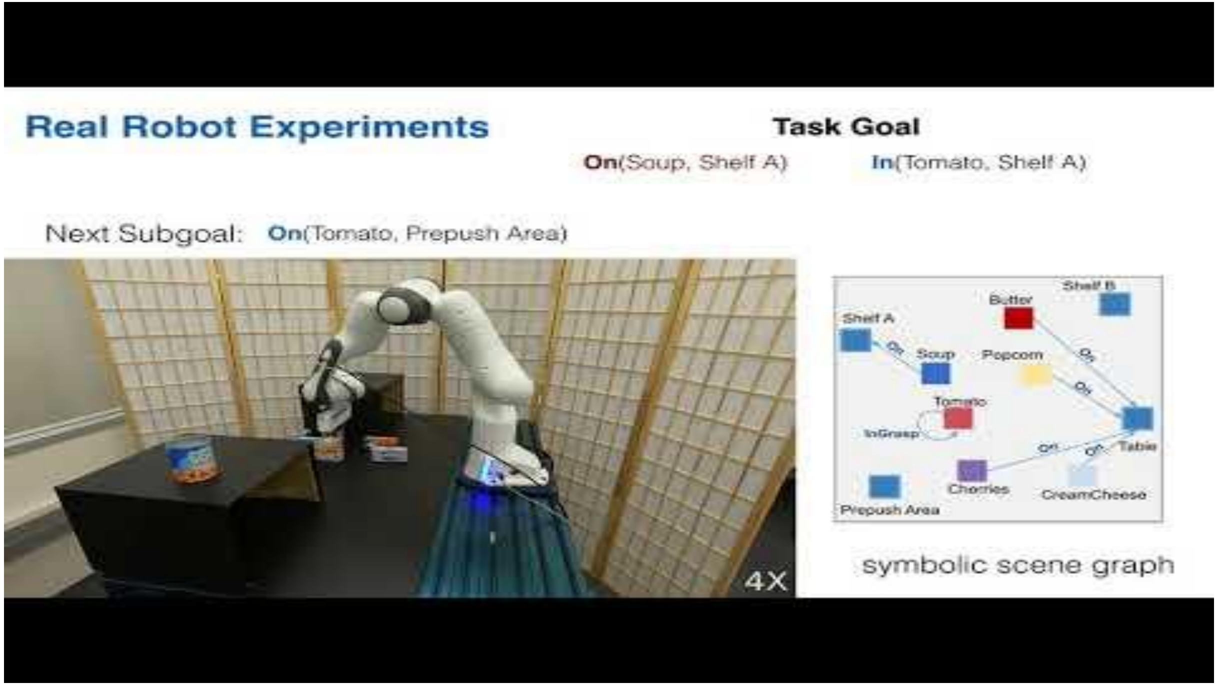

Source: Zhu et al.

Supervisor: Adrian Röfer - [Link to paper]

Questions