Leveraging Structural Information for Safe, Efficient, and Adaptable Robot Learning

Raffaello Camoriano
raffaello.camoriano@polito.it

Assistant Professor / RTDa @ Politecnico di Torino – Turin, Italy
Affiliated Researcher @ Istituto Italiano di Tecnologia – Genoa, Italy
Visual And Multimodal Applied Learning Laboratory - AI-Hub @ PoliTo - ELLIS Unit Turin

Visiting Scholar Seminar Series, Lund, Sweden
4/12/2025

Finanziato dall'Unione europea
NextGenerationEU

About Me

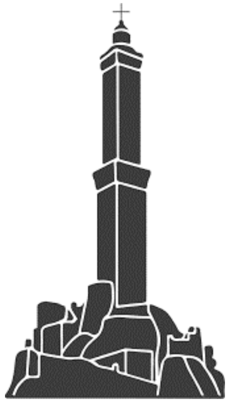

Genoa

2008 – 2013

B. Sc. in Computer Engineering
M. Sc. in Robotics Engineering

Robotics, ML,
Robot Grasping,
Underwater Robotics

2014 – 2017

Ph. D. in Machine Learning & Robotics

Large-scale Learning,
Humanoid Robotics,
Lifelong Learning

2017 - 2022

PostDoc

Reinforcement Learning,
Structured Prediction,
Large-scale Learning

2022 - now

Assistant Professor

3D Deep Learning,
Robot Learning,
Reinforcement Learning

Turin

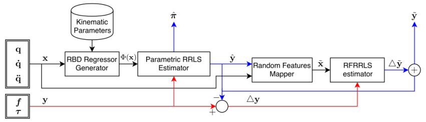

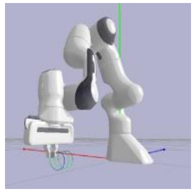

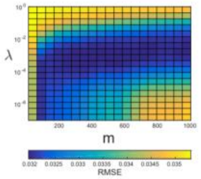

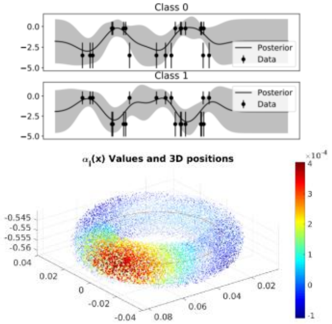

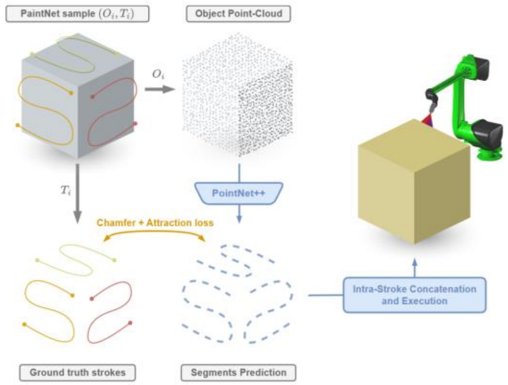

About Us

Barbara Caputo
Full
Professor

Tatiana Tommasi
Full
Professor

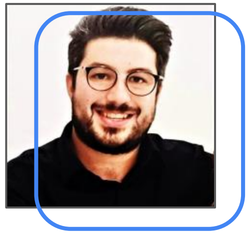

Giuseppe Averta
Assistant
Professor (TT)

Carlo Masone
Assistant
Professor

Raffaello Camoriano
Assistant
Professor

Francesca Pistilli
Assistant
Professor

- 6 faculties
- 16 PhD students
- 4 research interns (post-graduate scholars)
- 10 master students

About Us

Barbara Caputo
Full
Professor

Tatiana Tommasi
Full
Professor

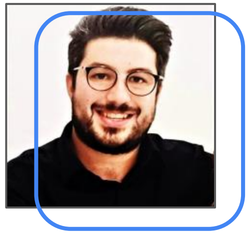

Giuseppe Averta
Assistant
Professor (TT)

Carlo Masone
Assistant
Professor

Raffaello Camoriano
Assistant
Professor

Francesca Pistilli
Assistant
Professor

- 6 faculties
- 16 PhD students
- 4 research interns (post-graduate scholars)
- 10 master students

Machine Learning Courses:
training more than 600 students every year

- Machine Learning and Deep Learning
  Msc Data Science and Engineering
- Advanced Machine Learning
- Robot Learning
  Msc Computer Engineering
- Data Analysis and Artificial Intelligence
  Msc Automotive Engineering
- Machine Learning for Mathematical Engineering
  Msc Mathematical Engineering
- AI for Geospatial Data
  Bsc Construction Engineering, Engineering and Management

Recent Publication Records
2025: 3 IEEE CVPR, 3 IEEE ICCV, 1 ICML, 1 ICLR, 1 T-RO, 2 IEEE IROS, 2 IEEE RAL, 1 TMLR
2024: 6 IEEE CVPR, 2 ECCV, 2 ICML, 1 ICLR, 1 IEEE T-PAMI, 2 IJCV
2023: 1 IEEE CVPR, 4 IEEE ICCV, 1 IJRR, 2 IEEE IROS, 3 IEEE WACV, 2 IEEE RAL
2022: 4 IEEE CVPR, 2 IEEE ICCV, 3 IEEE IROS, 5 IEEE WACV, 1 ICML, 1 NeurIPS, 1 IEEE RAL

Industrial Collaborations / Research Contracts
- ASI - Agenzia Spaziale Italiana
- EFORT Europe
- COMAU
- STMicroelectronics
- FCA-CRF
- Italdesign
- Reply
- …

EU Projects
- FAIR - Future Artificial Intelligence Research (PNRR)
- HD-MOTION - EU Seal Of Excellence (SOE)
- EXPAND - European Digital Innovation Hub (EDIH)
- EU-ELISE - European Network of AI excellence centers
- EU-ELSA - European Lighthouse on Secure and Safe AI

Alumni
Marco Ciccone – Vector Institute, Toronto
Gabriele Berton – Research Fellow at Amazon
Francesco Cappio Borlino – Research Scientist at Amazon
Dario Fontanel – Machine Learning Engineer at Snap Inc.
Silvia Bucci – Argotec
Mirko Planamente – ARGO Vision

Fabio Frattin – Head of Product & Co-Founder of Algor
Fabio Cermelli – CTO and Co-Founder of FocoosAI
Antonio Tavera – CEO and Co-Founder of FocoosAI
Chiara Plizzari – Assistant Professor at Bocconi University
Gabriele Tiboni – Postdoc at JMU Würzburg & TU Darmstadt
Eros Fanì - Machine Learning Researcher, Gensyn

National Ph. D. in AI
Dottorato Nazionale in IA

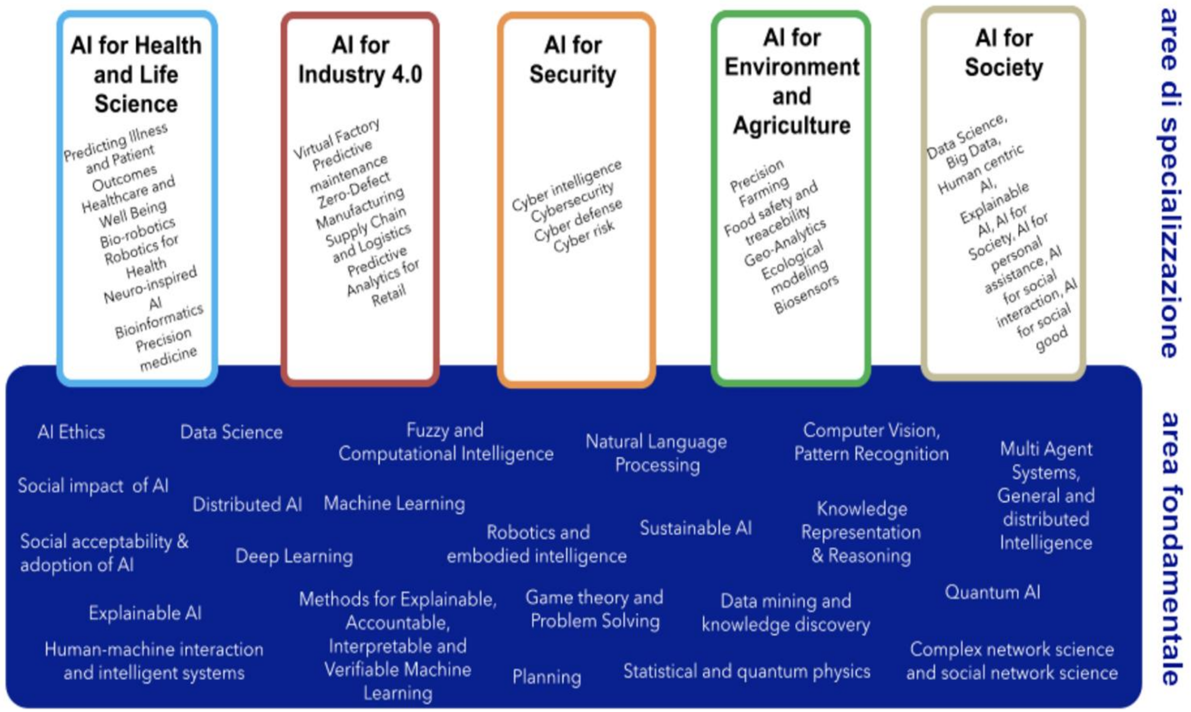

AI for Industry 4.0 - Lead University: Politecnico di Torino.

Participants Universities and Research Institutions: Politecnico di Torino, CNR, Università di Bologna, Politecnico di Milano, Università di Milano, Università di Milano Bicocca, Università di Padova, Università Ca' Foscari Venezia, Università di Verona **Associated Universities and Research Institutions**: Università dell'Aquila, Università di Ferrara, Università di Genova, Università del Molise, Università di Torino, Università Politecnica delle Marche.

ellis unit | TURIN

The European Laboratory for Learning and Intelligent Systems aims at creating new working environments for outstanding researchers to enable them to combine cutting-edge research paired with the creation of start-ups and industrial impact.

ELLIS Units: Alicante, Amsterdam, Barcelona, Berlin, Cambridge, Copenhagen, Darmstadt, Delft, Edinburgh, Lausanne (EPFL), Zürich (ETH), Freiburg, Genoa, Graz, Haifa (Technion), Heidelberg, Helsinki, Jena, Vienna (IST Austria), Leuven, Linz, Lisbon, Lviv, London (UCL), Madrid, Manchester, Milan, Modena (Unimore), Munich, Nijmegen, Oxford, Paris, Potsdam, Prague, Saarbrücken, Stuttgart, Tel Aviv, Trento, Tübingen, Turin, Sofia, Warsaw

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

Robust Models for Urban Scenes

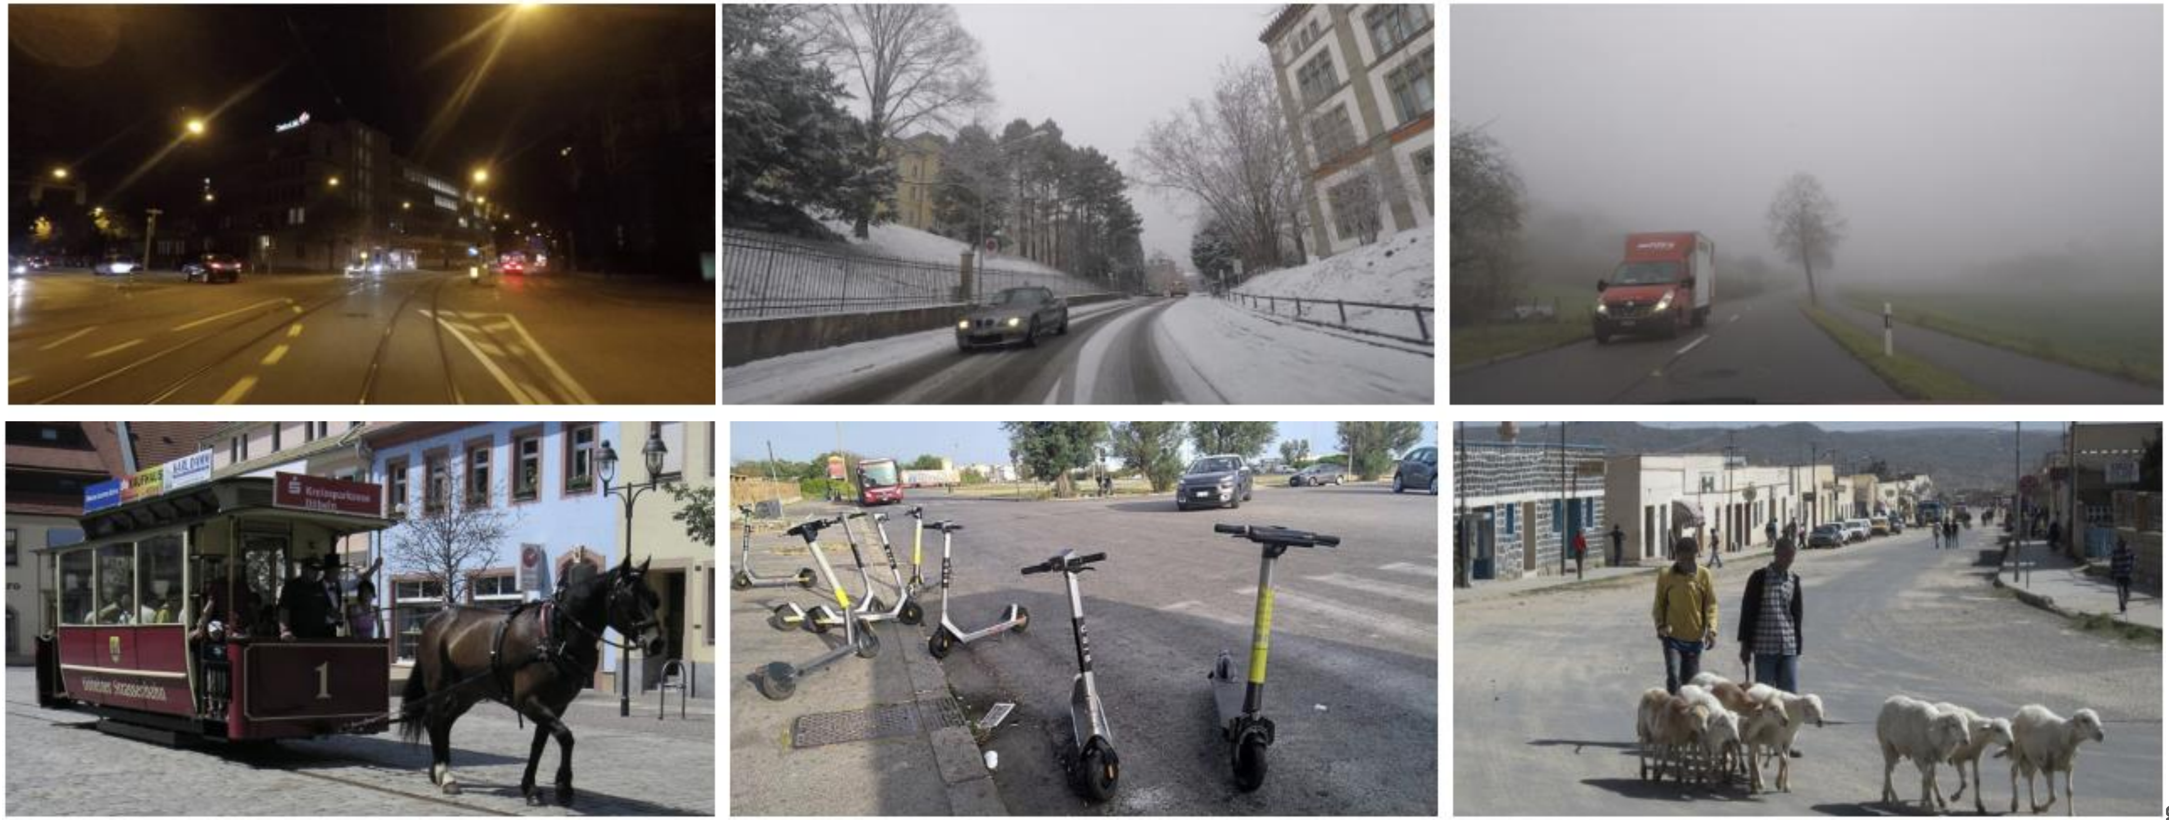

[Images from the Bravo Challenge at ECCV 2024]

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

Semantic Segmentation and Anomaly Detection

Mask2anomaly: Mask transformer for universal open-set segmentation, IEEE T-PAMI 2024
Unmasking Anomalies in Road-Scene Segmentation, IEEE CVPR 2023
Comformer: Continual learning in semantic and panoptic segmentation, IEEE CVPR 2023
Modeling the background for incremental and weakly-supervised semantic segmentation, IEEE T-PAMI 2021

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

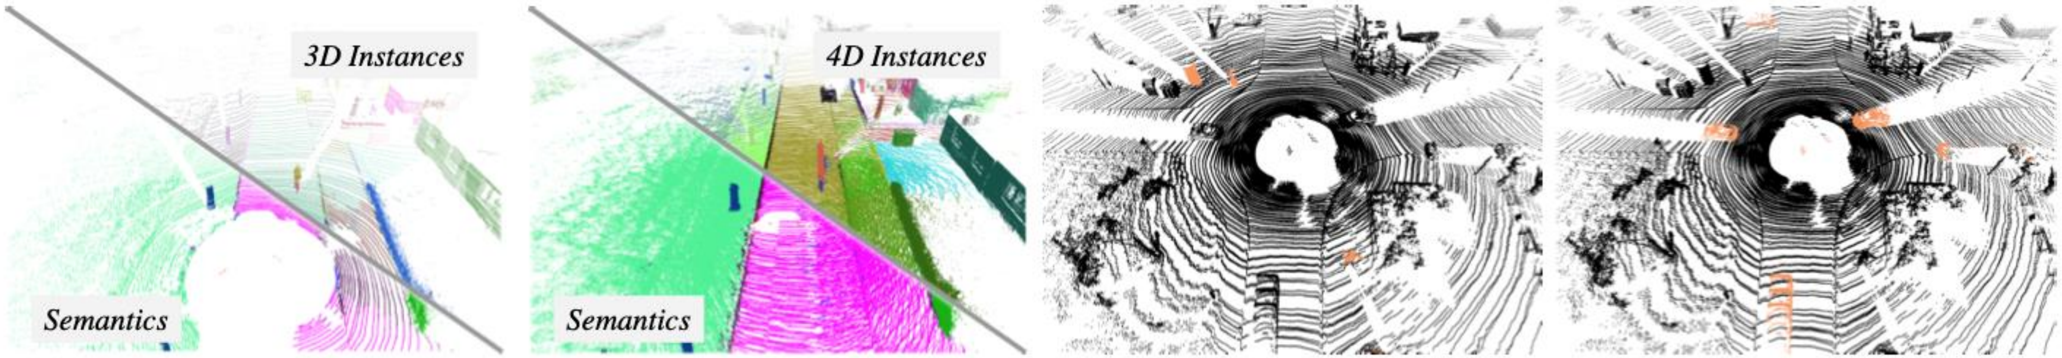

Prior work: Zero-Shot (3D) Lidar Panoptic Segmentation

This work: Zero-Shot 4D Lidar Panoptic Segmentation

Text prompts: {advertising stand}

Text prompts: {car}

[Images fromZero-Shote 4D Lidar Panoptic Segmentation - NVIDIA paper at CVPR 2025]

Moving towards the integration of 3D, 4D and language

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

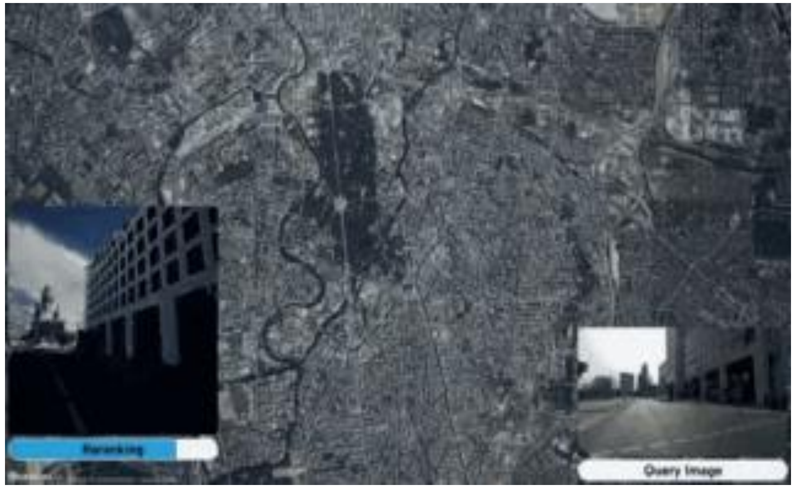

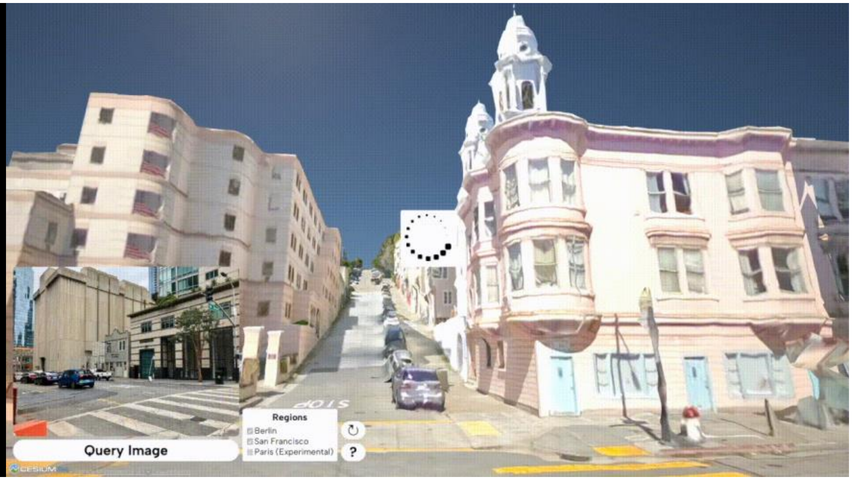

Geolocalization from 2D and 3D data

EarthLoc: Astronaut Photography Localization by Indexing Earth from Space, IEEE CVPR 2024
MeshVPR: Citywide Visual Place Recognition Using 3D Meshes, ECCV 2024
JIST: Joint Image and Sequence Training for Sequential Visual Place Recognition, IEEE RAL 2024
EigenPlaces: Training Viewpoint Robust Models for Visual Place Recognition, IEEE ICCV 2023

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

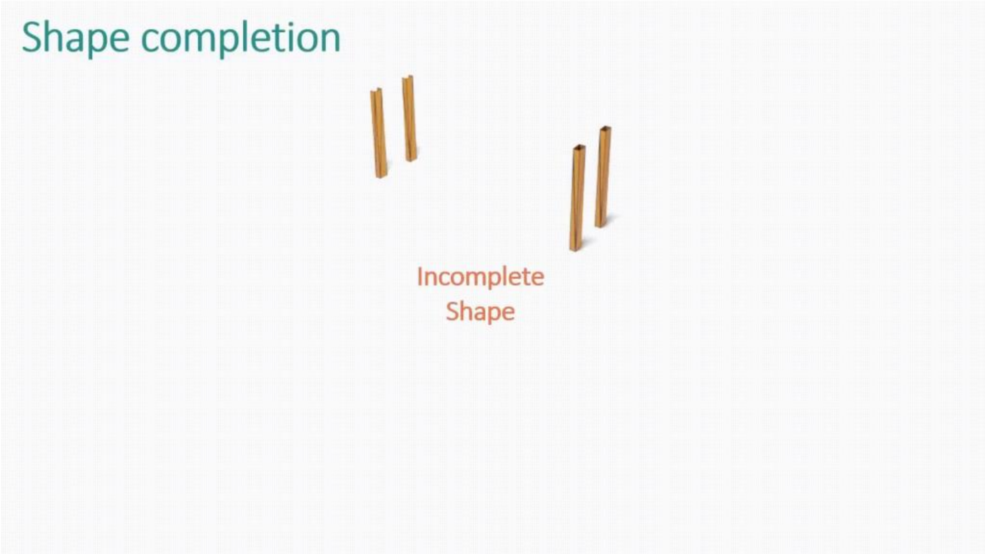

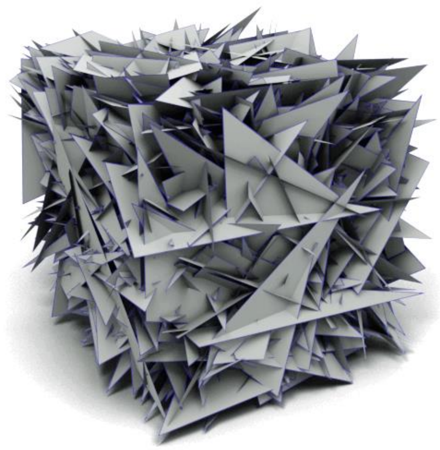

Generative AI on 3D data

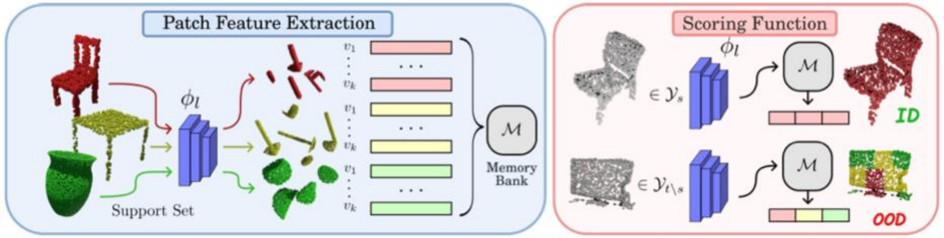

Defect and Novelty Detection on Point Clouds and Multi-View data

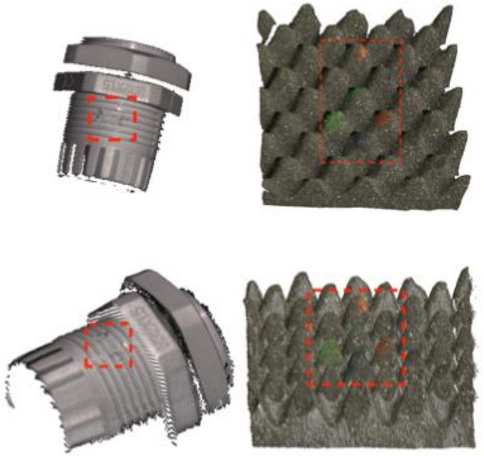

[Images from CVPR-W 2023]

Meshgpt: Generating triangle meshes with decoder-only transformers, IEEE CVPR 2024
Polydiff: Generating 3d polygonal meshes with diffusion models, ArXiv 2024
3D Semantic Novelty Detection via Large-Scale Pre-Trained Models, IEEE Access 2024
3dos: Towards 3d open set learning-benchmarking and understanding semantic novelty detection on point clouds, NeurIPS 2022

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

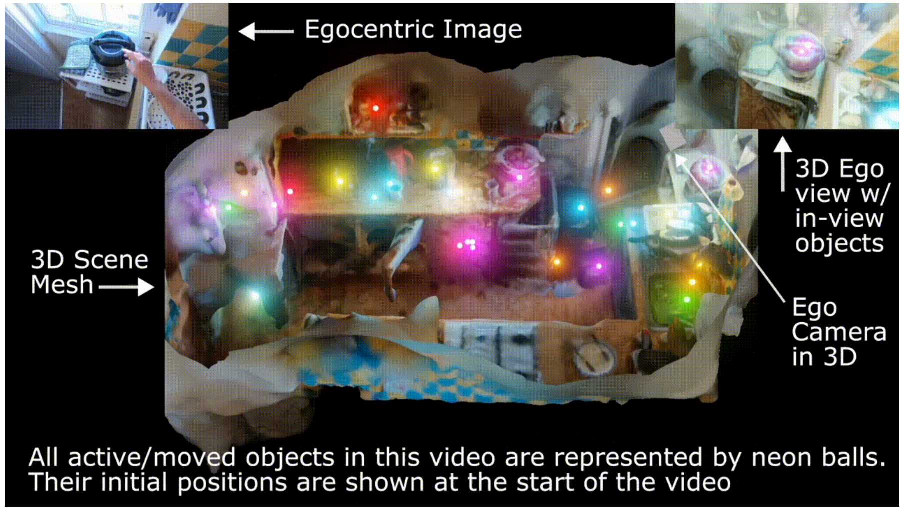

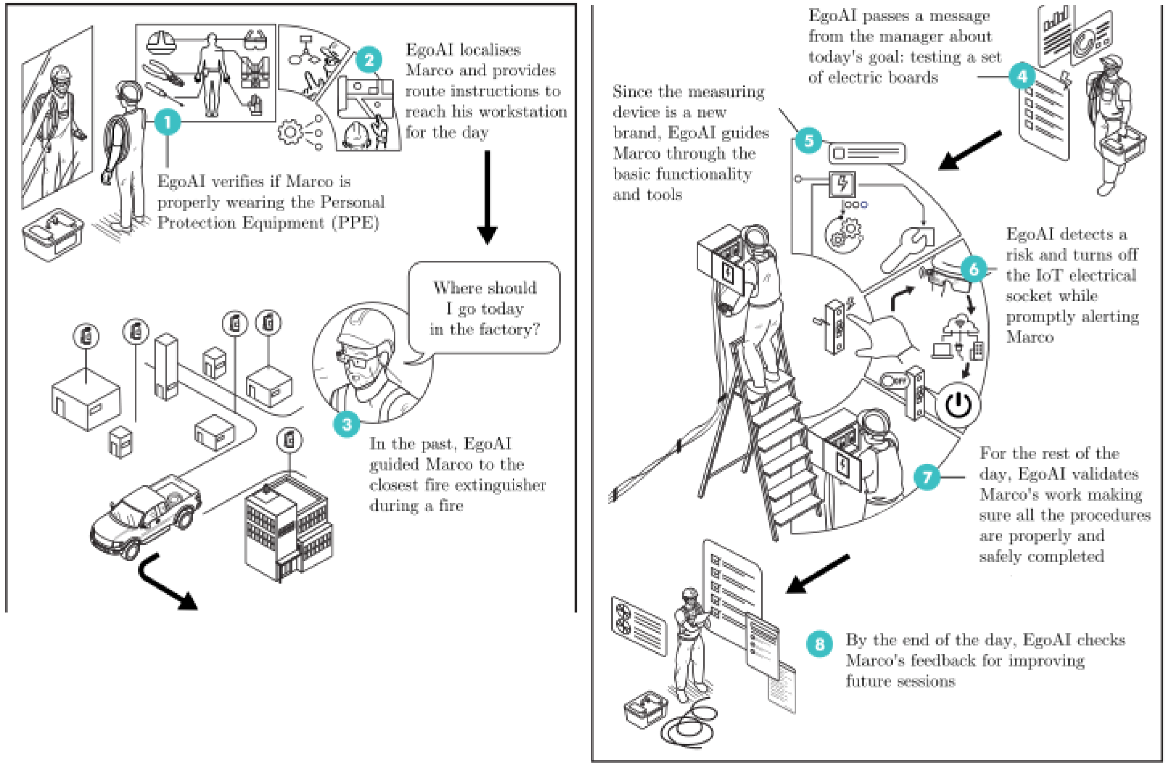

Egocentric Video Understanding

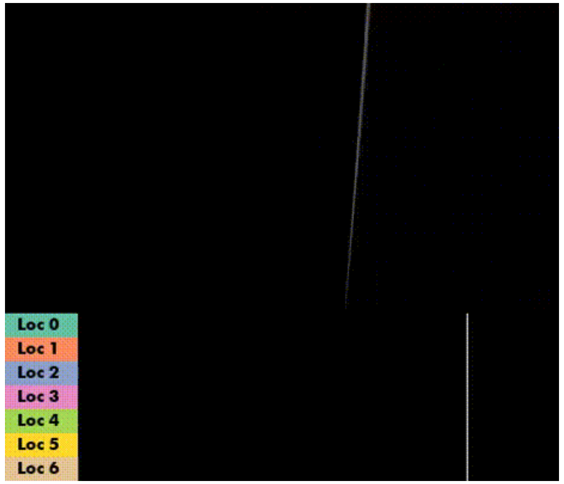

HiERO: understanding the hierarchy of human behavior enhances reasoning on egocentric videos, IEEE ICCV 2025
Amego: Active memory from long egocentric videos, ECCV 2024
A Backpack Full of Skills: Egocentric Video Understanding with Diverse Task Perspectives, IEEE CVPR 2024
Domain generalization using action sequences for egocentric action recognition, Pattern Recognition Letters 2025
Spatial Cognition from Egocentric Video: Out of Sight, Not Out of Mind, 3DV 2025
An Outlook into the Future of Egocentric Vision, IJCV 2024

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

Multi-modal (Video and Language)

SAMWISE in action

Prompt: *"The tall man in grey"*

Prompt: *"The man packing stuff"*

Prompt: *"The stick near the table"*

Prompt: *"The carpet"*

The male child moving leftward while guiding his bicycle

The bicycle moving back and forth in the living room

SAMWISE: Infusing Wisdom in SAM2 for Text-Driven Video Segmentation ,CVPR 2025

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

Distributed, Federated Learning

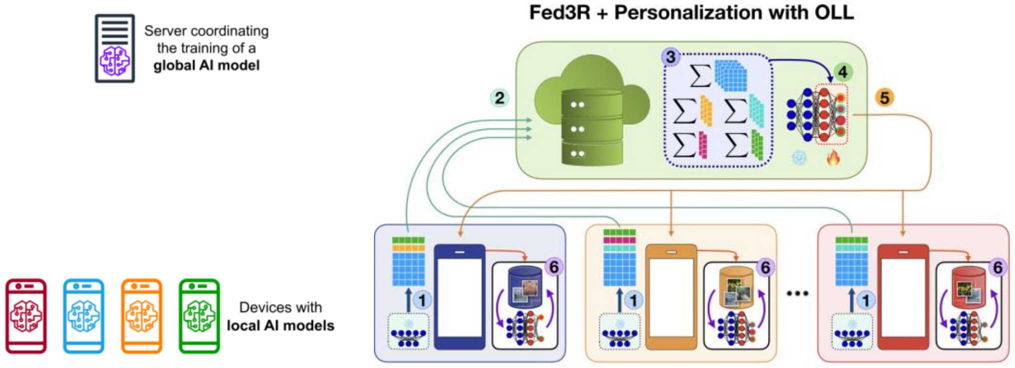

Pruning & Model Efficiency

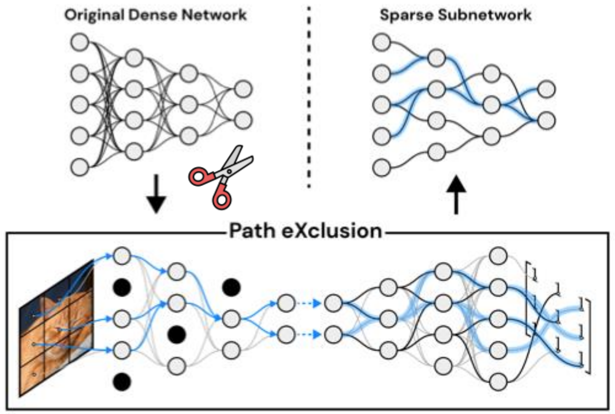

Model Safety & Unlearning

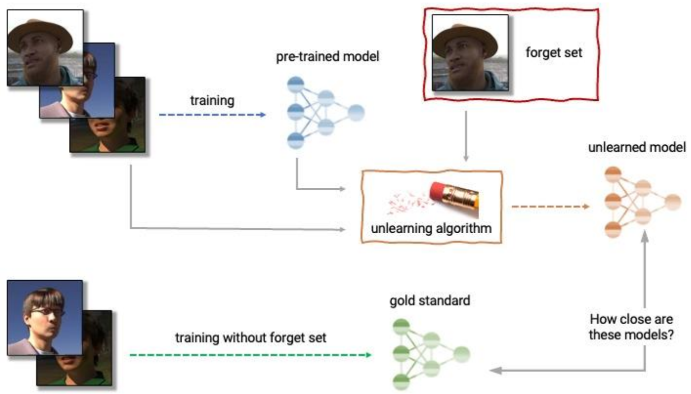

Theory Side of Machine Learning

Interaction-Aware Gaussian Weighting for Clustered Federated Learning, ICML 2025
Communication-Efficient Heterogeneous Federated Learning with Generalized Heavy-Ball Momentum, TMLR 2025
Beyond Local Sharpness: Communication-Efficient Global Sharpness-aware Minimization for Federated Learning, CVPR 2025
Improving Generalization in Federated Learning by Seeking Flat Minima, ECCV 2022
Accelerating Federated Learning via Sequential Training of Grouped Heterogeneous Clients, IEEE Access 2024
Finding Lottery Tickets in Vision Models via Data-driven Spectral Foresight Pruning, CVPR 2024
Efficient Model Editing with Task-Localized Sparse Fine-tuning, ICLR 2025

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

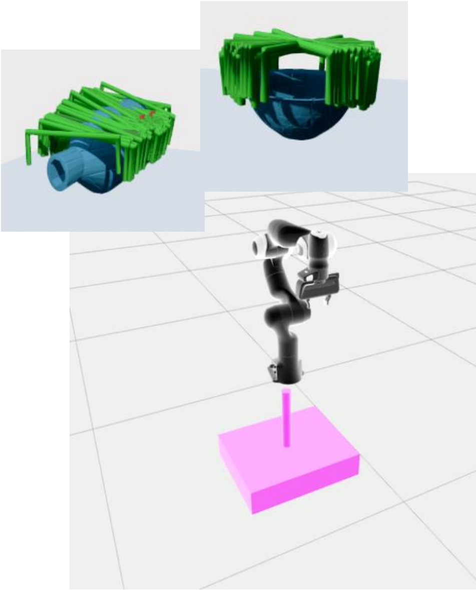

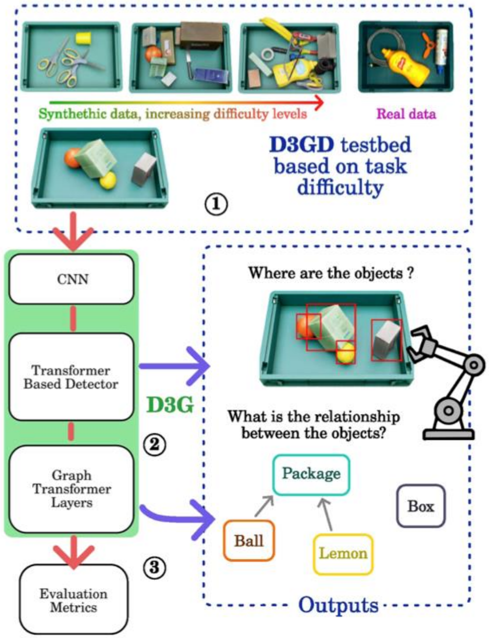

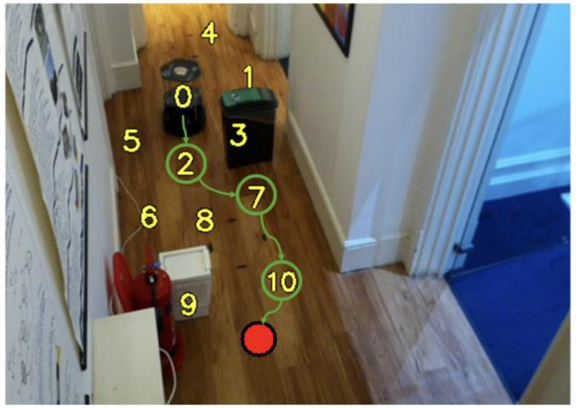

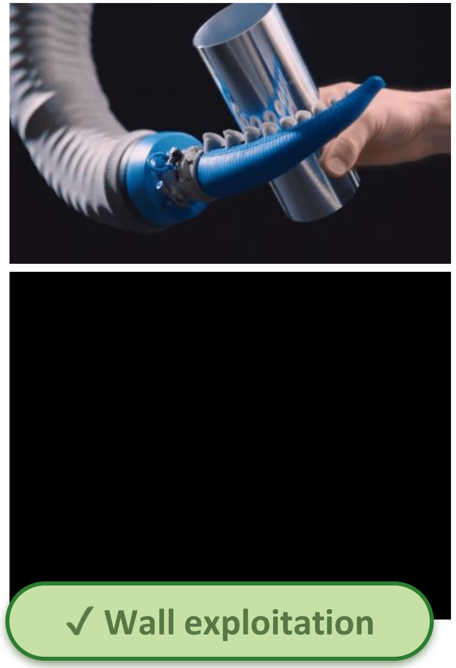

Robotics

A Modern Take on Visual Relationship Reasoning for Grasp Planning, IEEE RAL 2025
Select2Plan: Training-Free ICL-Based Planning through VQA and Memory Retrieval, ArXiv 2024
PaintNet: Unstructured Multi-Path Learning from 3D Point Clouds for Robotic Spray Painting, IEEE IROS 2023
Domain Randomization for Robust, Affordable and Effective Closed-Loop Control of Soft Robots, IEEE IROS 2023
End-to-end learning to grasp via sampling from object point clouds, IEEE RAL 2022

Mission

Develop **efficient and reliable algorithms** that allow artificial systems **to see, understand the surrounding world** and **learn autonomously.**

Focus: solve perception tasks based on visual (2D, 3D) and multimodal (vision, language) information.

Multi-Modal, Adaptive, Efficient and Robust Learning

- Spatial Intelligence
  - Semantic segmentation
  - Geo-localization
- Anomaly detection
- 3D learning
- Video Understanding
  - Ego and exocentric
- Robot Learning
  - Grasping
  - Painting
  - Locomotion
  - Soft-robotics
- Efficient and Distributed Learning
  - Federated Learning
  - Pruning and Distillation
- Un-learning (remove bias)

Structured Robot Learning

Key Enablers of Recent Machine Learning Successes

Large overparameterized models

Plenty of data

Growing computational capacity

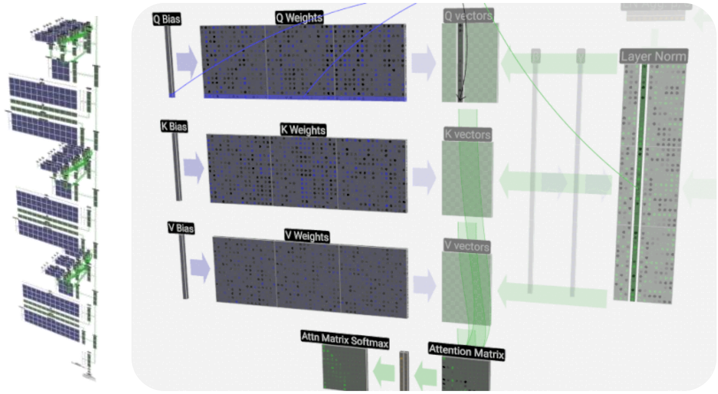

GPT-2

Autonomous Driving

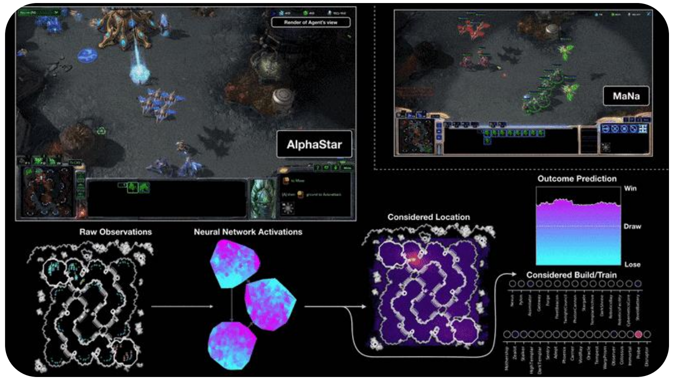

Deep Reinforcement Learning

Challenges for Resource-constrained Learning

Large overparameterized models
- Limited on-board memory
- Needs swift (re-)training to adapt

Plenty of data
- Data scarcity & domain shift
- Complex, structured data

Growing computational capacity
- Limited on-board computing

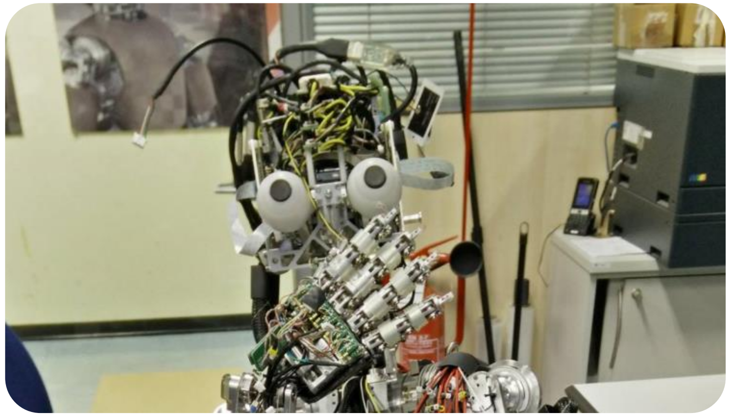

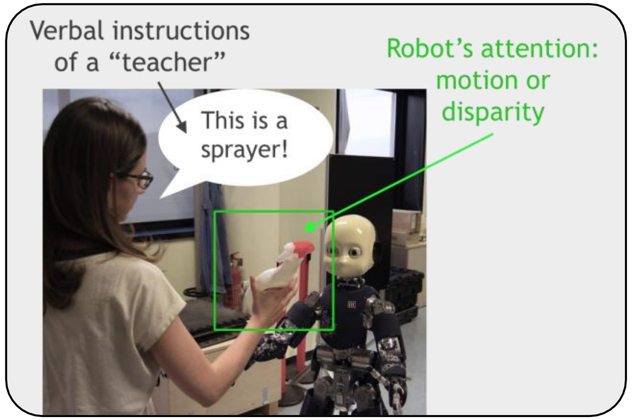

Challenges for Resource-constrained Learning

Large overparameterized models

- Limited on-board memory
- Needs swift (re-)training to adapt

arcity & domain shift

ex, structured data

Growing computa

on-board computing

Key Issues:

- Efficiency
- Structure

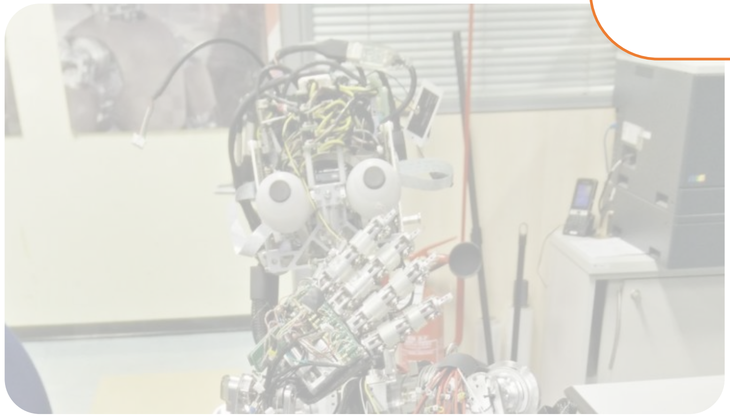

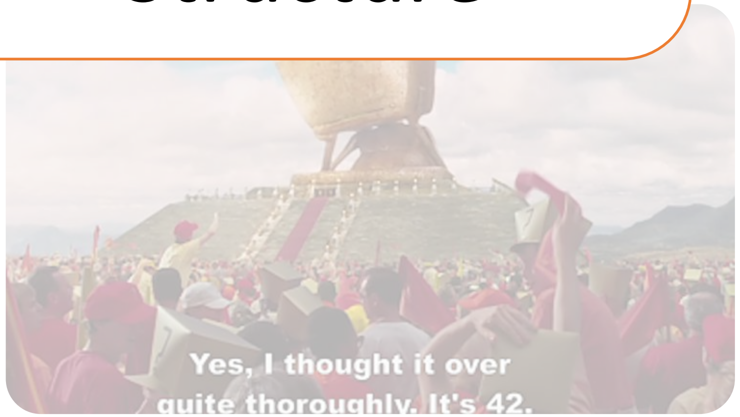

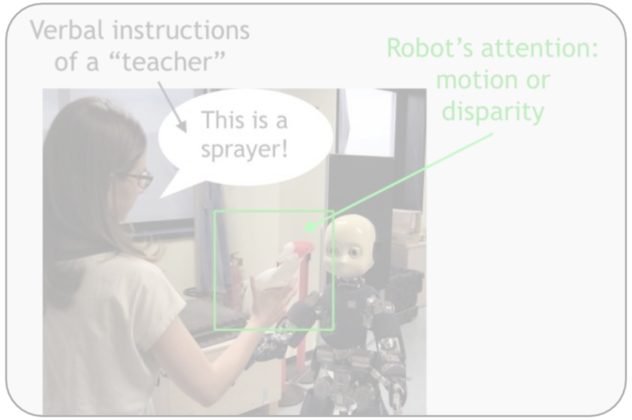

Robotics: A Playground for Efficient and Structured Learning

- **Goals**
  1. Learn robot models
  2. Learn actionable representations of the world
  3. Learn to control interaction to solve tasks
- **Requirements**
  - Efficiency
    - Limited on-board resources
    - Real-time prediction for seamless interaction
    - Data and labeling efficiency enable autonomy
  - Structure
    - Input/output spaces constrained by physics and geometry
    - Ensure safety
    - Learning on data streams in time, sequential structure
    - Environment dynamics evolves and is influenced by predictions/actions (e.g., RL)

Digit (credits: Agility robotics)

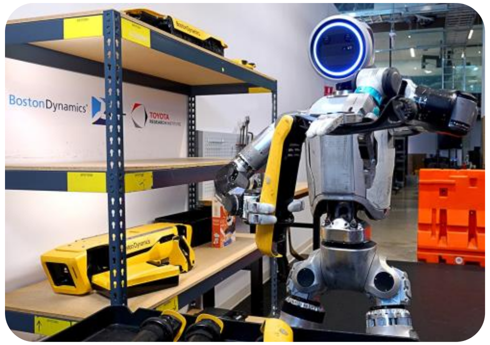

Atlas (credits: Boston Dynamics)

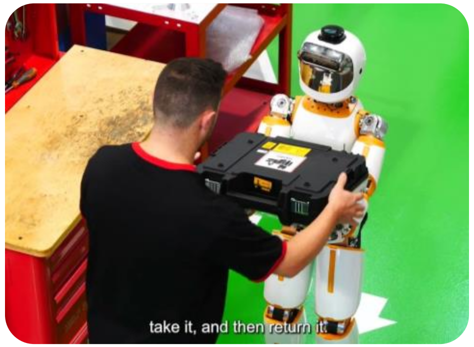

ErgoCub (credits: Istituto Italiano di Tecnologia & INAIL)

Robotics: A Playground for Efficient and Structured Learning

- **Goals**
  1. Learn robot models
  2. Learn actionable representations of the world
  3. Learn to control interactio
- **Requirements**
  - Efficiency
    - Limited on-board reso
    - Real-time prediction f
    - Data and labeling effic
  - Structure
    - Input/output spaces co
    - Ensure safety
    - Learning on data streams in time, sequential structure
    - Environment dynamics evolves and is influenced by predictions/actions (e.g., RL)

Applies to most Cyberphysical and Embodied Systems

Agility robotics)

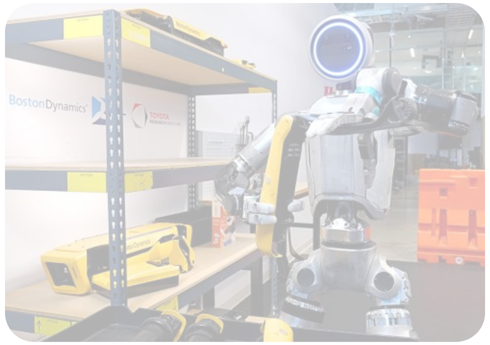

Atlas (credits: Boston Dynamics)

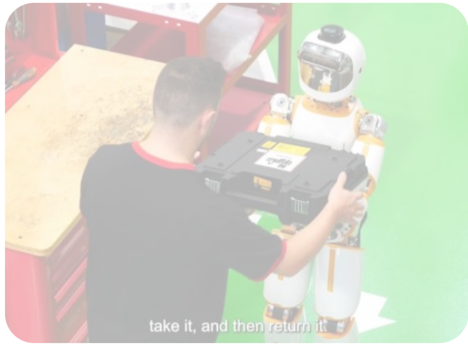

ErgoCub (credits: Istituto Italiano di Tecnologia & INAIL)

Supervised Learning

Dataset: $S_n = \{(x_i, y_i)\}_{i=1}^{\mathrm{n}}$

Problem: Given S_n find f_n s. t.

$$f_n(x_{new}) \approx y_{new}$$

f_n shall have low error on *future data*

Supervised Learning

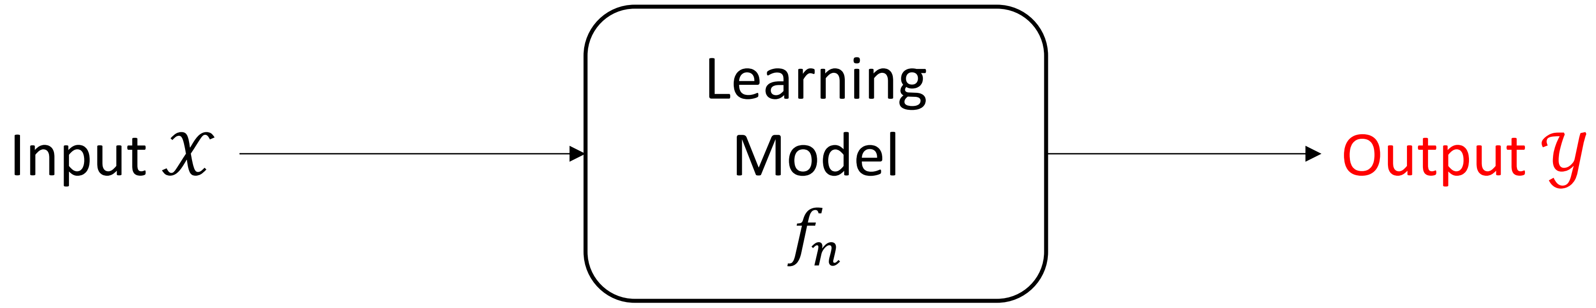

Output Spaces

Classification

$x_1 =$  $x_2 =$ 

$y_1 = 1$ $y_2 = 0$

Regression

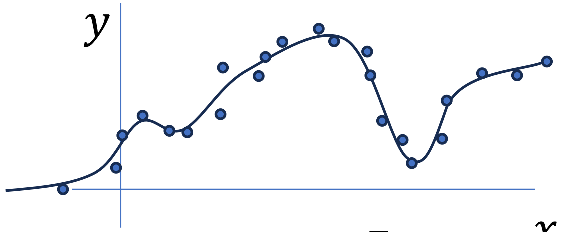

$\mathcal{Y} \in \mathbb{R}^T$

Structured Learning

$\mathcal{Y}$: set of structured outputs:

- Points on a manifold
- Histograms
- Graphs/trees
- Sequences
- …

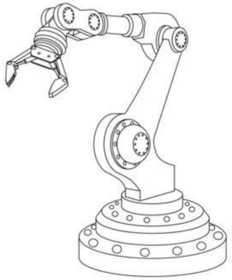

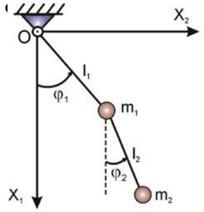

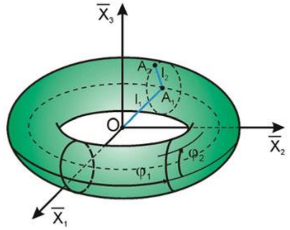

Ranking

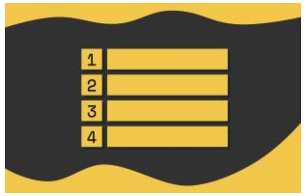

$\mathcal{Y} \in \mathcal{M} \subset \mathbb{R}^T$

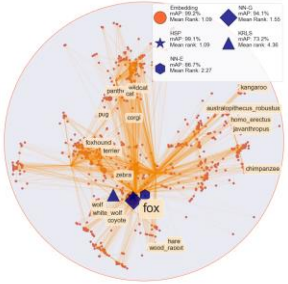

(a) WordNet mammals embedding

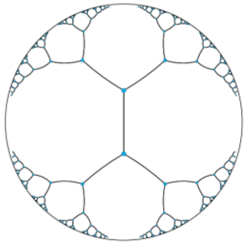

(b) Embedding of a tree in $\mathcal{B}^2$

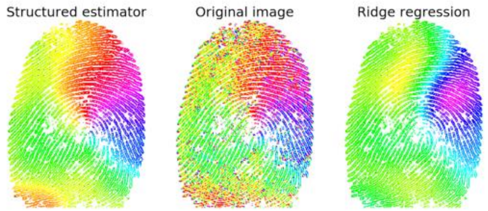

Nickel, Maximillian, and Douwe Kiela. "Poincaré embeddings for learning hierarchical representations." *Advances in neural information processing systems* 30 (2017).
Ciliberto, Carlo, Lorenzo Rosasco, and Alessandro Rudi. "A consistent regularization approach for structured prediction." NeurIPS 2016.
Marconi, Gian, Carlo Ciliberto, and Lorenzo Rosasco. "Hyperbolic manifold regression." AISTATS 2020.

Structured Robot Learning

- **Structured learning**:
  Useful approach for constrained learning in robotics
- Allows to encode desired properties of the output space or model
  - Physics priors and environment geometry [1,2]
  - Object geometry
  - Safety constraints [3]
  - Stability [4]

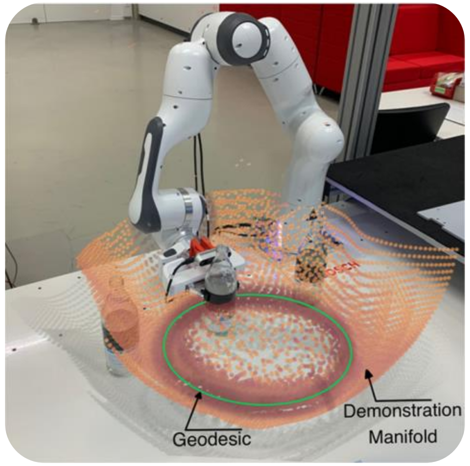

Learning the demonstration manifold [2]

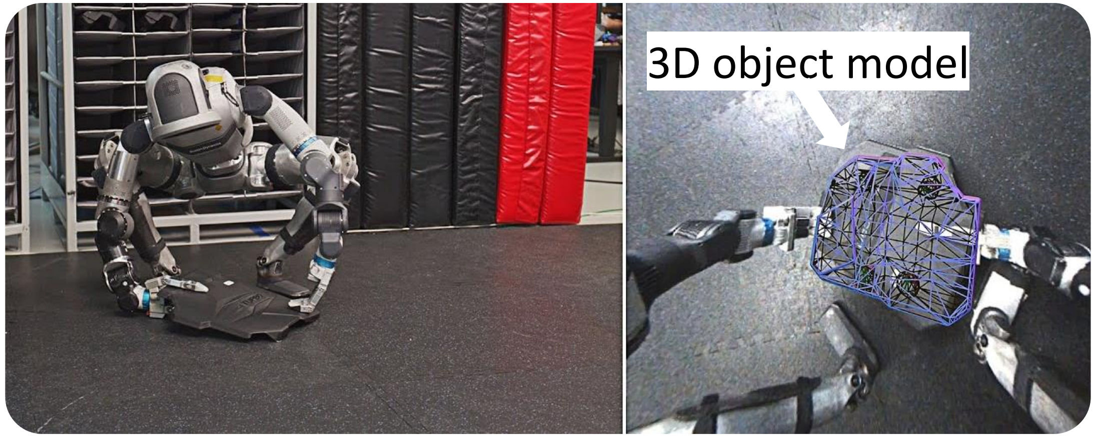

Bimanual object manipulation (credits: Boston Dynamics)

[1] Zeestraten, Martijn JA, et al. "An approach for imitation learning on Riemannian manifolds." *IEEE Robotics and Automation Letters* 2.3 (2017).
[2] Beik-Mohammadi, Hadi, et al. "Reactive motion generation on learned riemannian manifolds." *The International Journal of Robotics Research* (2023).
[3] Liu, P., Tateo, D., Ammar, H. B., & Peters, J. Robot reinforcement learning on the constraint manifold. In *Conference on Robot Learning* (2022, January).
[4] Mazzoleni, Mirko, et al., "A comparison of manifold regularization approaches for kernel-based system identification." *IFAC-PapersOnLine* (2019).

Structured Robot Learning

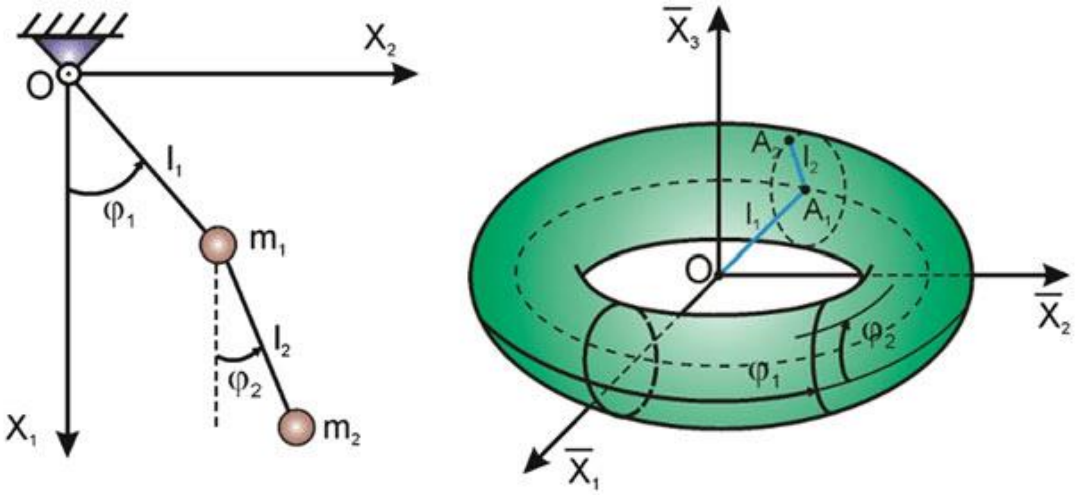

Robot Structure

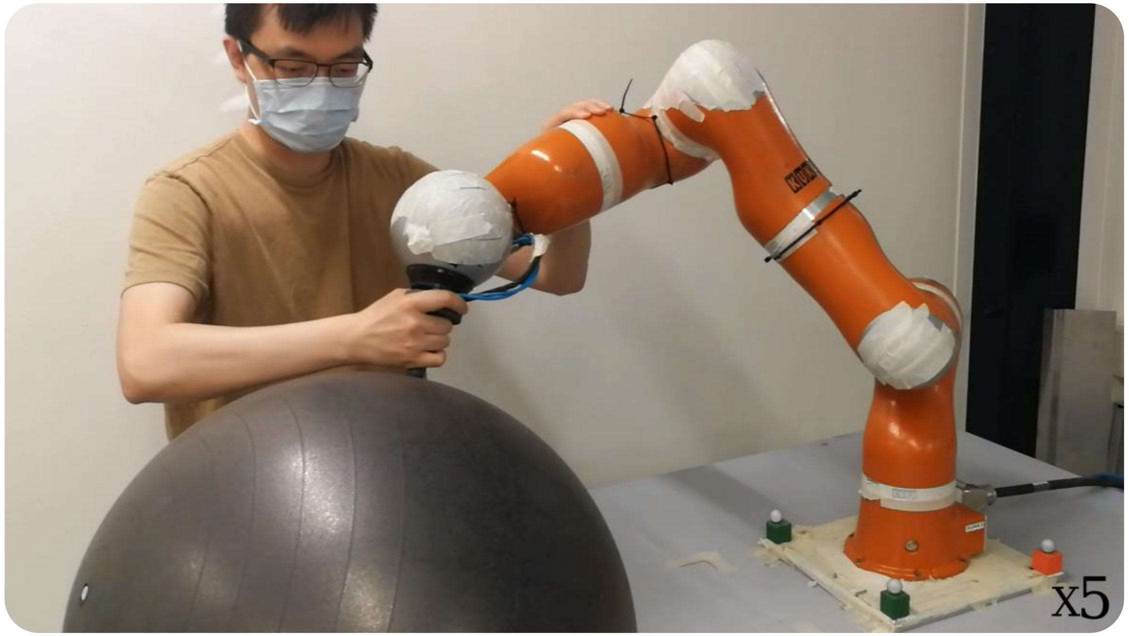

Object Structure

In both cases, leveraging structure improves efficiency and safety

Structured Inverse Kinematics Learning

- **Goal**: Learn robot inverse kinematics (e.g., for CLIKC)
  - Challenge: Model-based IK fails under <u>model misspecification</u>
- Learn on the robot manifold
- Can we learn accurate IK mapping from data?
- Yes, via structured prediction
  - Allows to use custom loss in Cartesian space
  - Flexible nonparametric kernel method

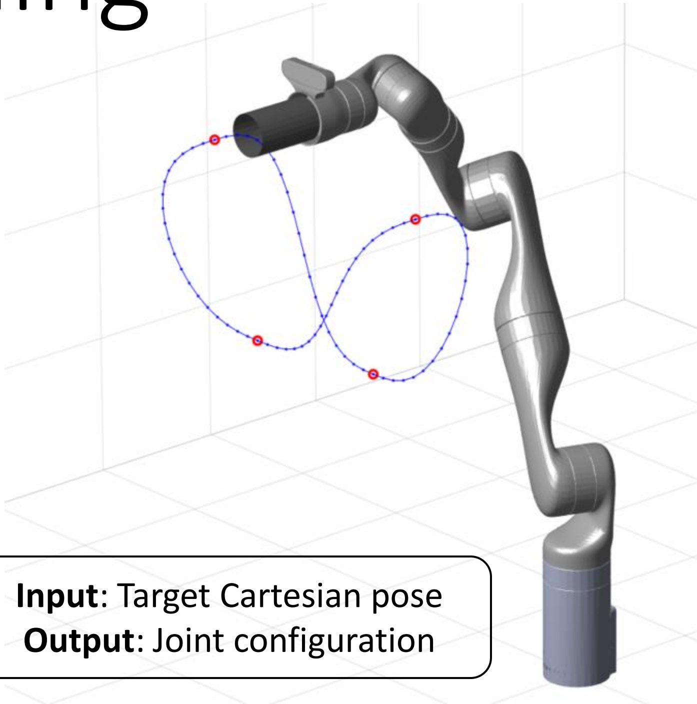

Input: Target Cartesian pose
Output: Joint configuration

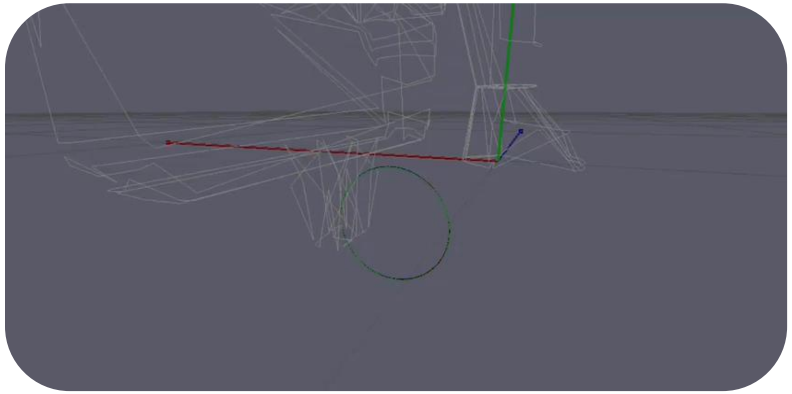

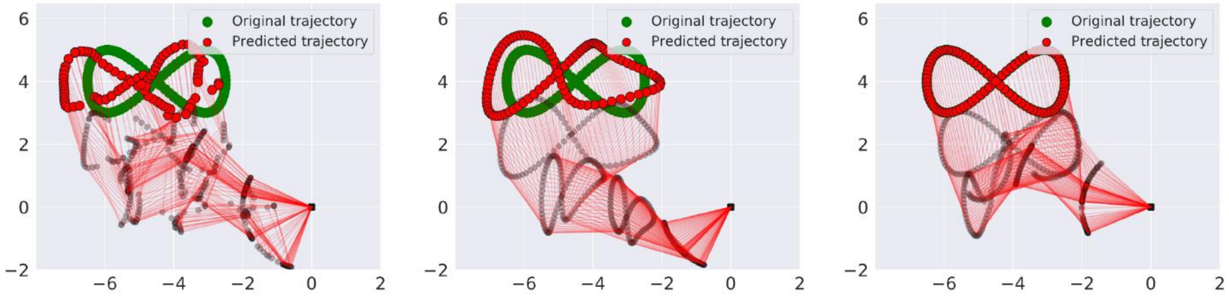

(a) One Class Structured SVM (b) CRiSP-R (c) CRiSP-FK

Gian Maria Marconi, R. C. , Lorenzo Rosasco, Carlo Ciliberto, "Structured Prediction for CRiSP Inverse Kinematics Learning with Misspecified Robot Models", IEEE RAL 2021

Inverse Kinematics

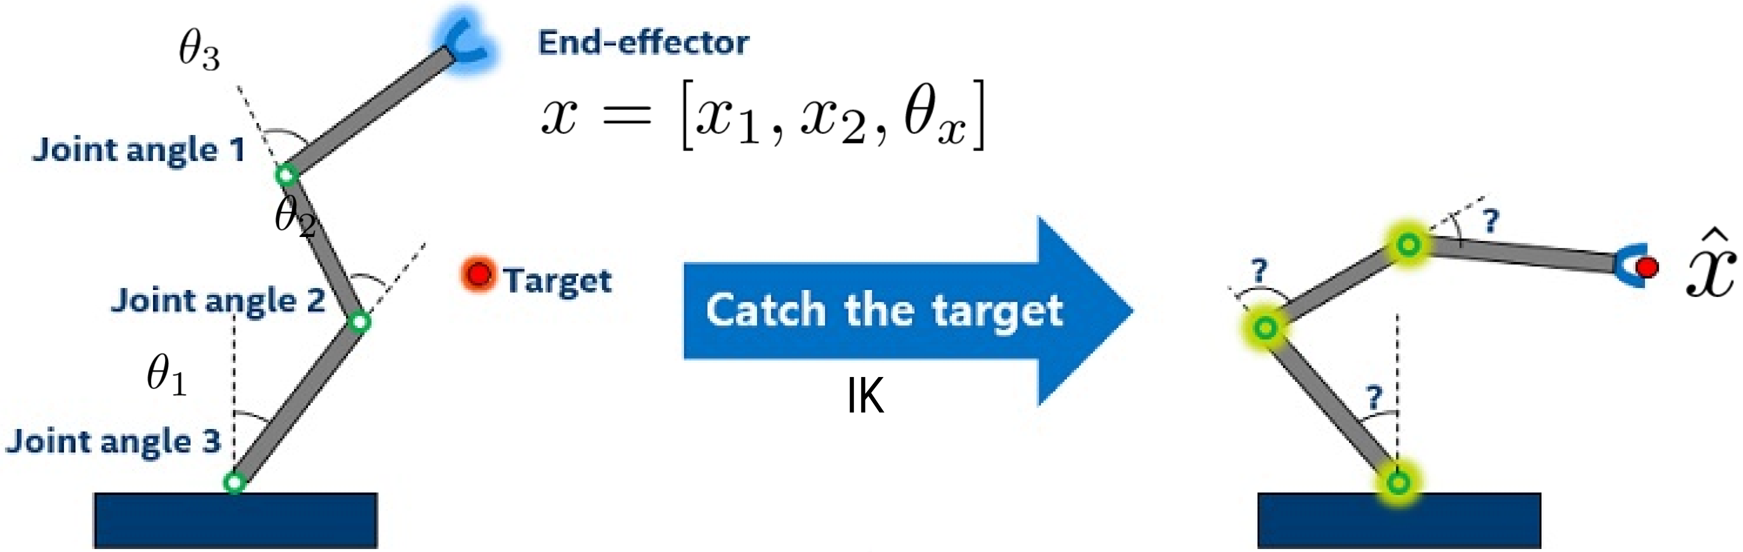

Configuration space ***Task space***

$$f : \mathbb{R}^d \times SO(d) \to [l_1, u_1] \times \ldots \times [l_J, u_J]$$

Goal: $\overset{\text{FK}}{g} \circ \overset{\text{IK}}{f}(x) \simeq x$

Nguyen-Tuong, Duy, and Jan Peters. "Model learning for robot control: a survey." *Cognitive processing* 12.4 (2011): 319-340.

Why Data Driven?

Analytical approaches:

- Need to be derived from scratch for every model
- Might be too complicated or unfeasible for some robot structures
- Can be very sensitive to errors in the mechanical model (e.g. non-rigid links, human body)

CRiSP: Kernel-based Structured IK Learning

Structured kernel method for learning inverse kinematics from data

- Output space
  - Safety constraints (joint limits)
  - Complex nonlinear structure

- Robust to model misspecification

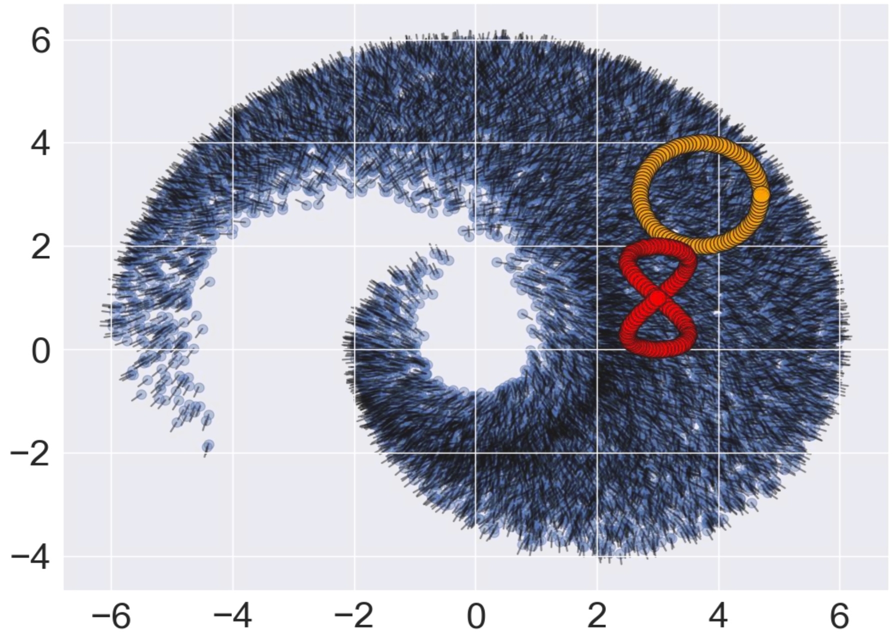

Example of CRiSP dataset

Marconi G. M.*, **C. R.*,** Rosasco L., Ciliberto C. Structured Prediction for CRiSP Inverse Kinematics Learning with Misspecified Robot Models, *IEEE Robotics and Automation Letters (RA-L) & IEEE ICRA 2021*

Challenge 1: Structured Output Space

Task space

Configuration space

$$f : \mathbb{R}^d \times SO(d) \to [l_1, u_1] \times \ldots \times [l_J, u_J]$$

$$[l_j, u_j] \subset [0, 2\pi)$$

Joint limits

Challenge 2: Loss in Joint Space

$$f(x) = \langle w, x \rangle$$

$$w = (X^\top X + \lambda I)^{-1} X^\top Y$$

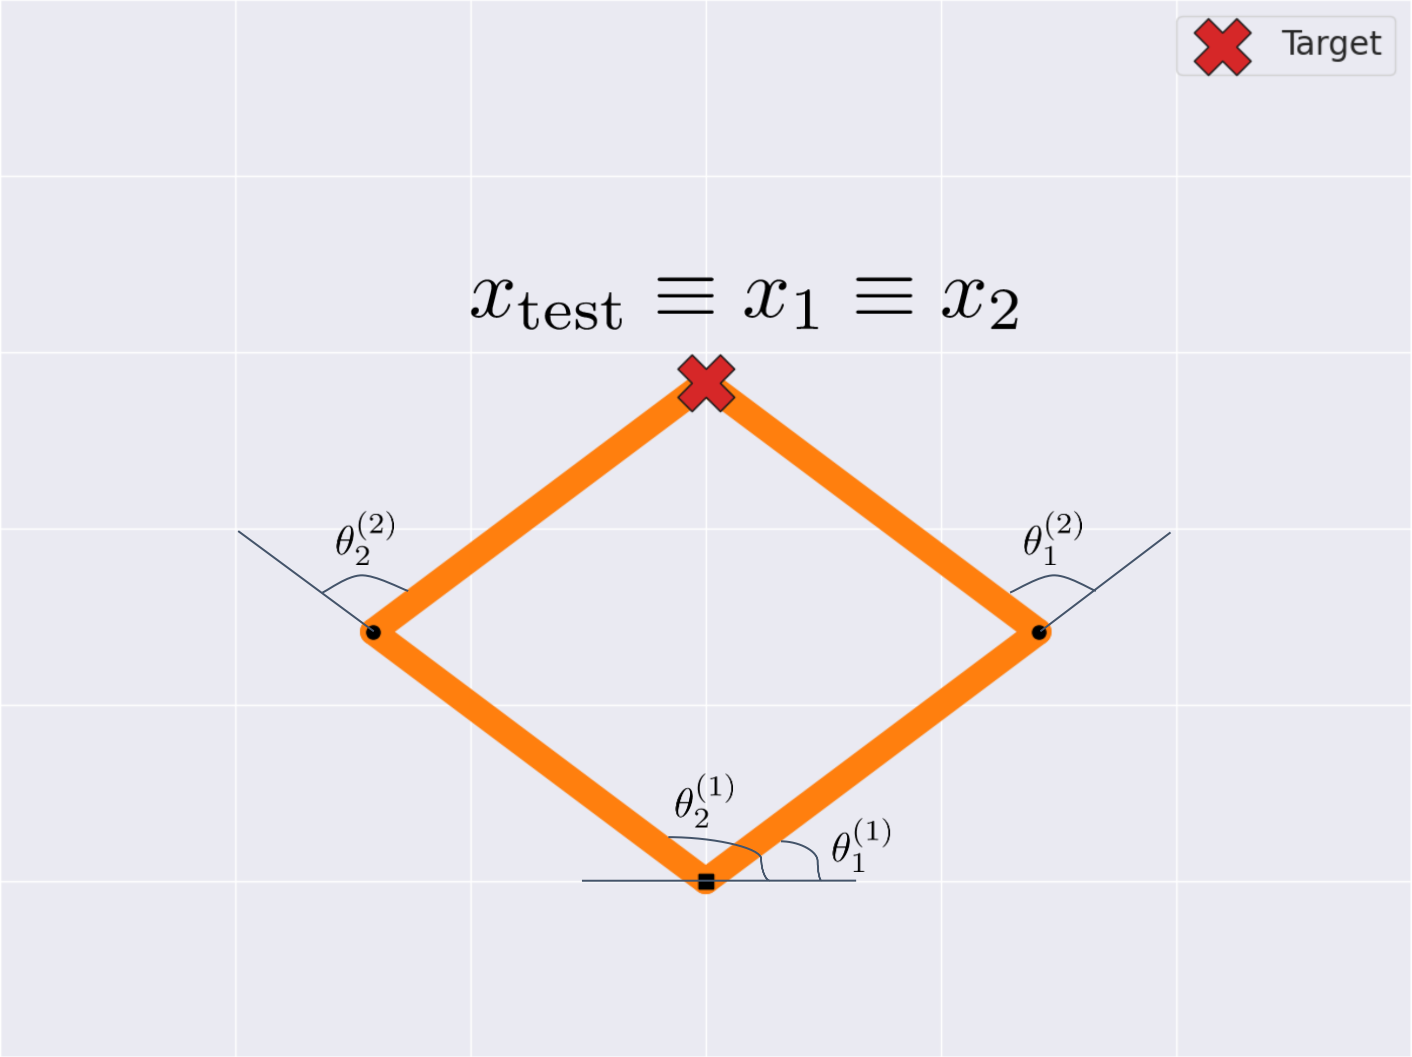

$$y_1 = [\theta_1^{(1)}, \theta_1^{(2)}]$$

$$y_2 = [\theta_2^{(1)}, \theta_2^{(2)}]$$

$$X = \begin{bmatrix} -x_1- \\ -x_2- \end{bmatrix}$$

Challenge 2: Loss in Joint Space

$$f(x_{\text{test}}) = [\hat{\theta}^{(1)}, \hat{\theta}^{(2)}]$$

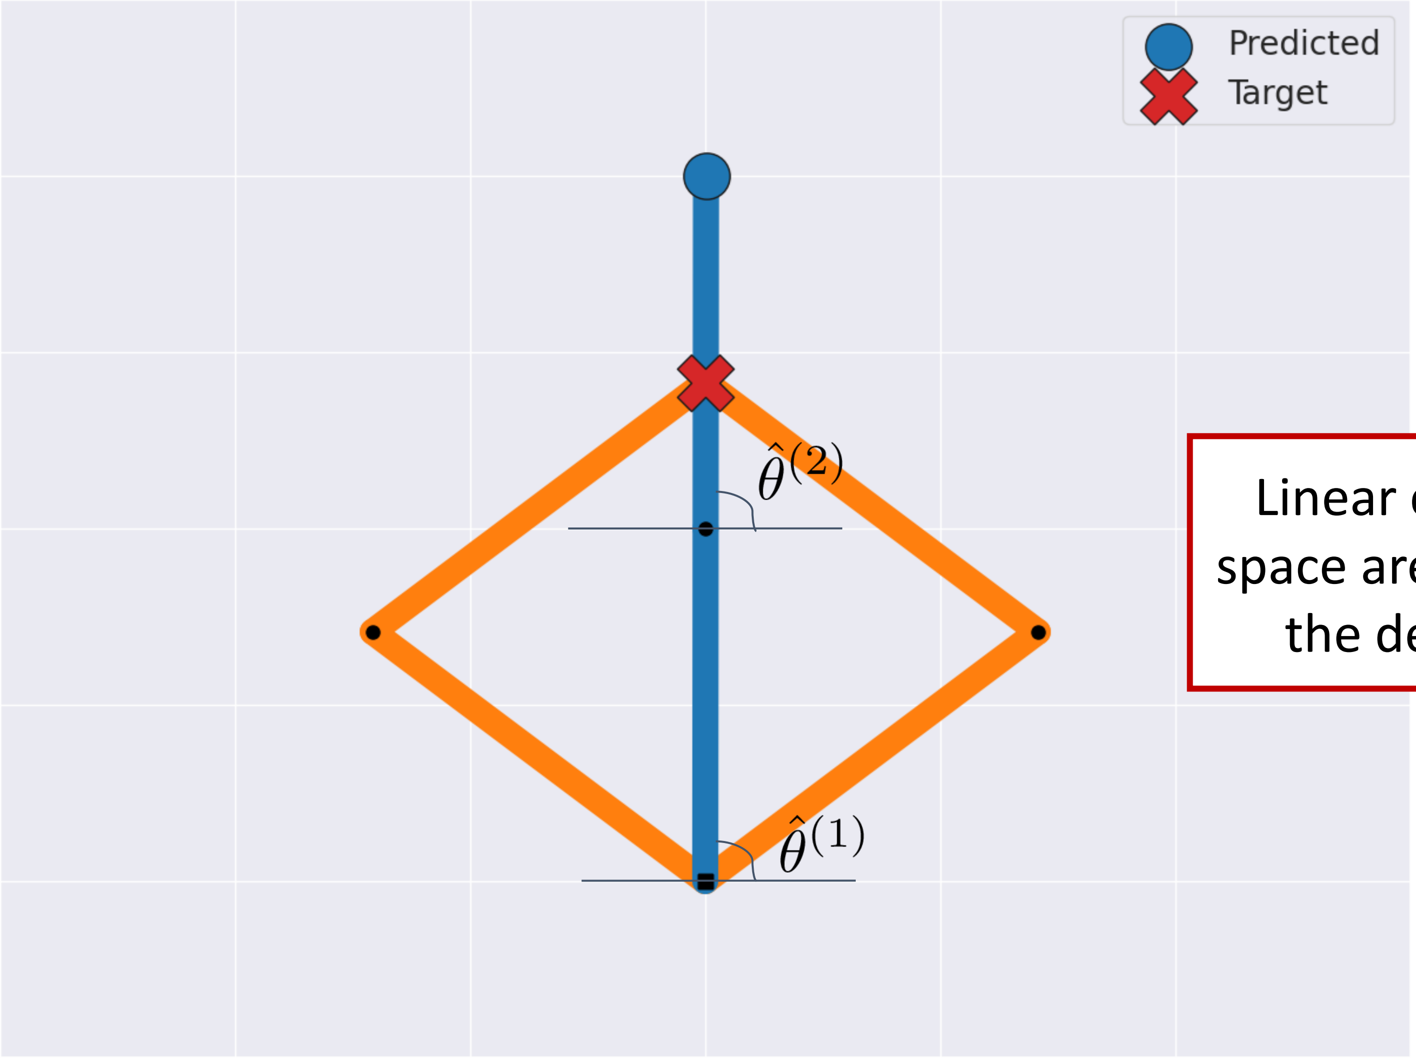

Linear combinations in joint space are not representative of the desired configuration

CRiSP-FK with Forward Kinematics Loss

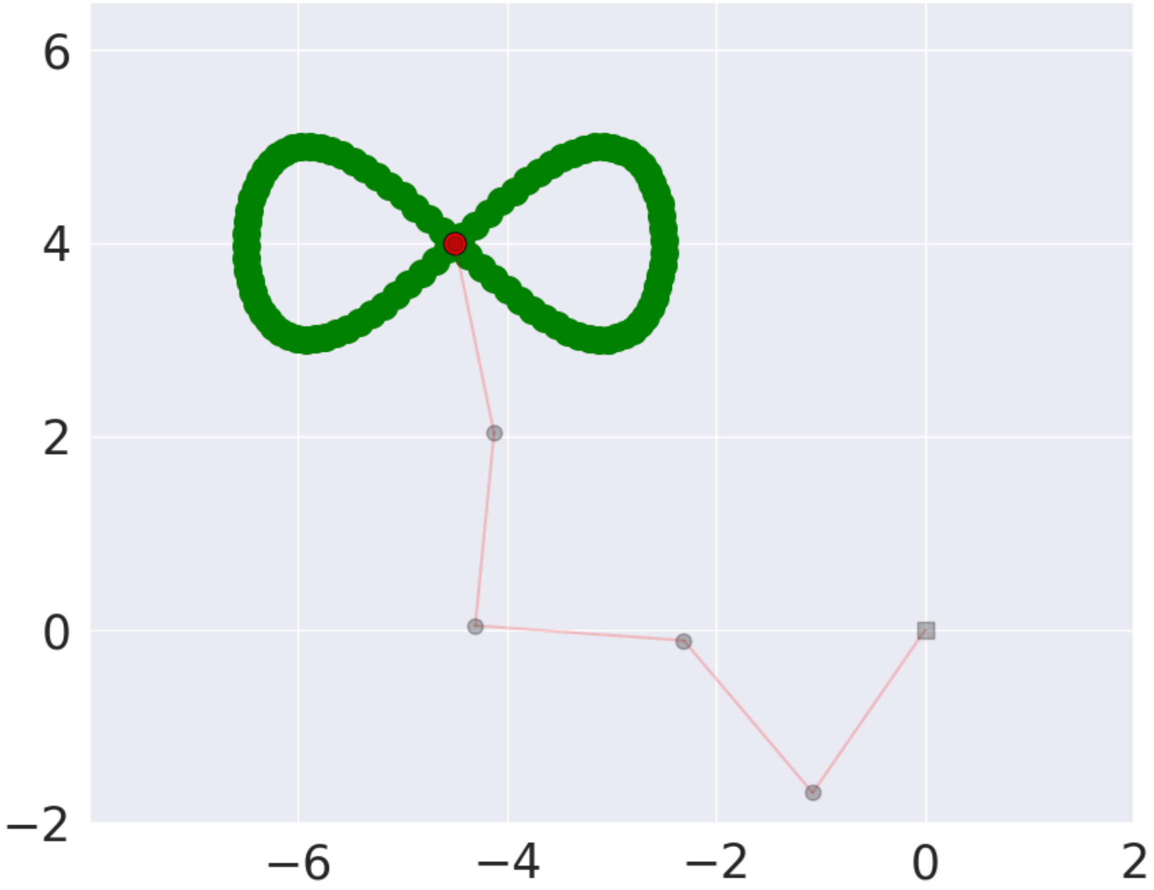

1) Training: Compute $\alpha_i : \mathbb{R}^d \to \mathbb{R}$

$$\alpha(x) = (K + n\lambda I)^{-1} K_x$$

2) Prediction: Find minimizing y
e.g., L-BFGS constrained minimization

$$f(x) = \arg\min_{y \in \mathcal{Y}} \alpha_i(x) \Delta\big(\tilde{g}(y), \tilde{g}(y_i)\big)$$

$$\triangle(y, y_i) = \underbrace{\|\tilde{g}_p(y) - \tilde{g}_p(y_i)\|^2}_{\text{Position Loss}} + \underbrace{d_O(\tilde{g}_o(y), \tilde{g}_o(y_i))^2}_{\text{Orientation Loss}}$$

Baseline Comparisons

CRiSP w/ Joint Loss Minimization

CRiSP-FK
(Proposed)

One-Class Structured SVM

Neural Network

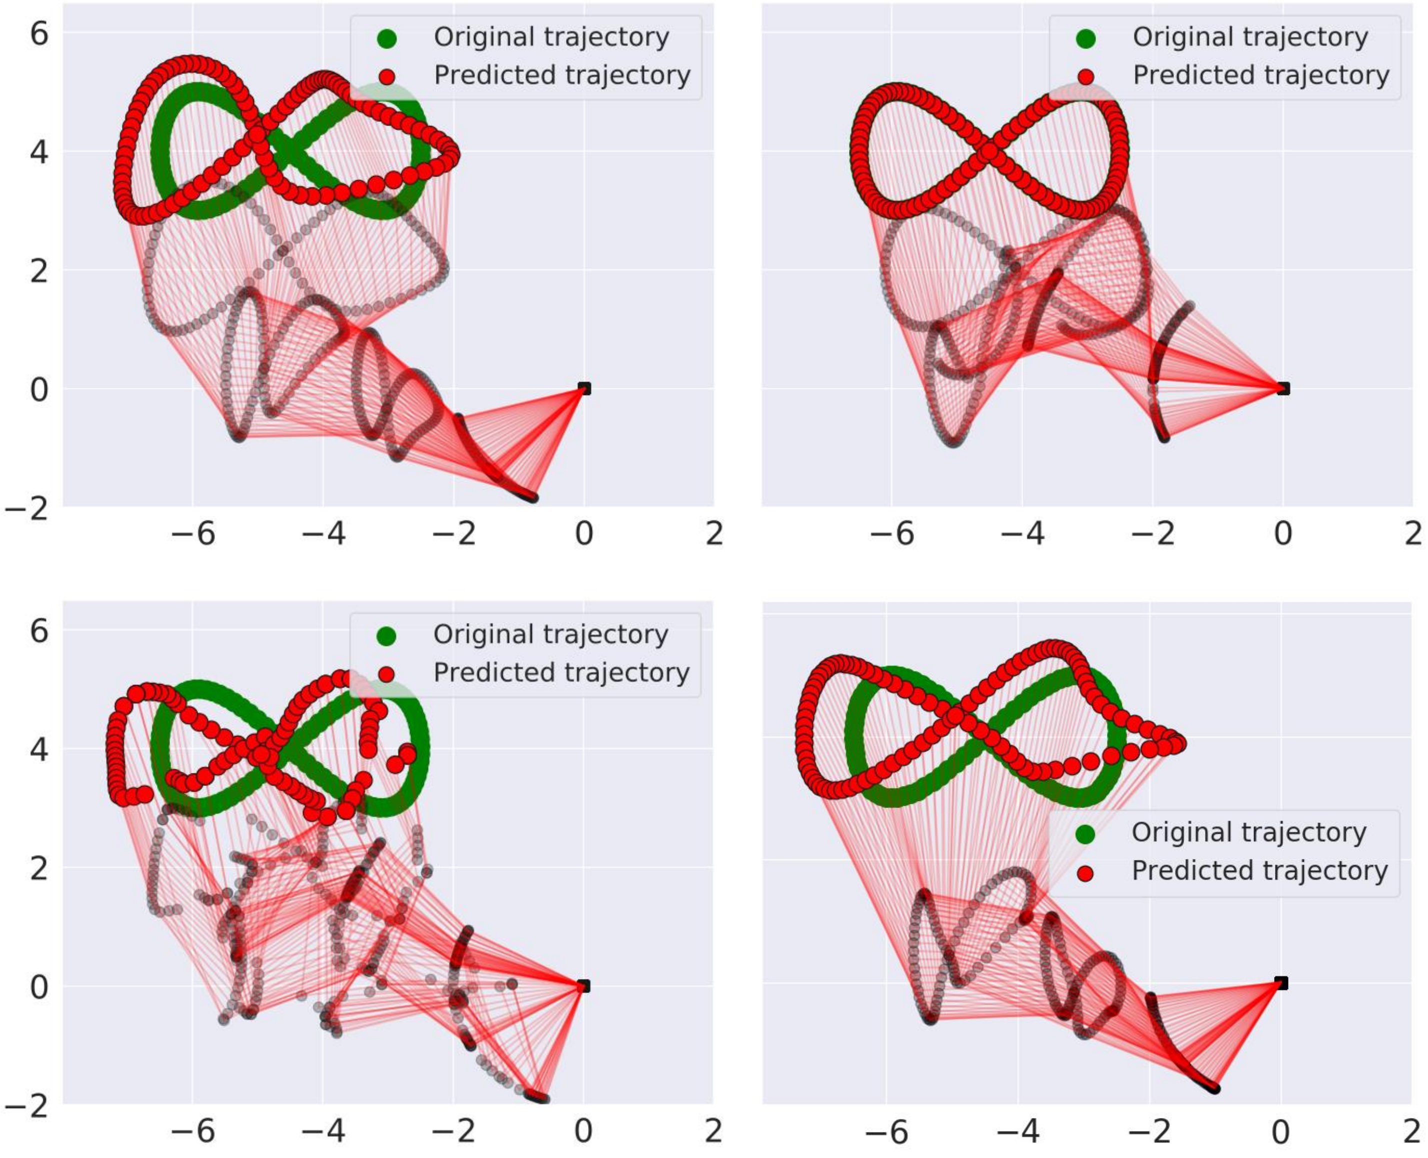

CRiSP-based Trajectory Reconstruction on a 7-DoF Panda Manipulator

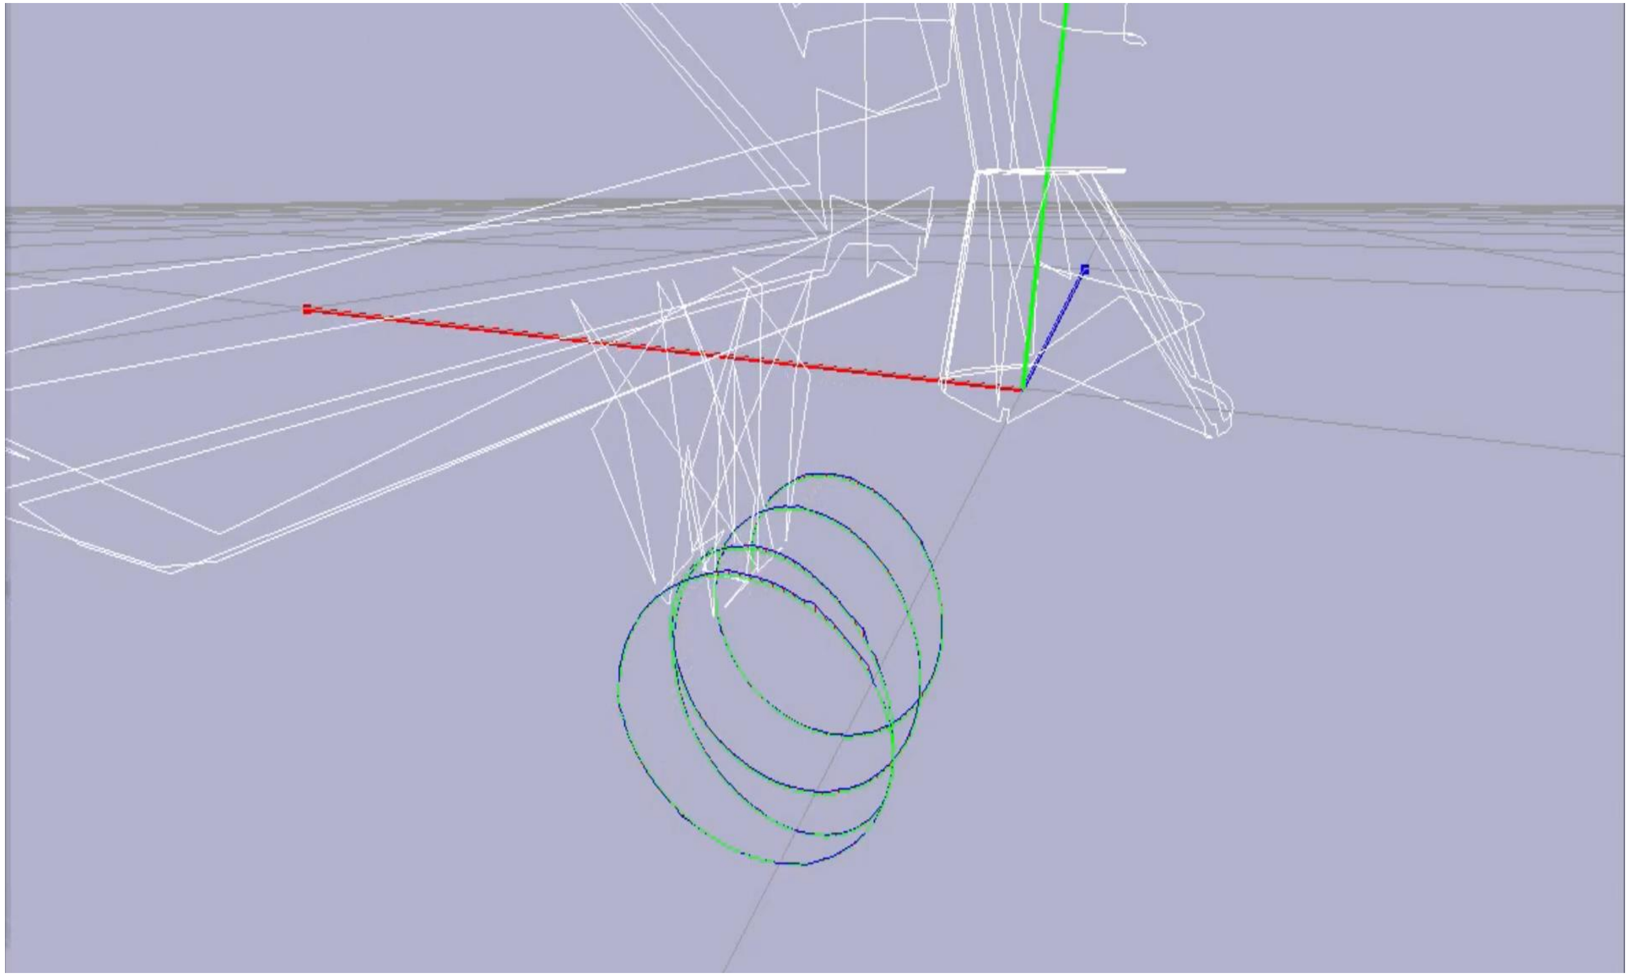

Structured Imitation Learning

- **Goal:** Safely learn trajectories on non-Euclidean surfaces
- Low-dimensional manifold: Fewer examples to learn policy
- Applications: Surface polishing, robotic surgery, …
- **We propose a kernelized structured prediction method for supervised imitation learning (SIL) on manifolds**

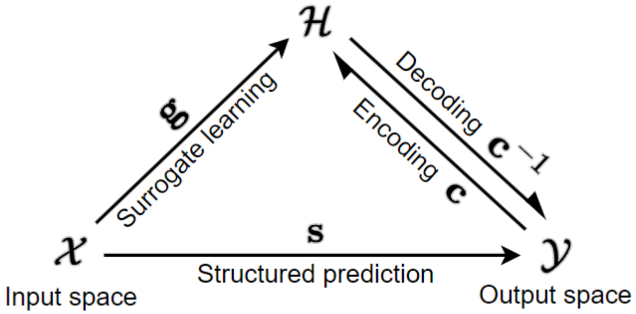

Figure 2. Schematic illustration of the surrogate approach to structured prediction.

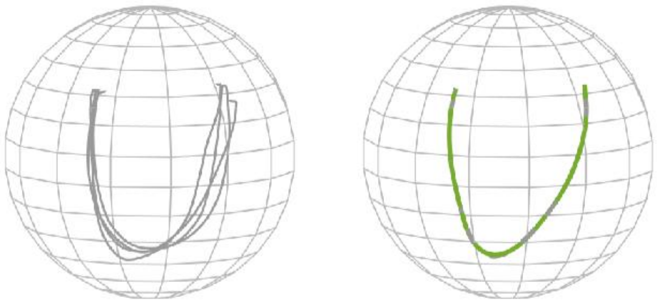

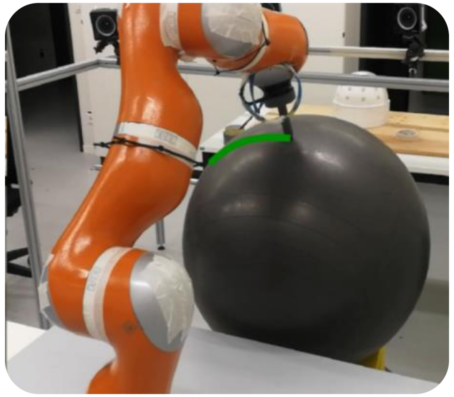

Surrogate Structure-encoding Loss

$$D_{\mathrm{KL}}(\tilde{\mathbf{y}}_n, \tilde{\mathbf{y}}) = \frac{1}{2}\int\left(\log\frac{|\mathbf{\Sigma}|}{|\mathbf{\Sigma}_n|} - \mathrm{Log}_{\boldsymbol{\mu}_n}(\mathbf{y})\mathbf{\Sigma}_n^{-1}\mathrm{Log}_{\boldsymbol{\mu}_n}(\mathbf{y}) + \mathrm{Log}_{\boldsymbol{\mu}}(\mathbf{y})\mathbf{\Sigma}^{-1}\mathrm{Log}_{\boldsymbol{\mu}}(\mathbf{y})\right)\tilde{\mathbf{y}}_n\, d\mathbf{y}.$$

A Duan, I Batzianoulis, R Camoriano, L Rosasco, D Pucci, A Billard "A Structured Prediction Approach for Robot Imitation Learning", IJRR 2023

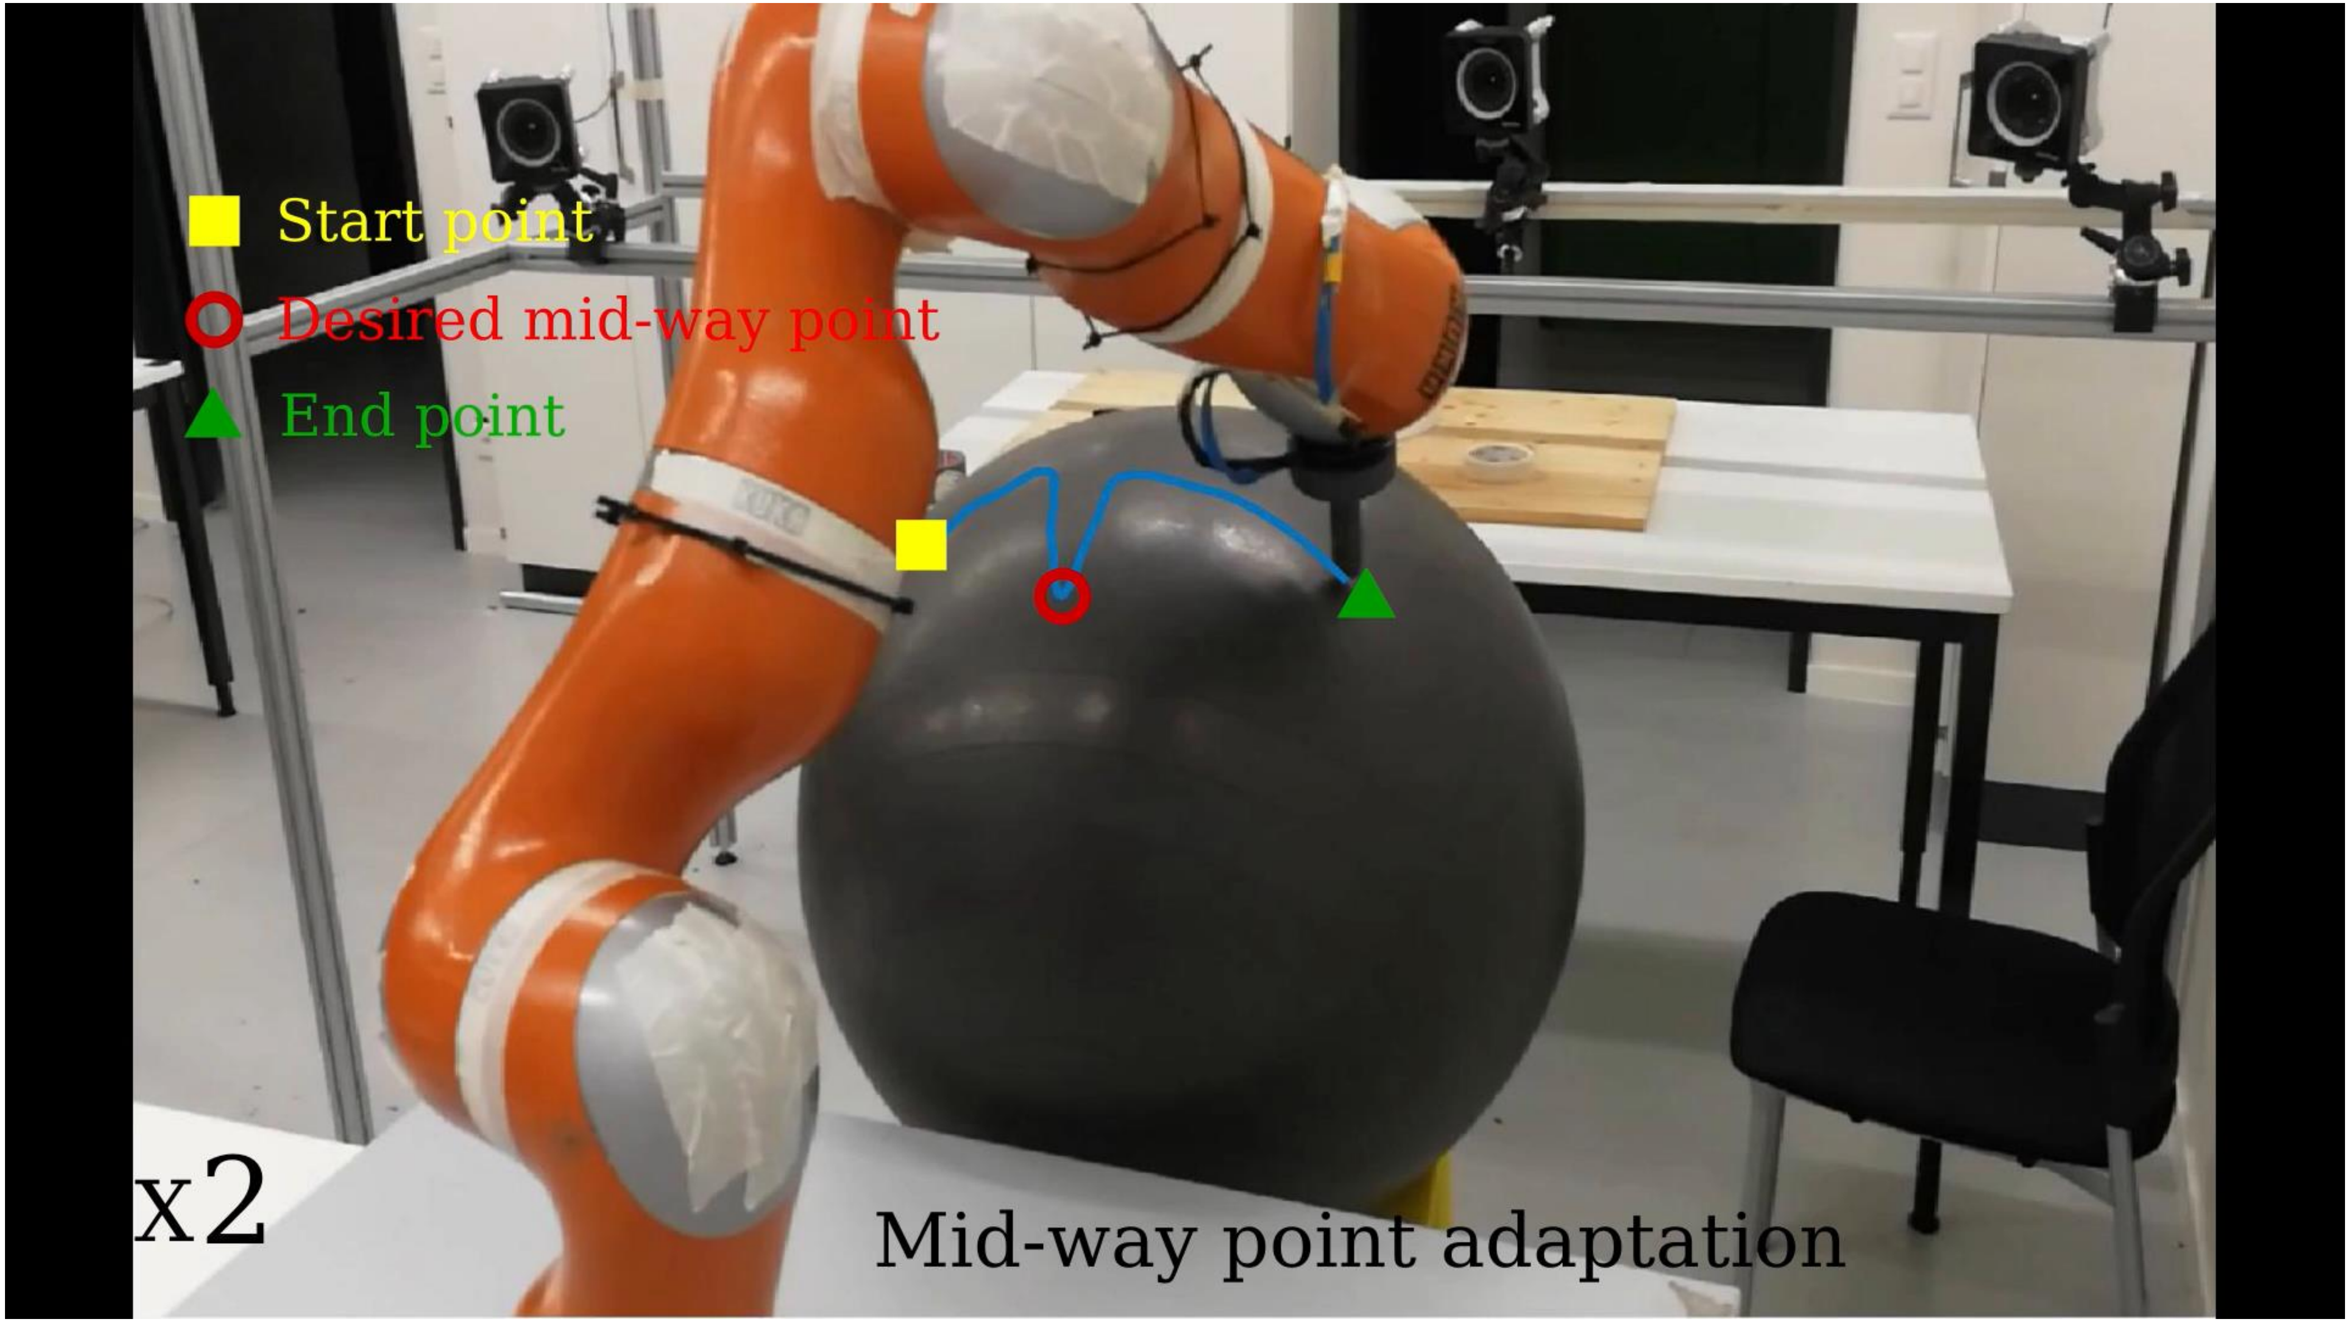
Start point
Desired mid-way point
End point
x2
Mid-way point adaptation

Extension: Learning (Quasi)periodic Trajectories on Manifolds

- Quasi-periodic trajectories are central to many applications
  - Healthcare (laser treatment)
  - Manufacturing
  - Coverage path planning
  - …
- Extension of SIL to QP trajectories
  - Structured prediction
  - **Periodic kernels**
  - Spherelet approximation
  - Adaptation to new target meshes
- Real-world experimental validation

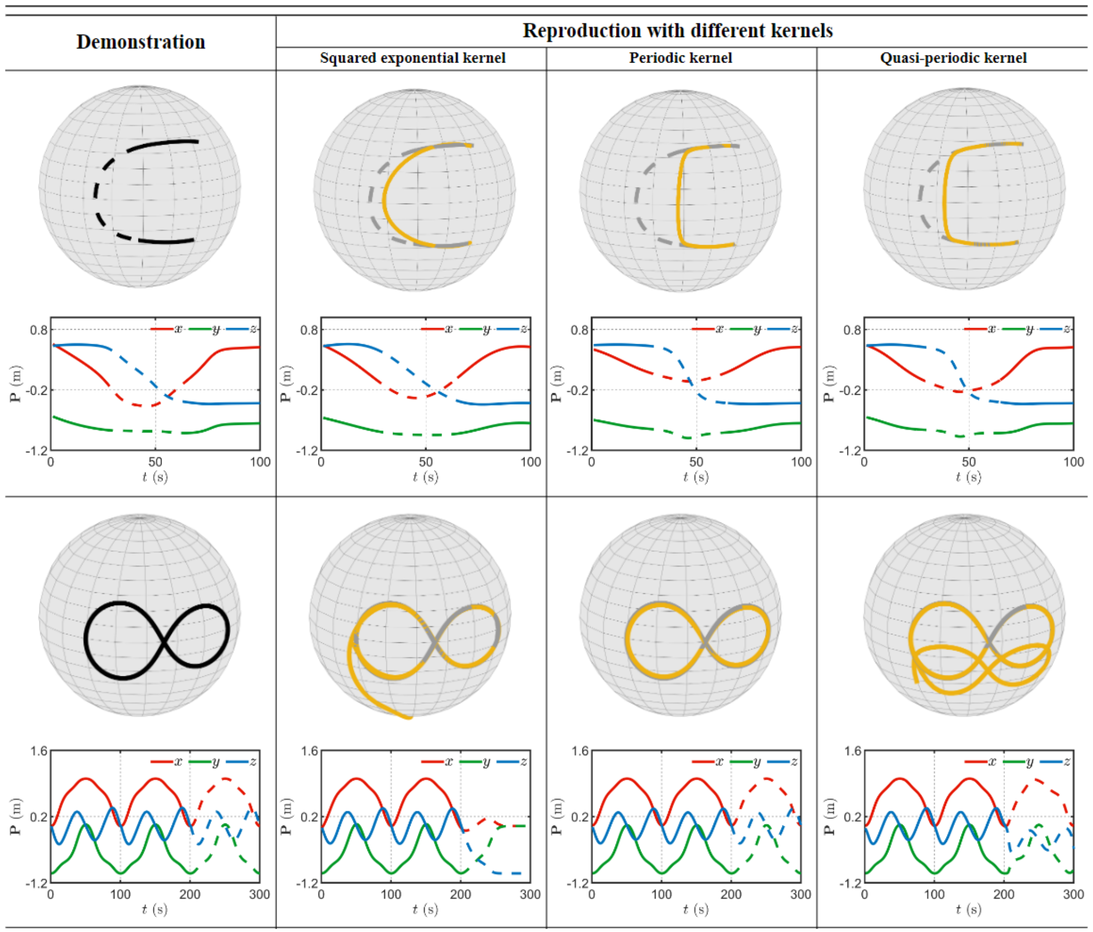

Duan, A. et al. "Learning Rhythmic Trajectories with GeometricConstraints for Laser-Based Skincare Procedures" IEEE T-RO 2025

Learning (Quasi)periodic Trajectories on Manifolds

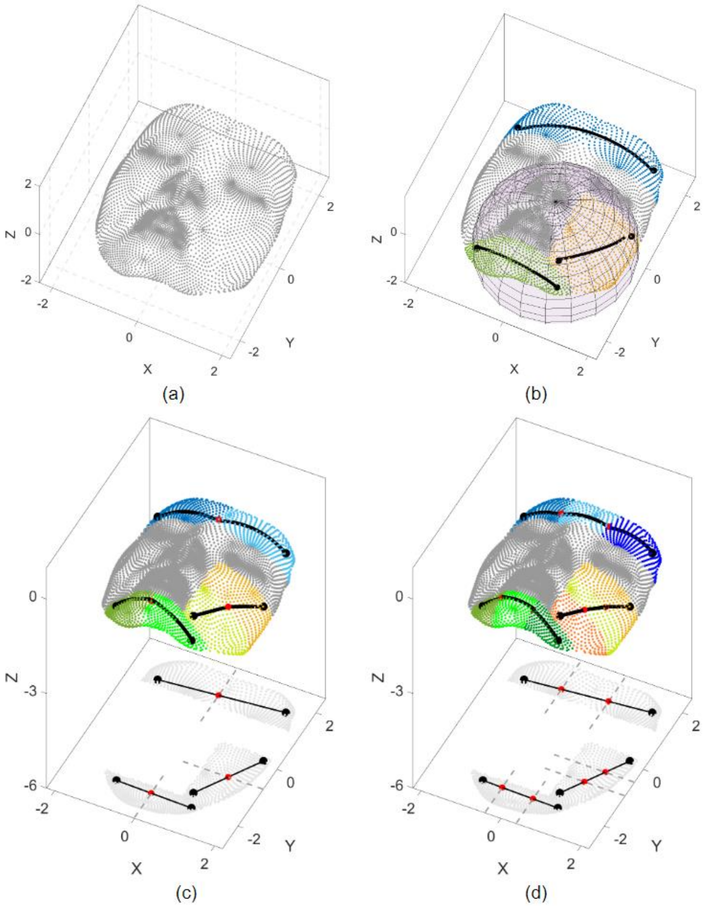

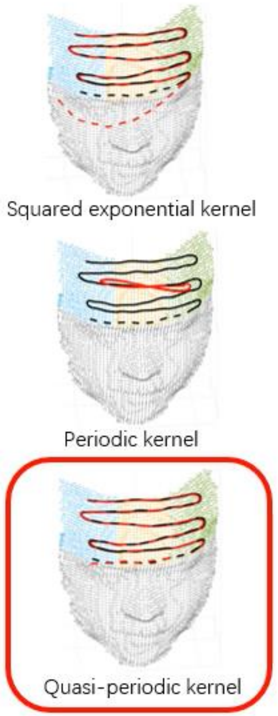

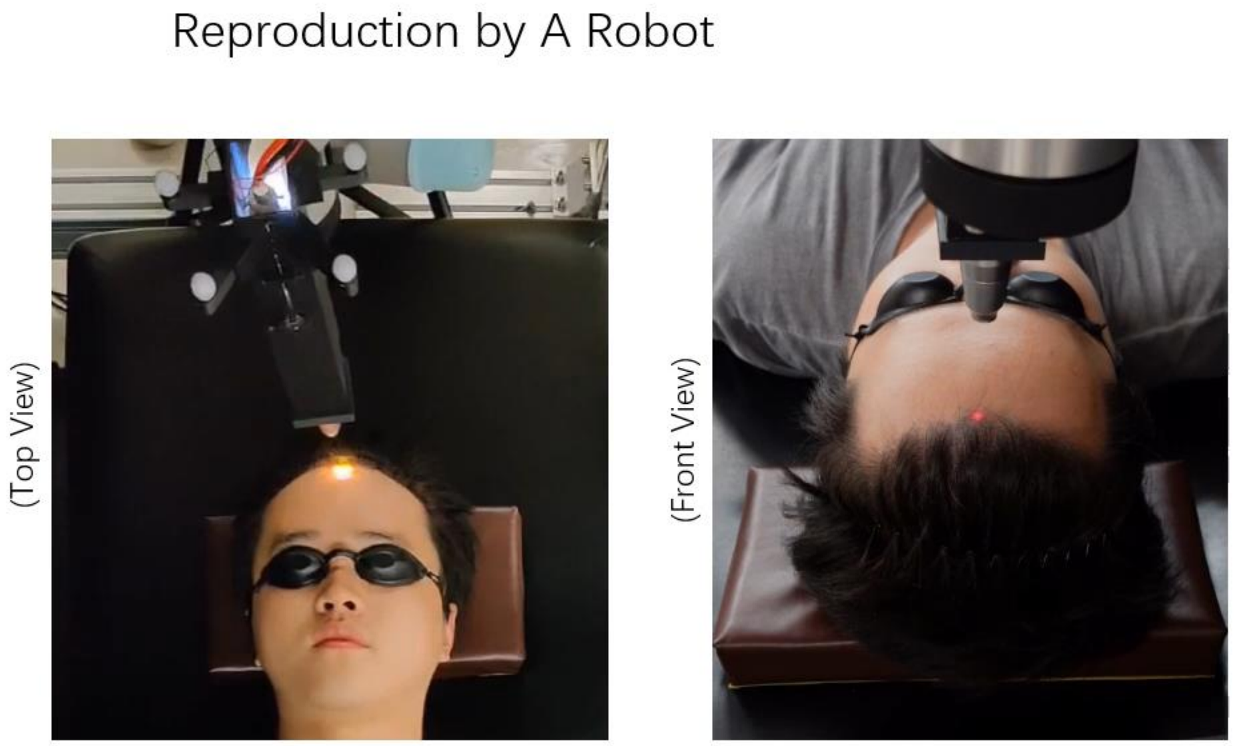

Duan, A. et al. "Learning Rhythmic Trajectories with GeometricConstraints for Laser-Based Skincare Procedures" IEEE T-RO 2025

Adaptation to new faces

Goal: equip the robot with the capacity to execute treatment on new subjects

Probabilistic nonrigid registration technique

- Human facial shapes are captured by key facial features such as control points (e.g., chin, eyebrows).
- Probabilistic trajectory adaptation is performed based on the extracted rules

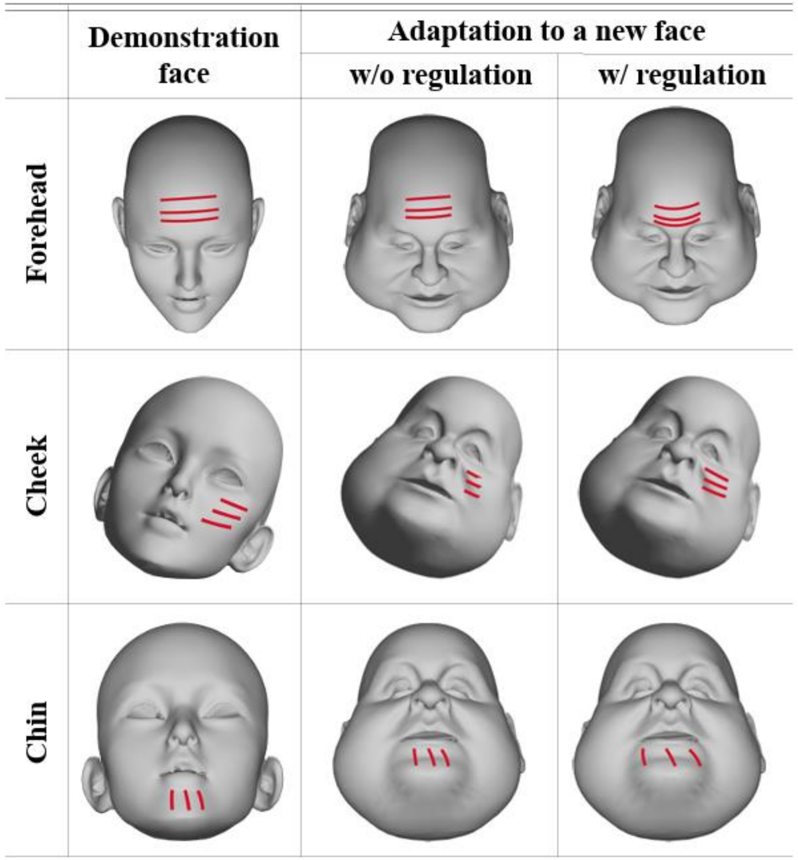

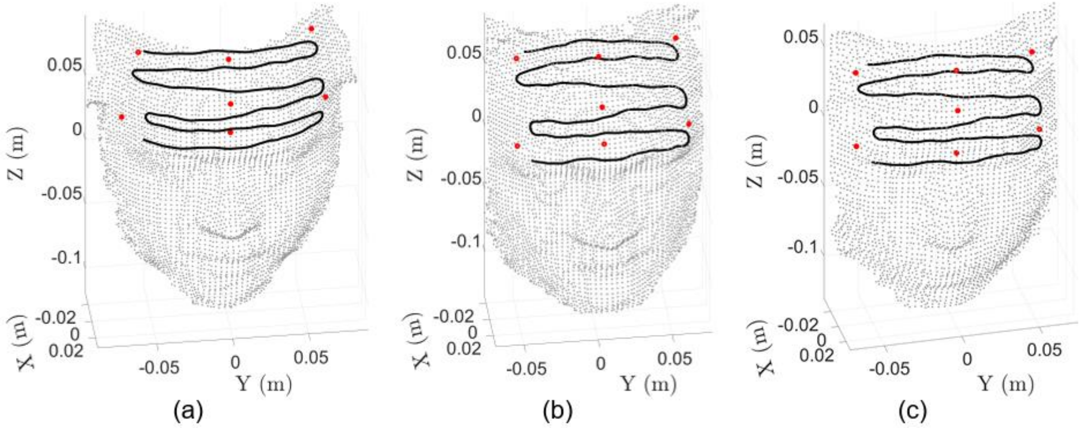

Duan, A. et al. "Learning Rhythmic Trajectories with GeometricConstraints for Laser-Based Skincare Procedures" IEEE T-RO 2025

Learning Long-horizon Path Generation on Point-cloud Representations

- **Goal: interacting with 3D objects**
- Point-cloud measurements (no explicit model)
- Object-conditioned planning
- *Long-horizon* tasks
- Leverage 3D Deep Learning architectures for shape completion
- **Extend to path generation**

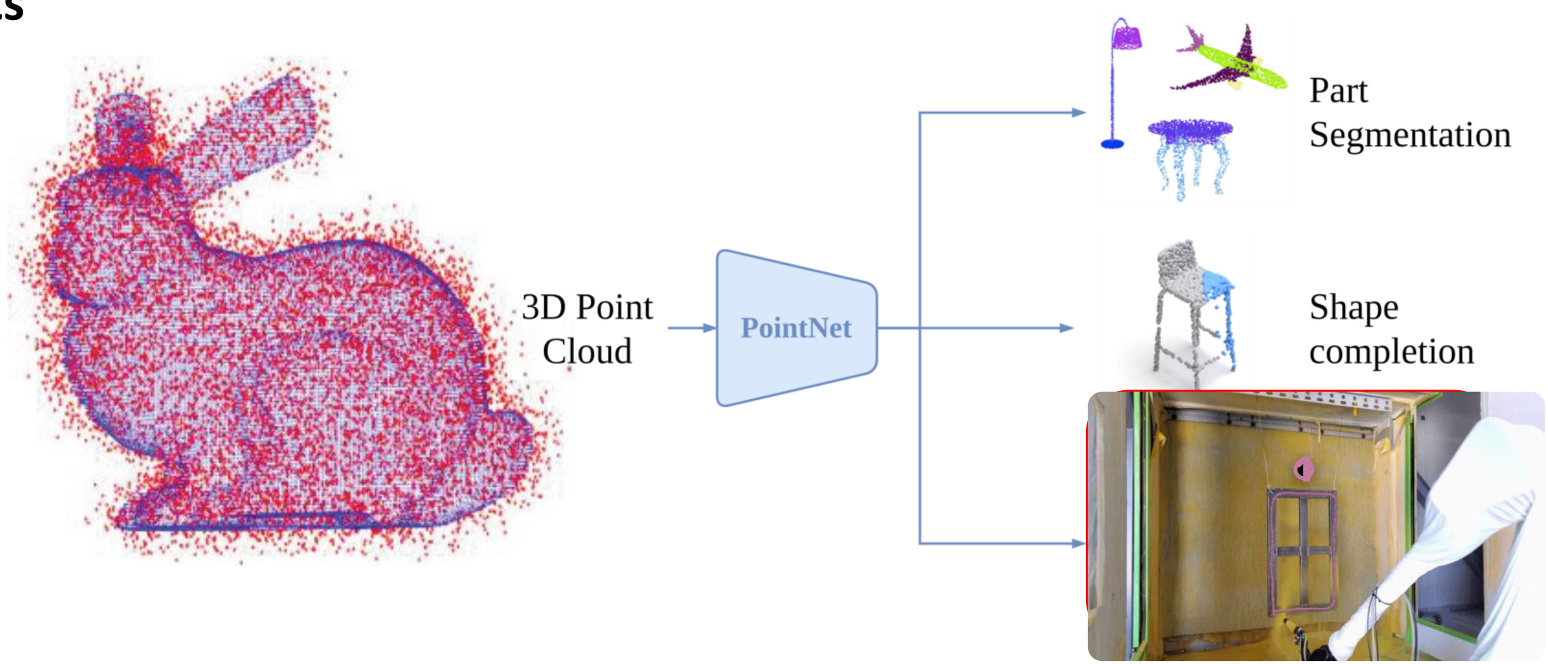

Motion generation

Tiboni, Gabriele, Raffaello Camoriano, and Tatiana Tommasi. "PaintNet: Unstructured Multi-Path Learning from 3D Point Clouds for Robotic Spray Painting.“ IROS 2023.
Tiboni, Gabriele, Raffaello Camoriano, and Tatiana Tommasi. "MaskPlanner: Learning-Based Object-Centric Motion Generation from 3D Point Clouds." Under review

Incremental Robot Learning

Inverse Dynamics Problem

Learn the mapping from

- Joint positions
- Joint velocities
- Joint accelerations

To

- Forces/torques

$$\boldsymbol{\tau} = ID(q, \dot{q}, \ddot{q})$$

Useful for:

- Accurate motion and locomotion control
- Contact detection
- Sim-to-real transfer

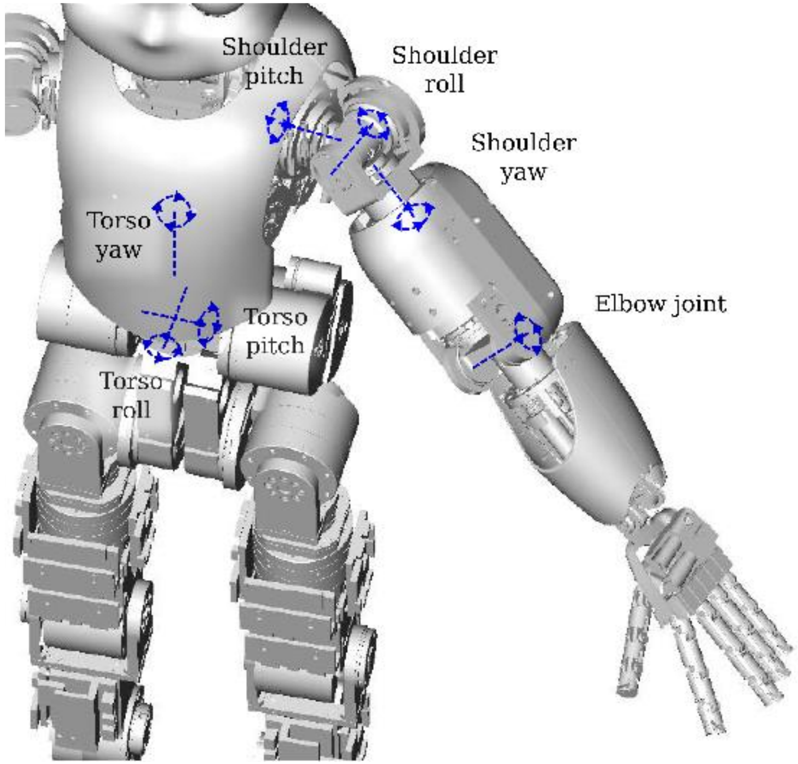

Parametric Modeling

Based on rigid body dynamics (RBD)

$$\boldsymbol{\tau} = M(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + g(\mathbf{q}) = \Phi(\mathbf{x})\boldsymbol{\pi}$$

Parameters $\boldsymbol{\pi}$ estimated from data, or computed from CAD model

Pros

- Interpretability
- Good performance in entire workspace

Cons

- Non-rigid dynamics not captured

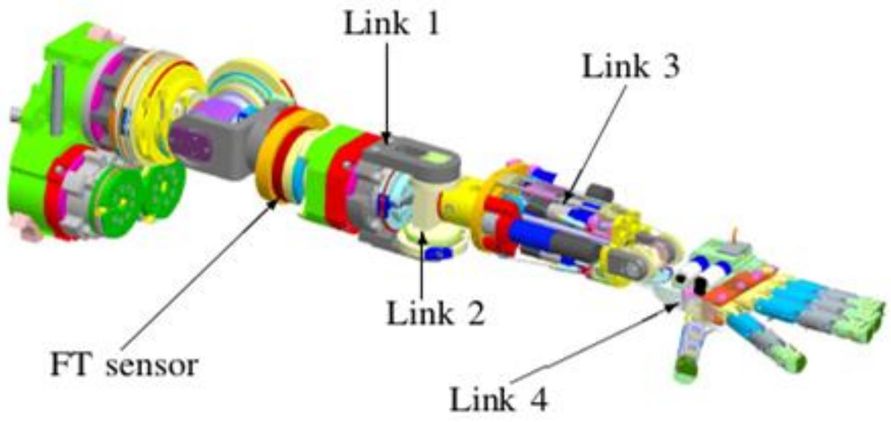

Nonparametric Modeling

Models ID as Black-box function, e.g., KRR, NN

Fully data-driven, minimal assumptions

Pros

- Models non-linear effects, not just RBD
- Higher accuracy

Cons

- Performance depends on distribution of data in workspace
- Model is not interpretable

Incremental Learning

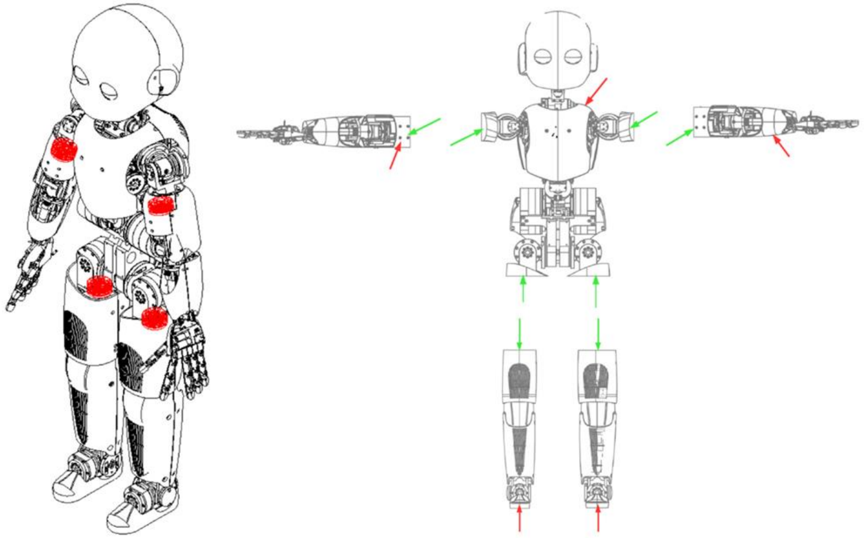

Robot dynamics can be ***time-varying***

- Wear & tear
- Temperature drift

→ Batch model performance may decline

- **Solution:** Use an *incremental* model

→ Recursive Regularized Least Squares (RRLS)
→ Random Features for non-linear kernel approximation [1]

[1] Gijsberts, Arjan, and Giorgio Metta. "Incremental learning of robot dynamics using random features." IEEE ICRA, 2011.

Incremental Semiparametric Dynamics Learning

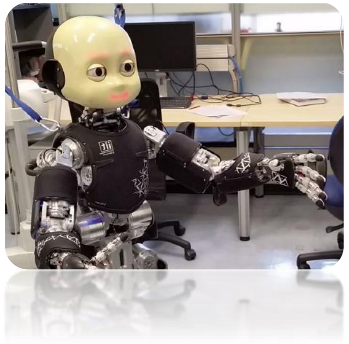

iCub

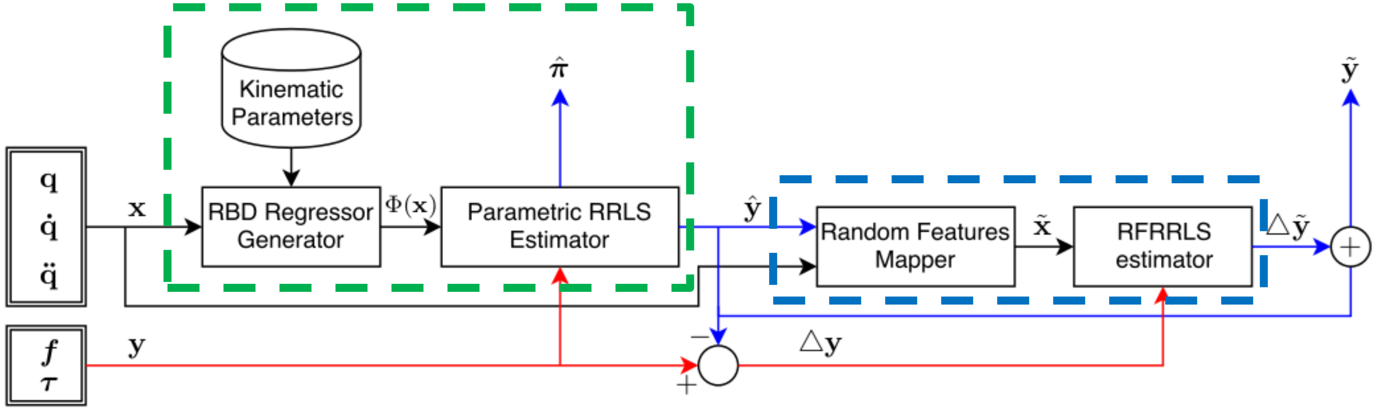

Parametric Model

Nonparametric Model

Idea 1: Exploit prior knowledge about system's physics (parametric SysID)

Idea 2: Nonparametric kernel learning on residuals

Incremental with *O(d)* update complexity

R. C., S. Traversaro, L. Rosasco, G. Metta, and F. Nori. "Incremental Semiparametric Inverse Dynamics Learning". *IEEE ICRA*, 2016.

Results

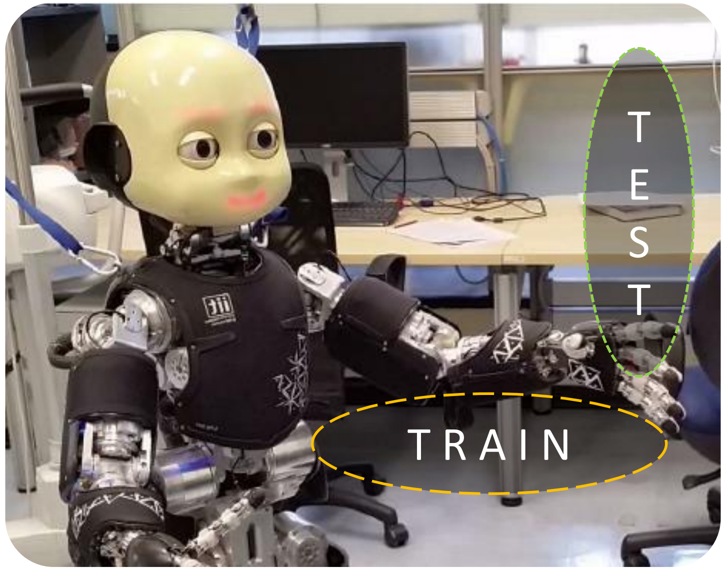

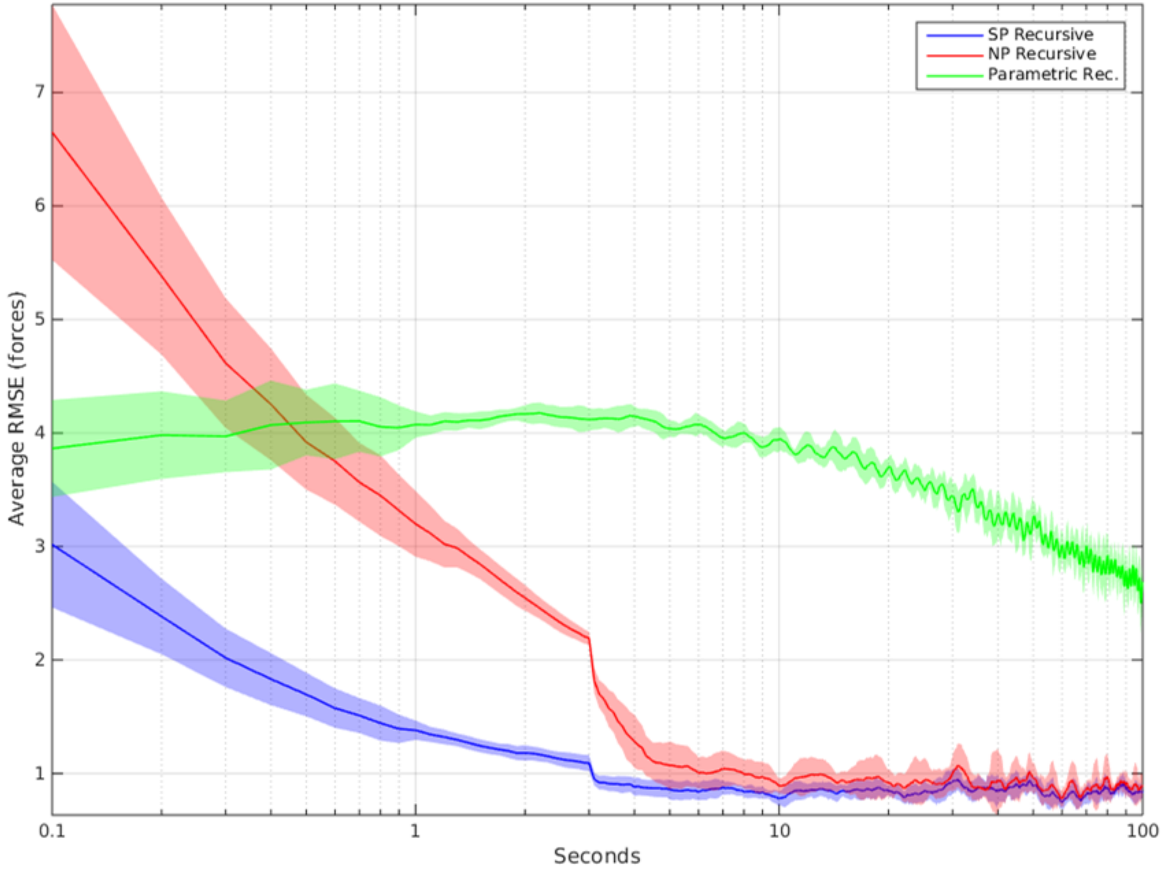

- ✓ Higher accuracy
- ✓ Better generalization
- ✓ Efficient model updates

Follow-up works, incorporating uncertainty quantification:

Romeres, D., et al. "Online semi-parametric learning for inverse dynamics modeling." *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016.

Romeres, D., et al. "Derivative-free online learning of inverse dynamics models." *IEEE Transactions on Control Systems Technology* 28.3 (2019): 816-830.

Incremental Classification for Prosthesis Control

- **Goal:** restoring controllable movements in limited-mobility patients

- Hannes Hand (+ wrist, + elbow)
  - Low-cost
  - 5+ DoFs
  - Myoelectric sensors (60+)

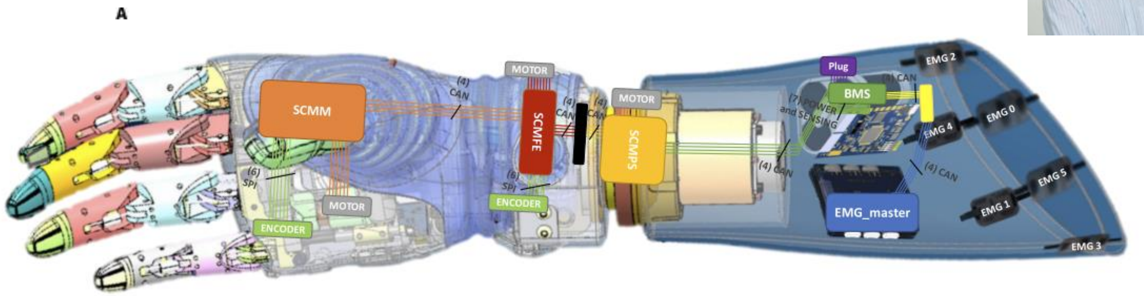

Figure 5.3: *Hannes* system with electronic parts

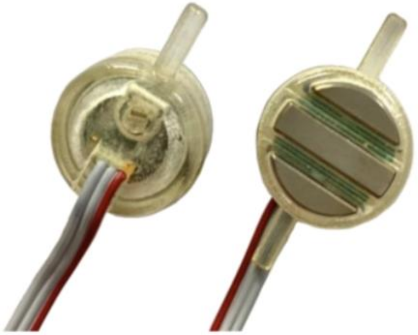

Figure 5.6: IIT sEMG electrodes

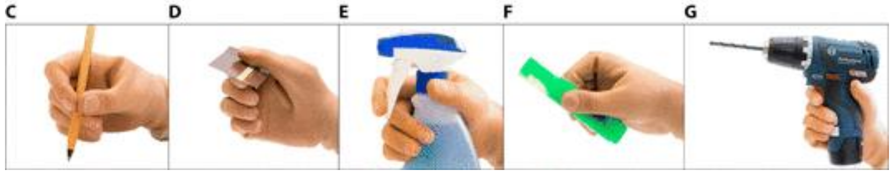

Incremental Classification for Prosthesis Control

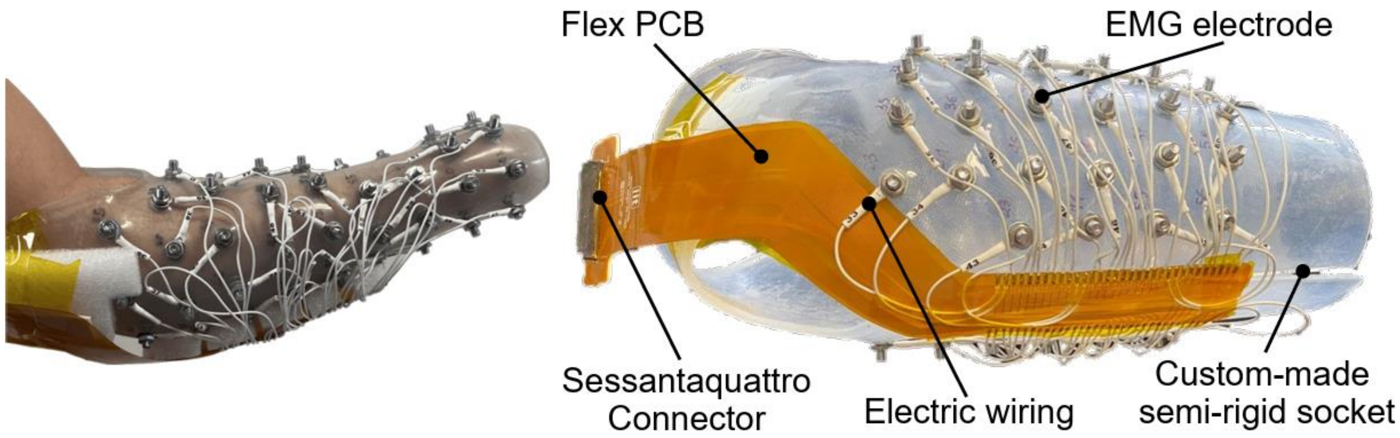

- **Classification problem:**
  - Map signals to useful movements
  - Input: myoelectric time series
  - Output: 6+ classes

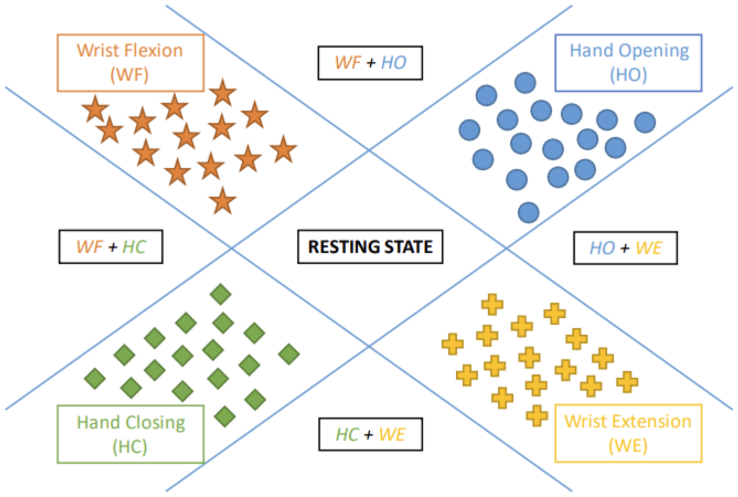

Figure 5.13: Combination of four classes and the common resting state

Incremental Classification for Prosthesis Control

- **Classification problem:**
  - Map signals to useful movements
  - Input: myoelectric time series
  - Output: 6+ classes

- **Warning:** I/O mapping shifts in time

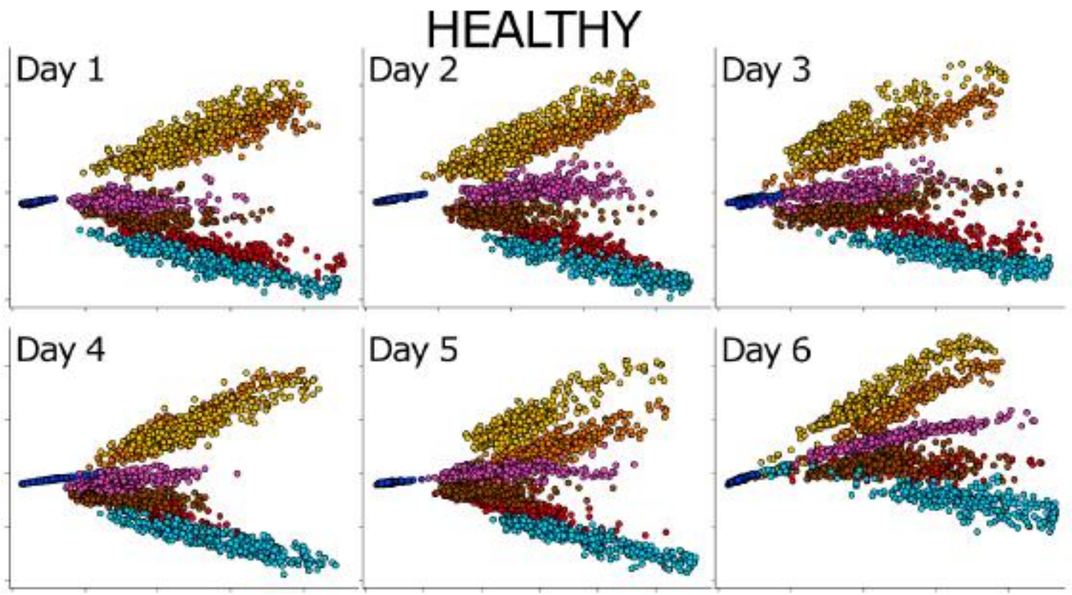

Incremental Classification for Prosthesis Control

- **Classification problem:**
  - Map signals to useful movements
  - Input: myoelectric time series
  - Output: 6+ classes
- **Warning:** I/O mapping shifts in time
- **Our solutions:**
  Incremental/Continual Learning
  Domain Adaptation

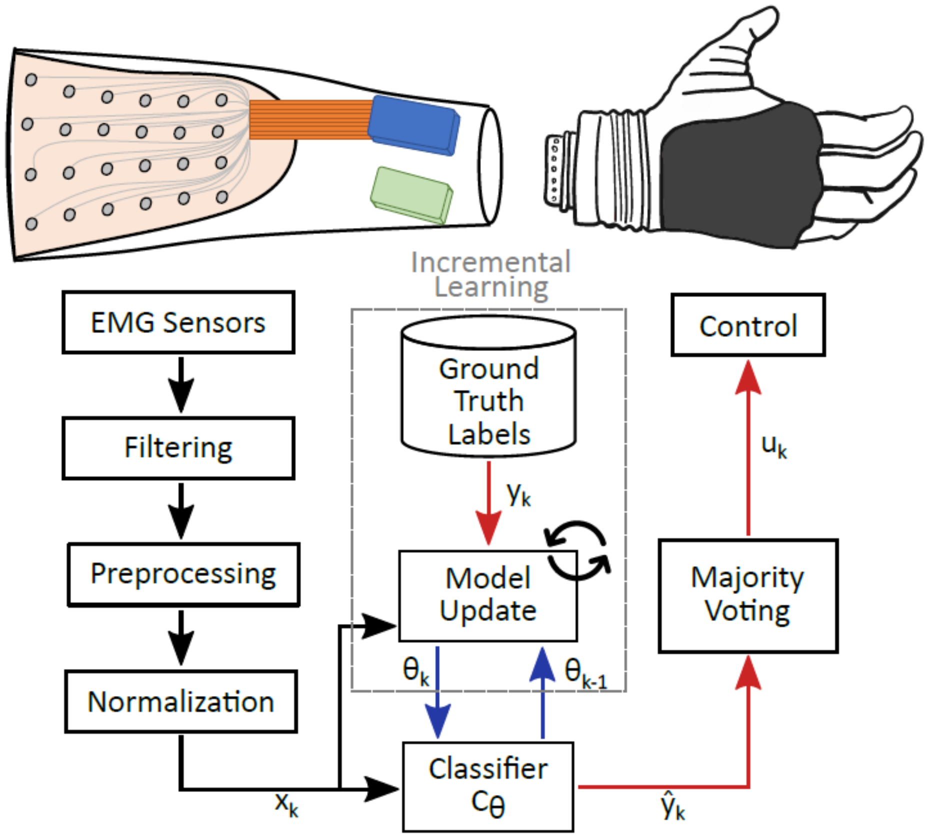

Incremental Classification for Prosthesis Control

Results across days:

- Batch methods yield degrading performance
- Incremental RFRLSC adapts to shifting distribution

More accurate control and improved ergonomy

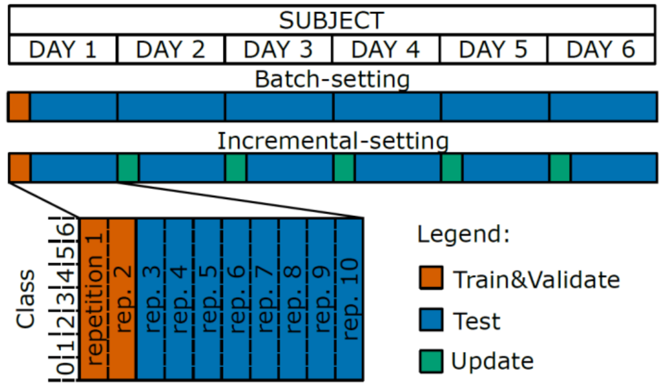

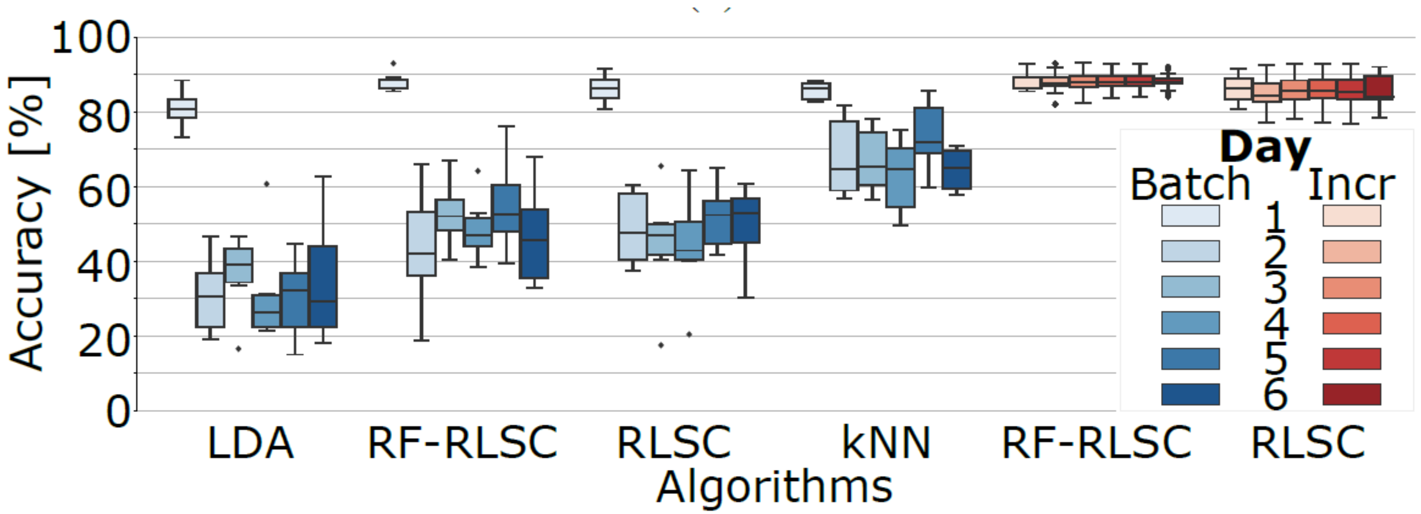

Di Domenico, Dario, et al. "Long-Term Upper-Limb Prosthesis Myocontrol via High-Density sEMG and Incremental Learning." *IEEE Robotics and Automation Letters* (2024).

Results on Amputee Subject



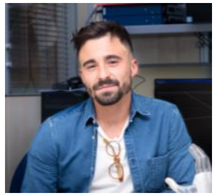

Accuracy across days

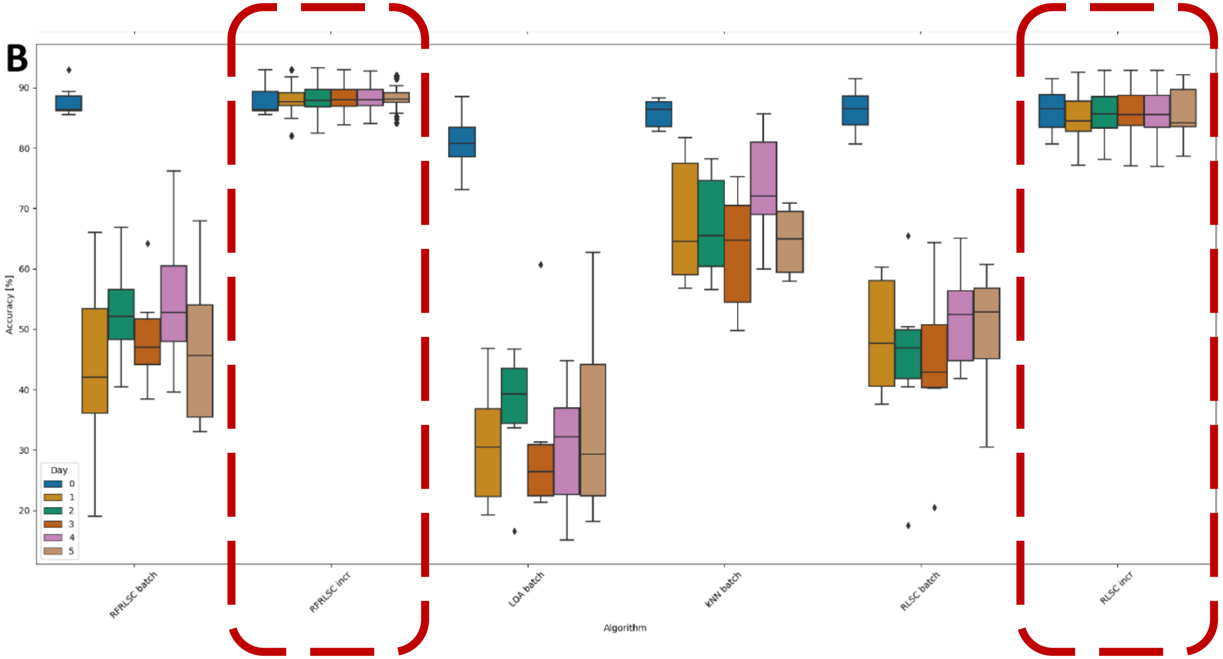

Incremental Kernel Classifier

Incremental Linear Classifier

Linear vs. kernelized classifier

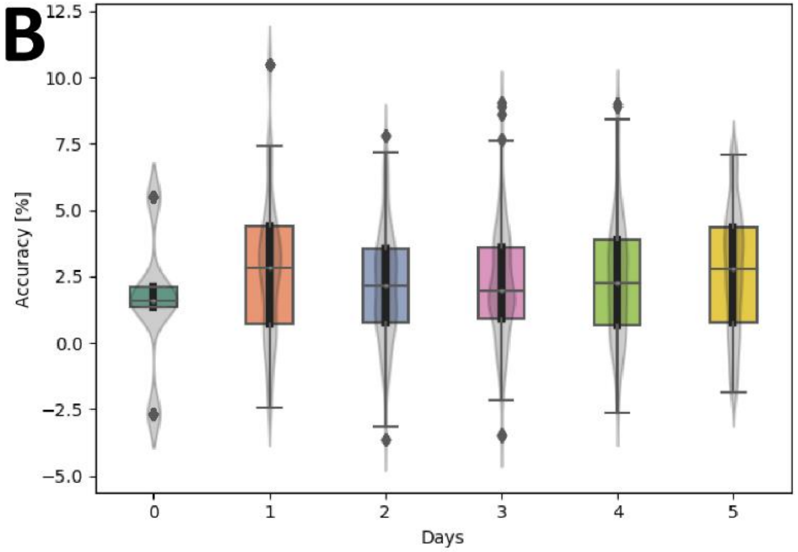

Kernel (RF) nonlinear mapping helps

Learning Humanoid Locomotion

Learning High-dimensional Robot Controllers

Complex structure
- Non-IID data
- Agent-Environment interaction (MDP)

Efficiency Real-time prediction requirements

We investigate Reinforcement and Imitation Learning methods in high-dimensional humanoid robotics settings

Reinforcement Learning

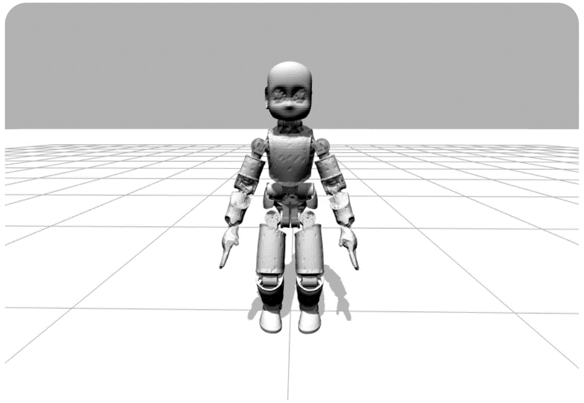

(Supervised) Imitation Learning

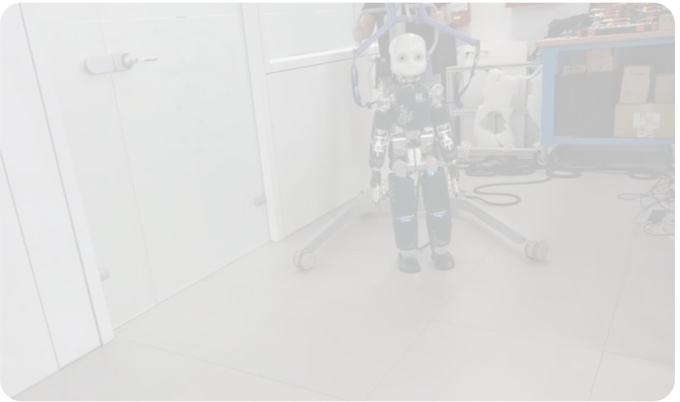

Reinforcement Learning
for Humanoid Push Recovery

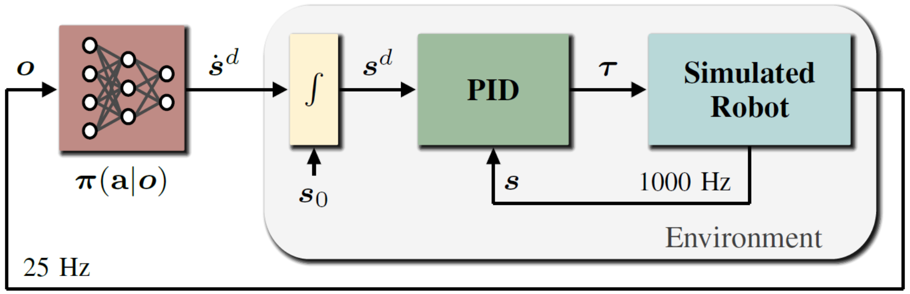

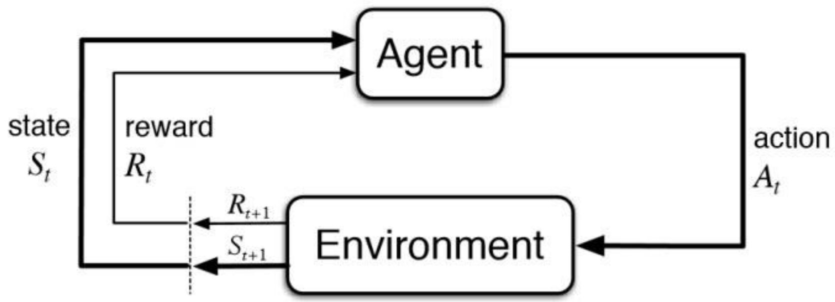

Model-free Deep Reinforcement Learning (DRL) for high-dimensional humanoid balancing and push recovery

- Single policy → Many whole-body strategies, including momentum-based
- DRL can be successful for whole-body robot control
- Extensive evaluation of robustness and generalization in simulation

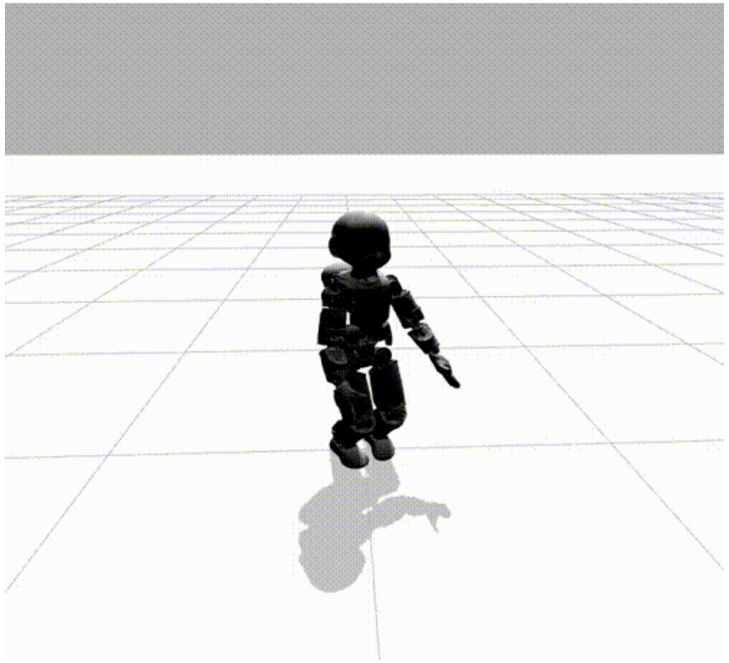

Ferigo, D.*, **R. C.***, et al. "On the Emergence of Whole-body Strategies from Humanoid Robot Push-recovery Learning." *IEEE RA-L & Humanoids* (2021).

Classical Approaches

- Control-theoretic methods significantly improved the state-of-the-art push-recovery performances of humanoids.
- Still, they present some limitations:
  - They usually encode a single strategy. Counteracting diverse perturbations requires designing multiple controllers and complex switching rules
  - Expensive robot- and task-specific tuning of the controllers and switching system
  - Simplified models and hard-coded strategies constrain behaviors and performance
  - MPC-based methods are computationally expensive, limiting real-time deployment

Our Approach

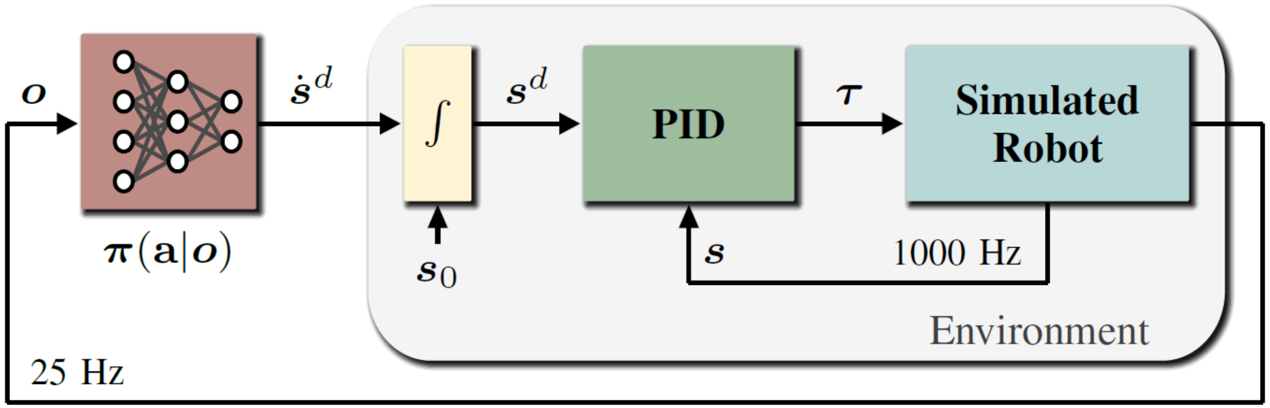

- Model-free Deep Reinforcement Learning (DRL) for high-dimensional humanoid balancing and push recovery

- Single policy → Many whole-body strategies, including momentum-based
- We show DRL can be successful for whole-body robot control **(23 DoF)**
- Parsimonious observation space
- Reward design from first principles in robot control
- Increased robustness via domain randomization
- Extensive evaluation of robustness and generalization in simulation

Observations

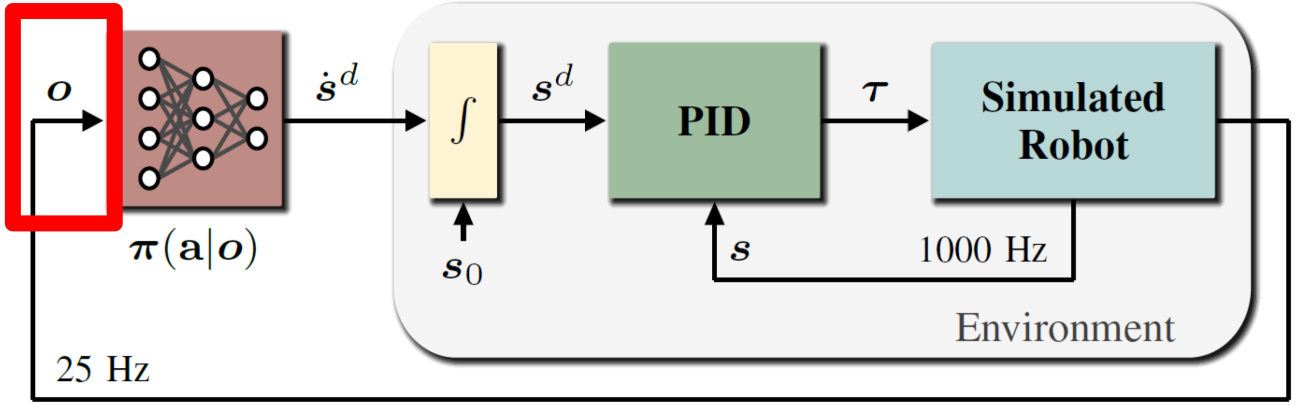

Observation space design:

- Low dimensionality ($o \in \mathbb{R}^{62}$)
- Feasible to obtain on simulated or real robots:

- ➢Joint positions
- ➢Joint velocities
- ➢Base height
- ➢Base orientation
- ➢Center of Mass velocity
- ➢Feet contacts configuration
- ➢Feet positions
- ➢Feet Center of Pressure forces

Actions

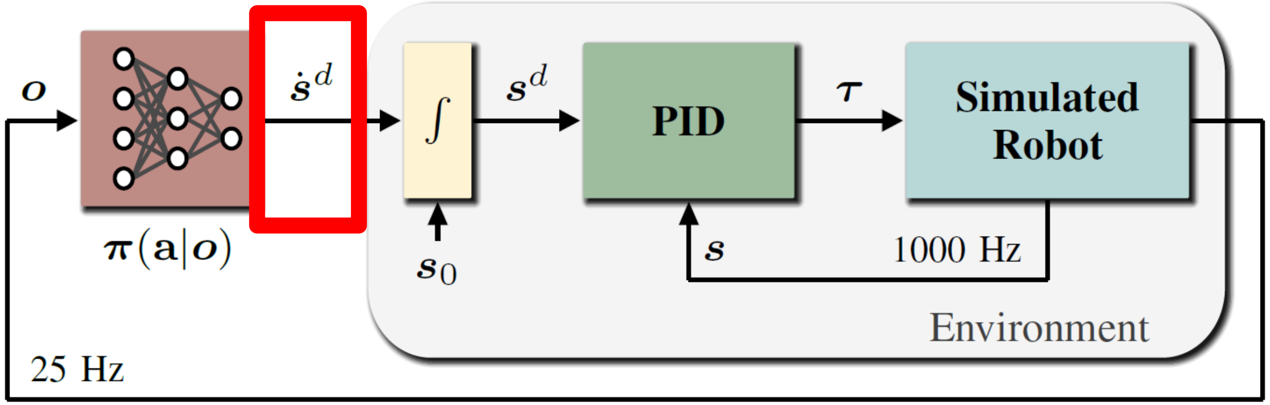

Policy outputs desired joint velocities:

- Smoother motion
- Euler integration required to keep position control
- 23 DoF
- Policy: 25 Hz
- Low-level control (PIDs): 1kHz

Reward

Linear combination of subtask components, each with a target value

Steady-state   Transient   Regularizers

- Postural Task
- CoM Projection
- CoM Horizontal Velocity
- CoM Vertical Velocity
- Whole-body Momentum
- Links in Contact

- Feet in Contact
- Feet Contact Forces
- Feet Orientation
- Feet CoP
- Joint Velocities
- Joint Torques

Generalization and Robustness

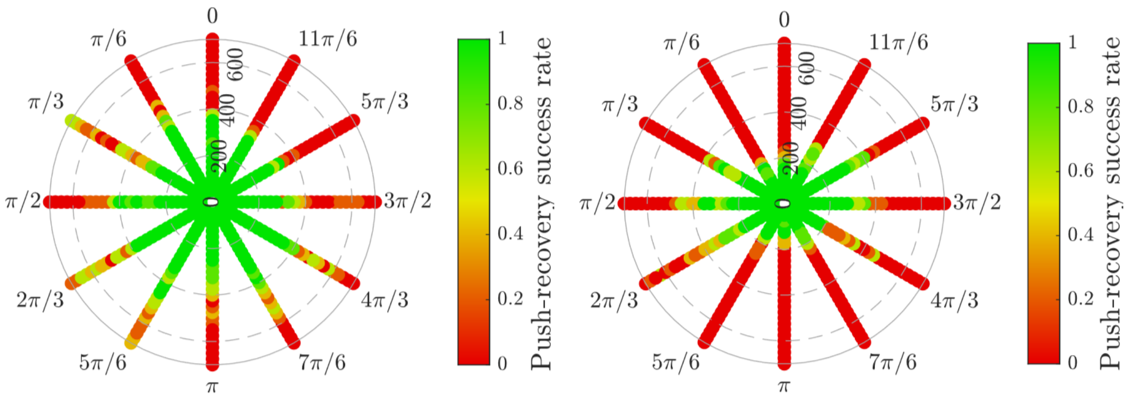

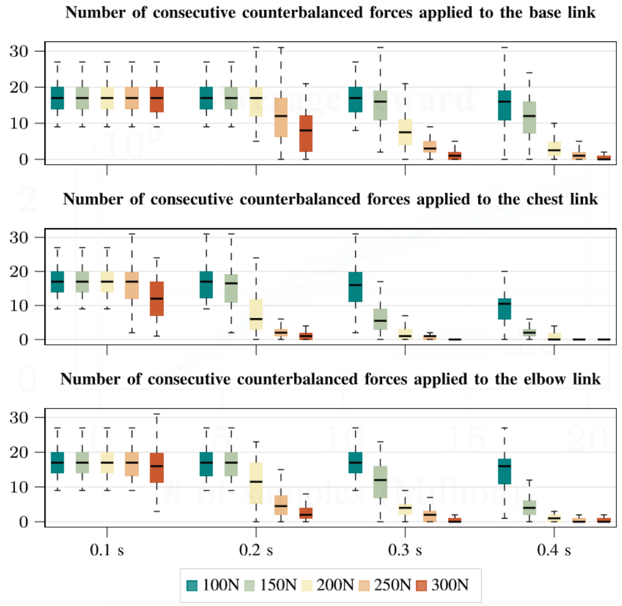

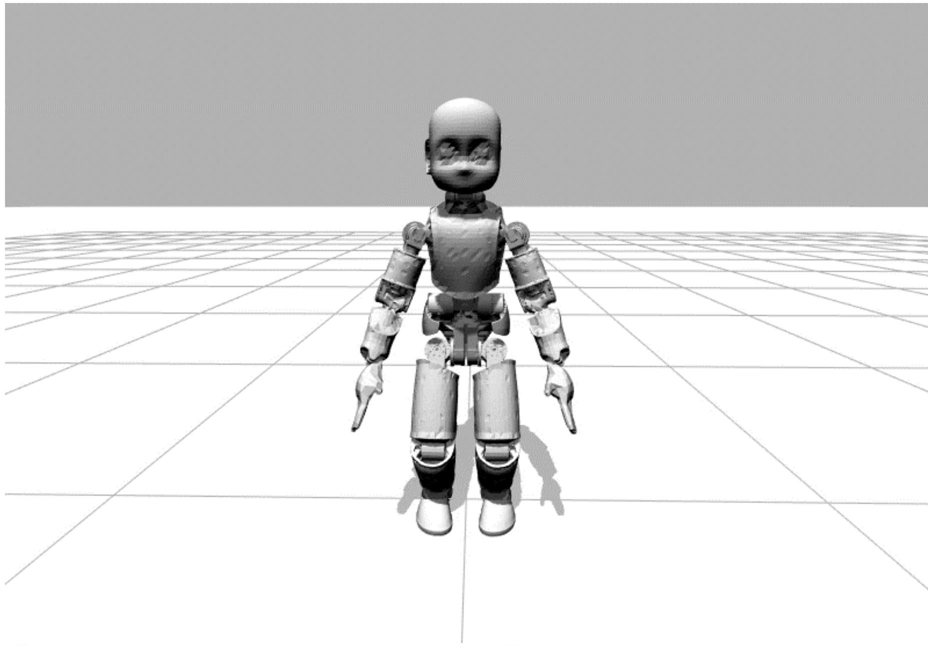

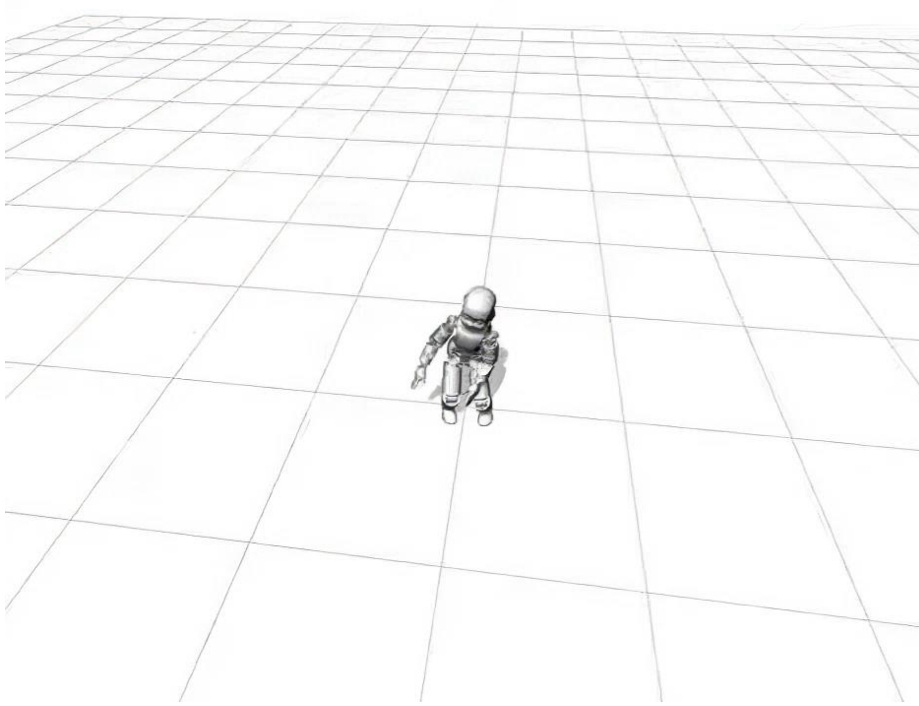

Learning Performance

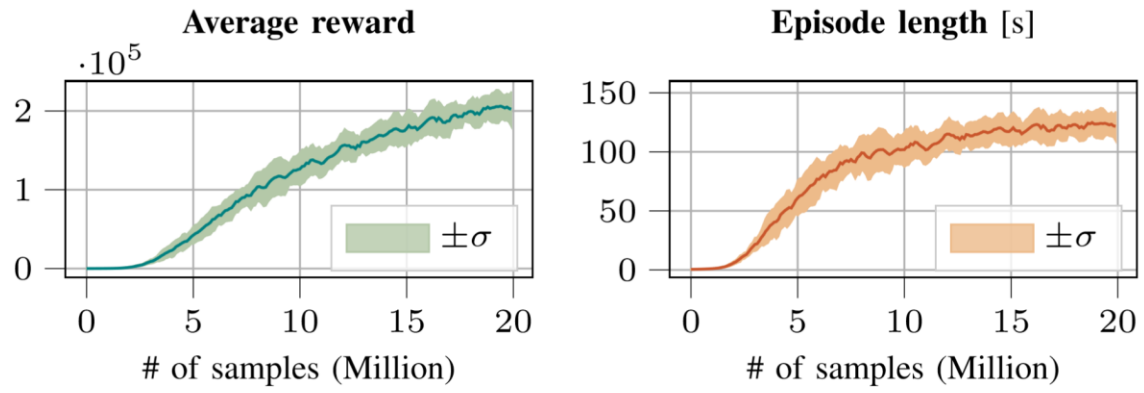

- Consistent reward and episode length growth
- Low variance (11 runs)
- Converges in 20M samples (quite slow back then)
- Sim-to-real challenges

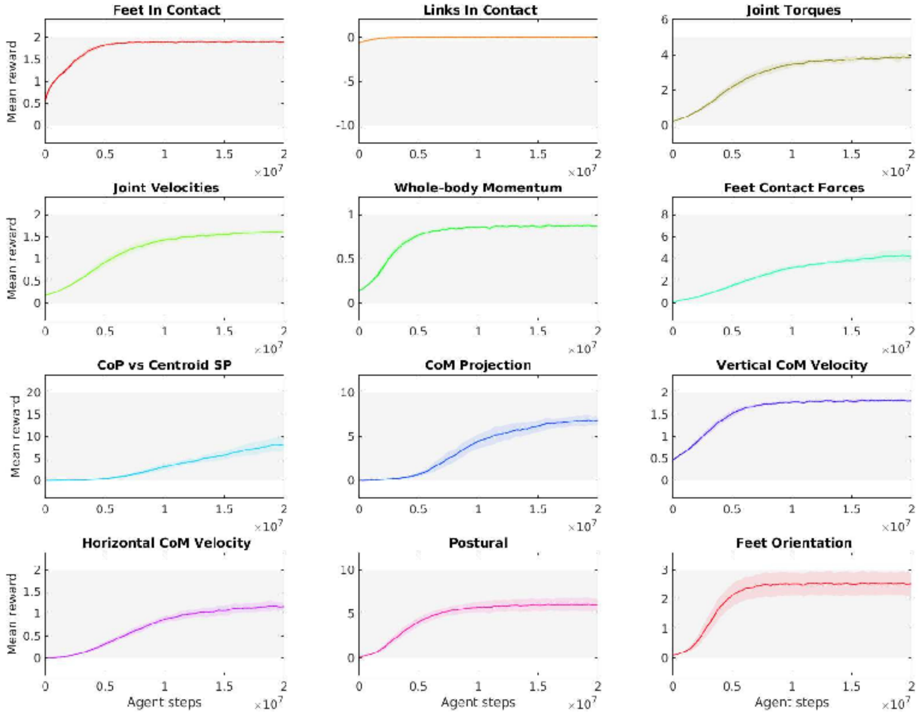

Breakdown by reward term

Learning High-dimensional Robot Controllers

Complex structure
- Non-IID data
- Agent-Environment interaction (MDP)

Efficiency Real-time prediction requirements

We investigate Reinforcement and Imitation Learning methods in high-dimensional humanoid robotics settings

Reinforcement Learning

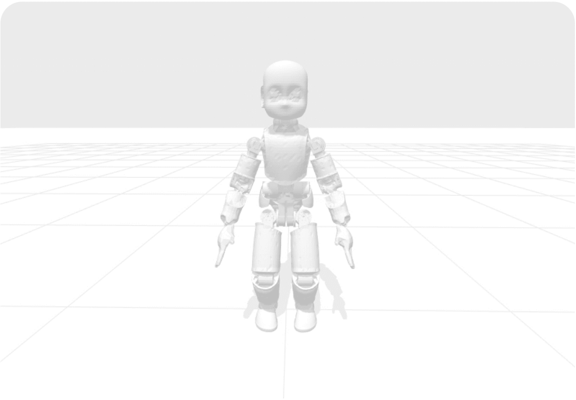

(Supervised) Imitation Learning

ADHERENT

Humanoid Locomotion

High-dimensional trajectory generation: **open problem**

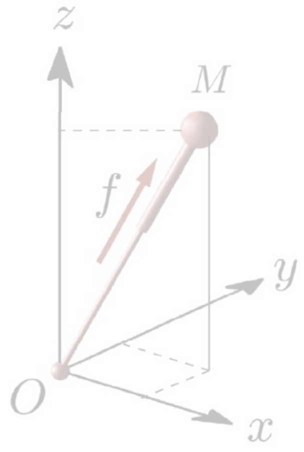

Schperberg et al., “Reducing Motion Perturbation for a Bipedal Robot using Model Predictive Control”, 2019

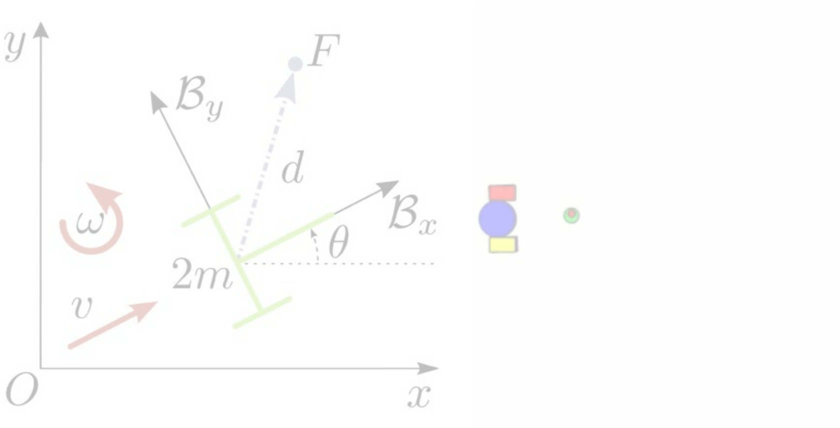

Dafarra et al., “A control architecture with (
Position and Torque Controlled Walking of

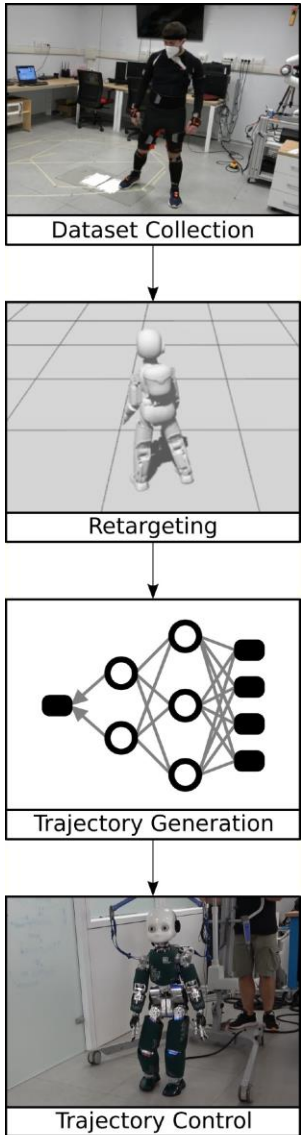

Character Animation

Recent breakthrough: **ML for character animation**

Holden et al., “Phase-functioned neural networks for character control”, *ACM Transactions on Graphics*, 2017

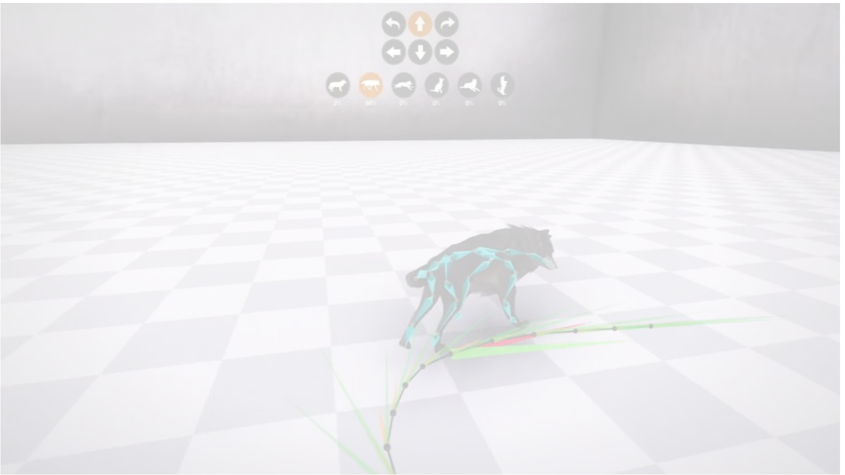

ural Networks for Quadruped Motion
raphics, 2018

Viceconte P., **C. R.**, Romualdi G, et al., “ADHERENT: Learning Human-like Trajectory Generators for Whole-body Control of Humanoid Robots”, IEEE RA-L, 2022

ADHERENT System Architecture

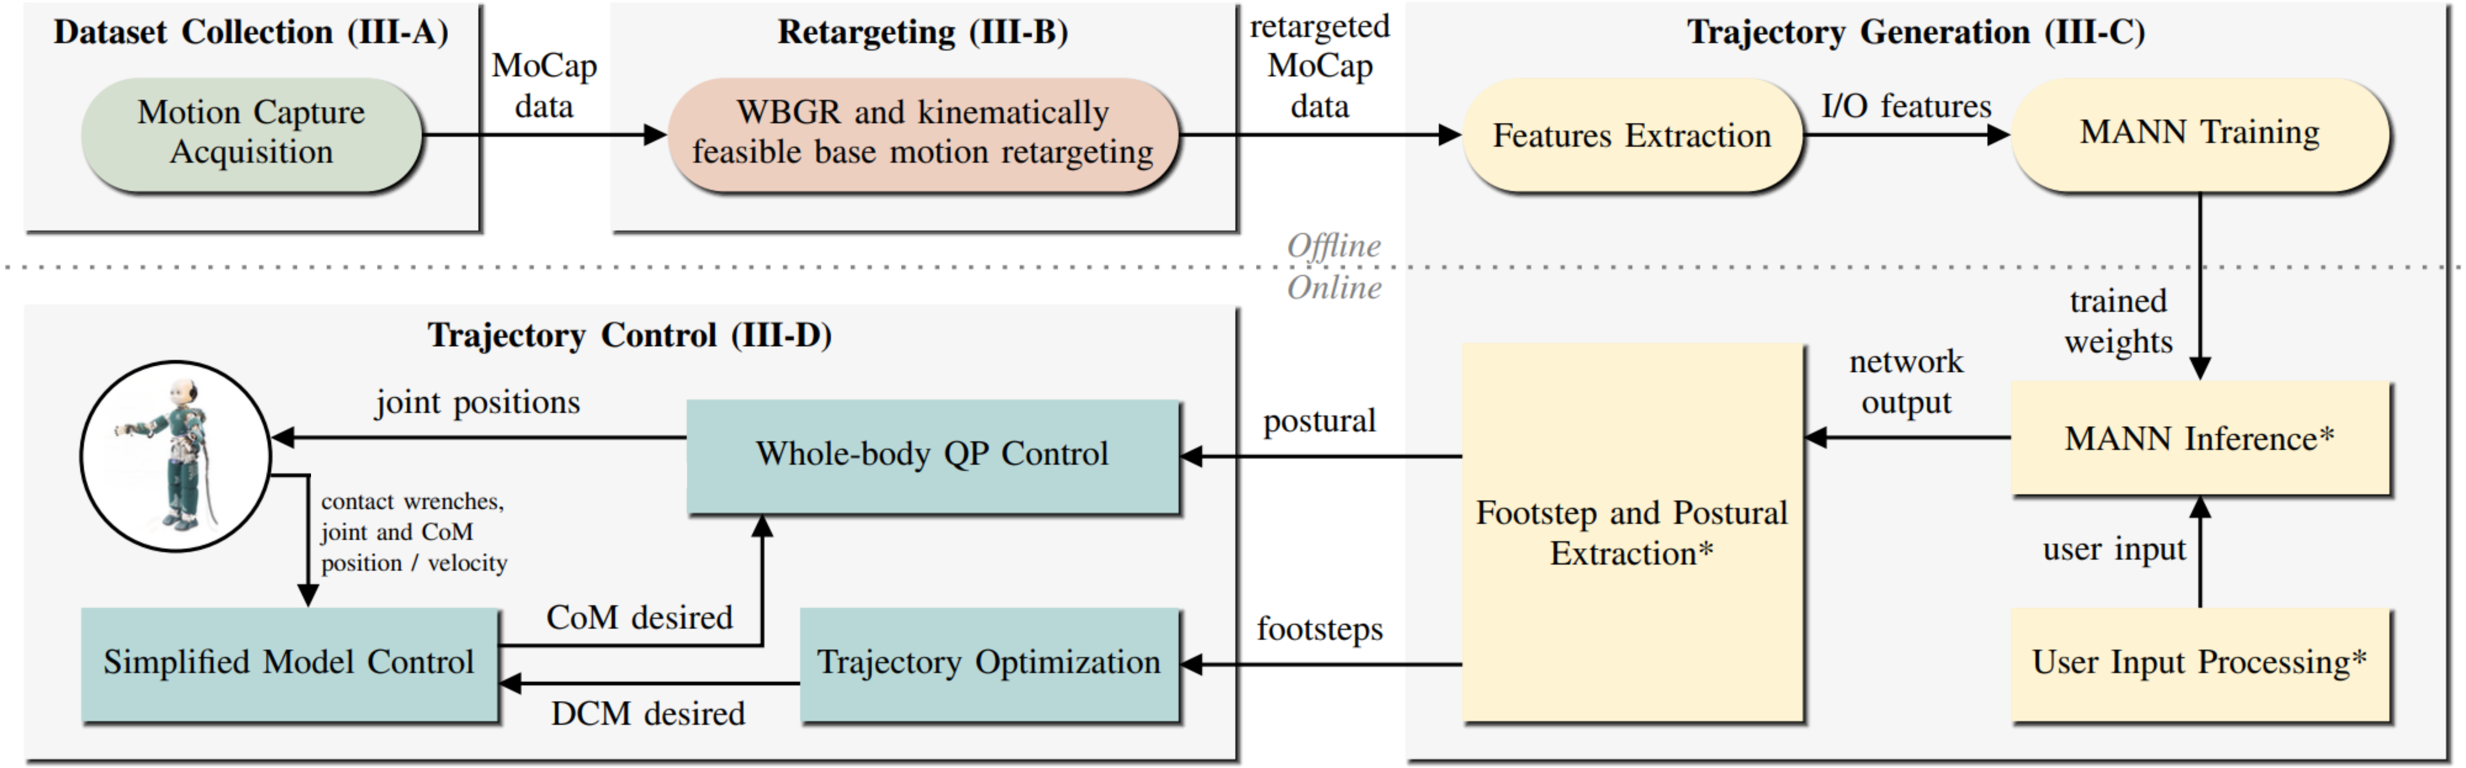

Training Data Collection & Retargeting

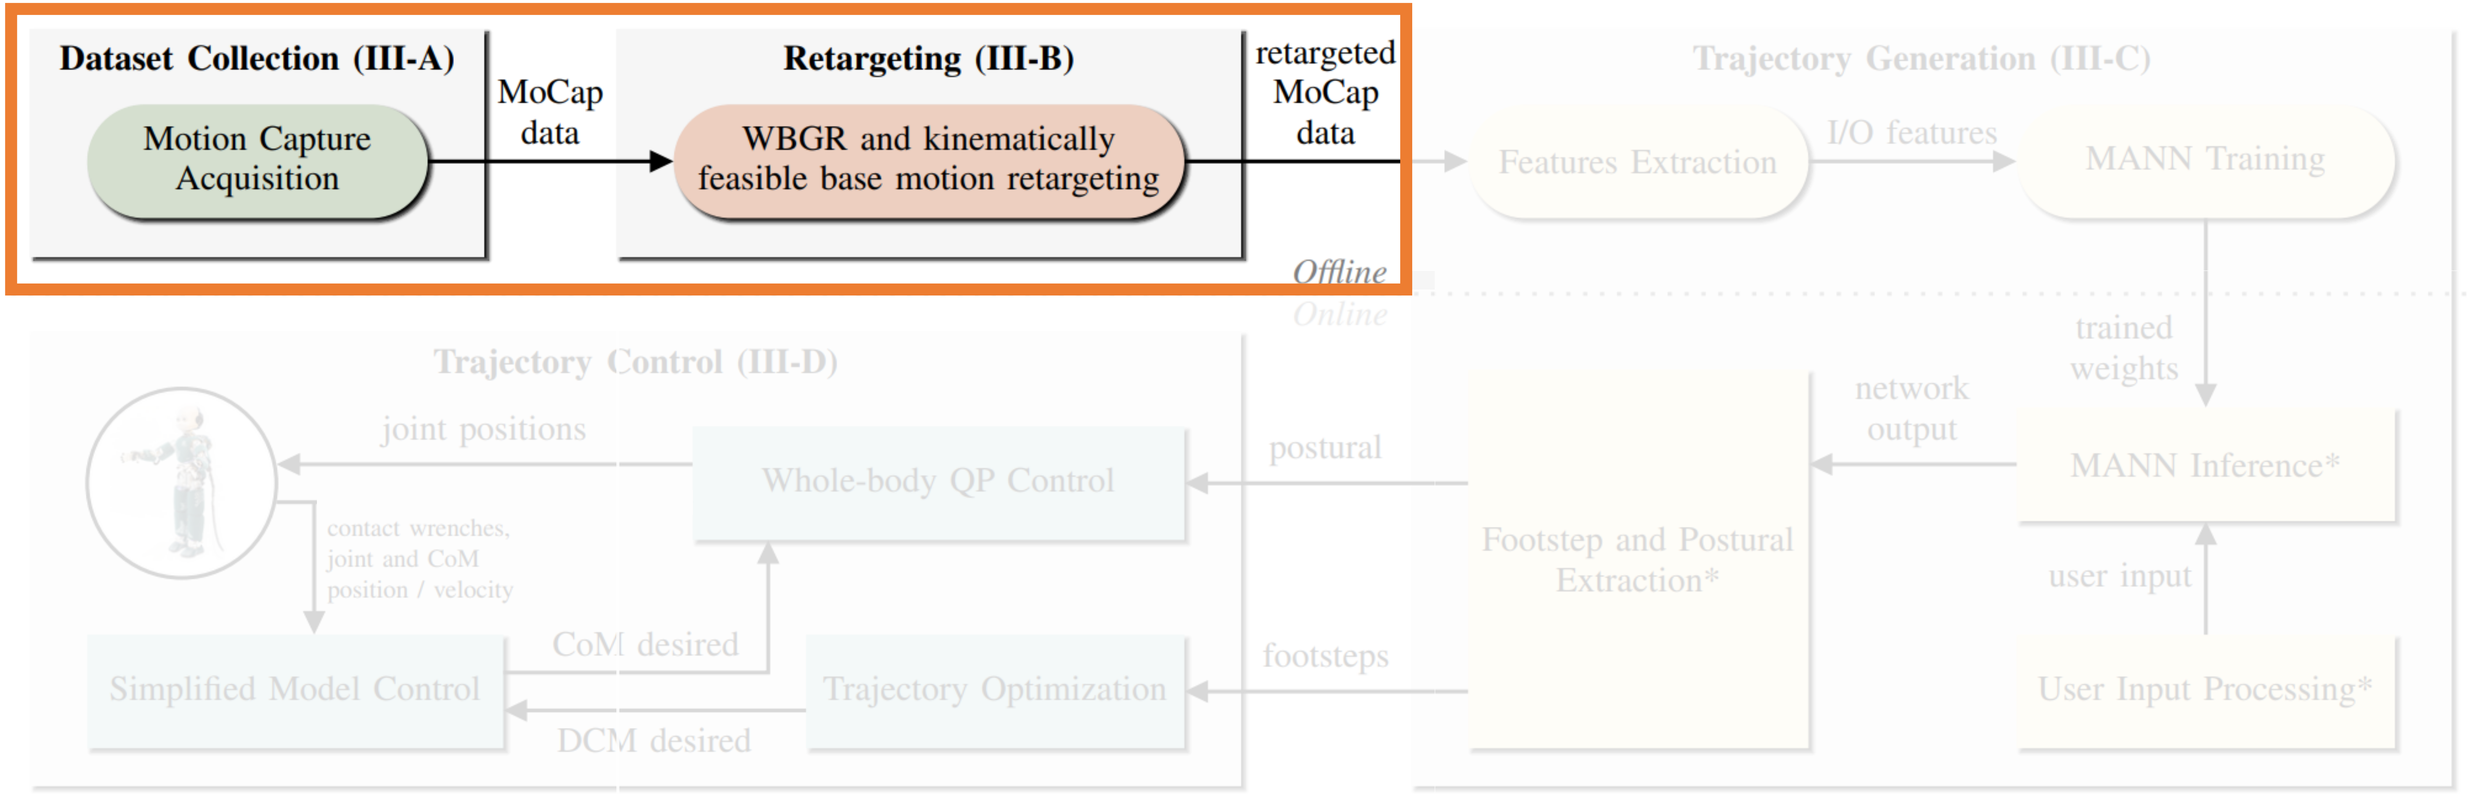

WBGR + Kinematically-feasible Base Retargeting

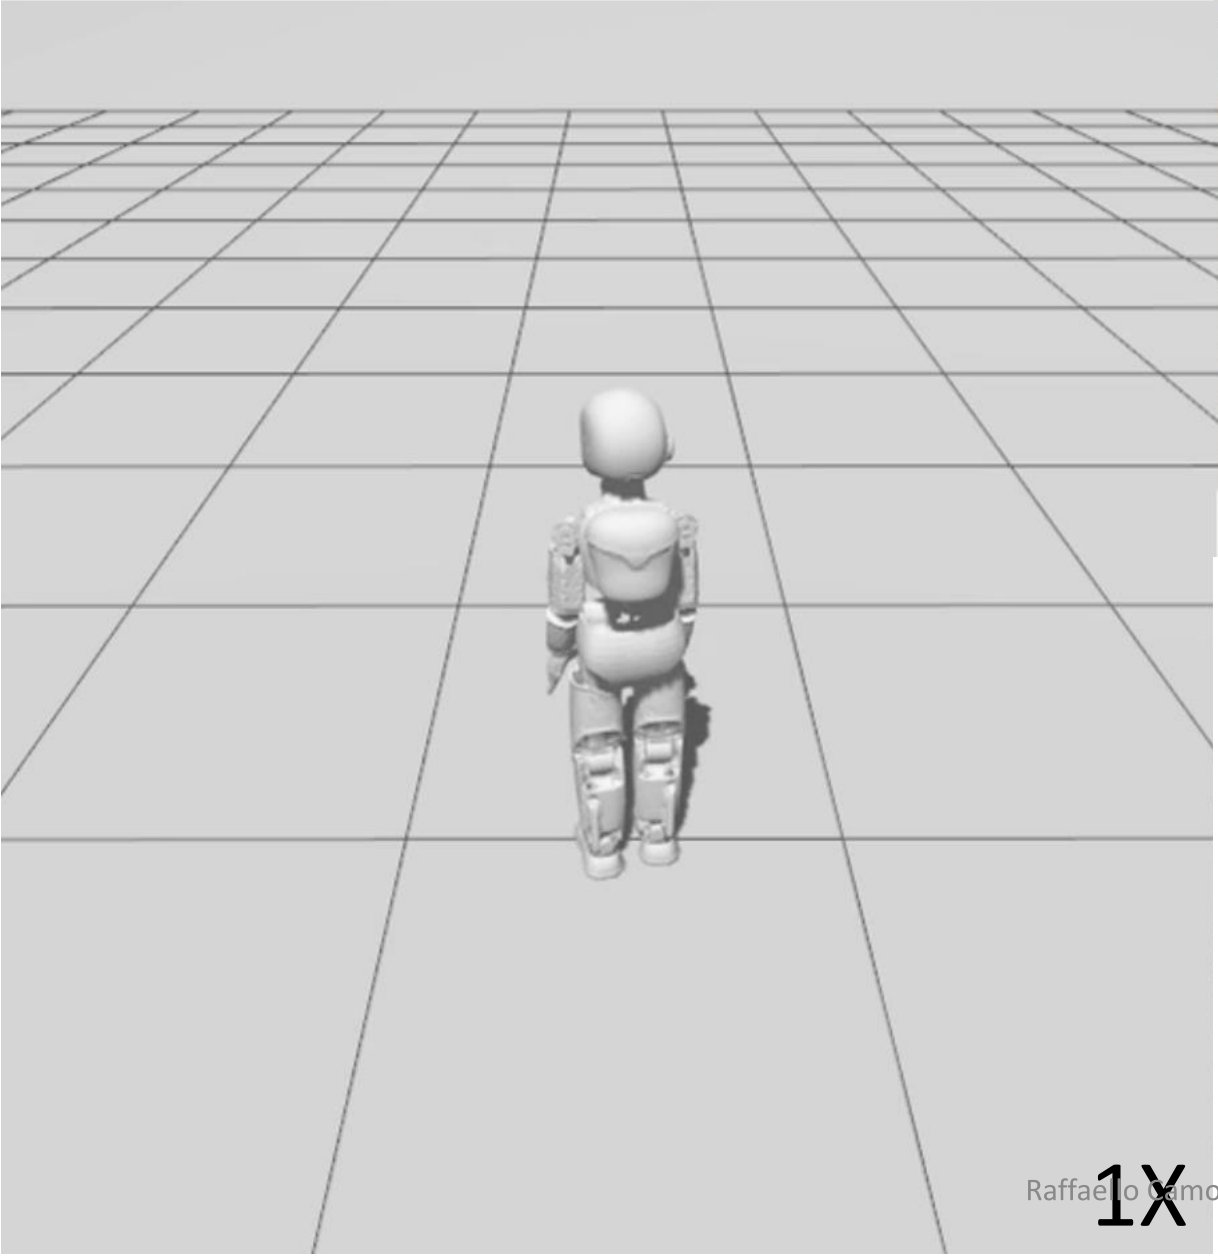

1X

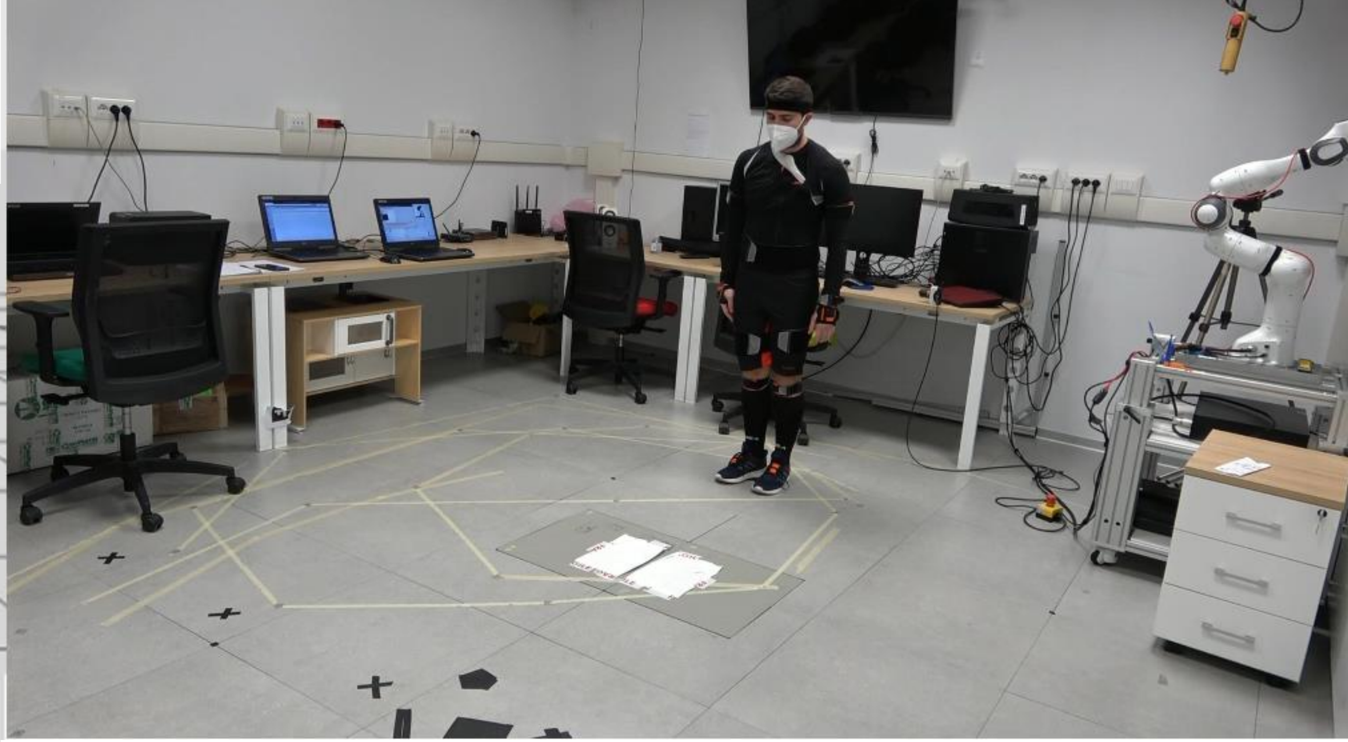

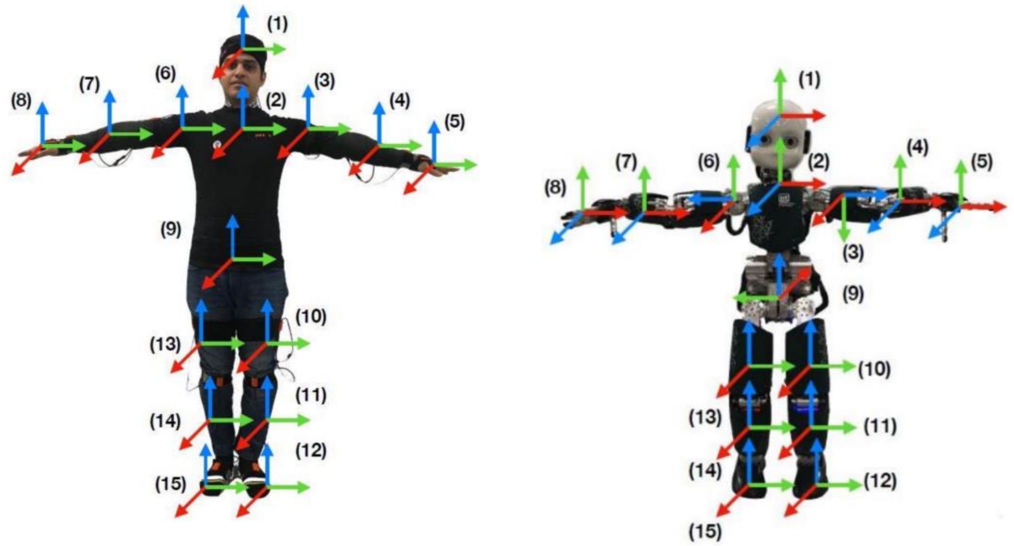

Human/robot frames associations for WBGR

Final Retargeted MoCap Data

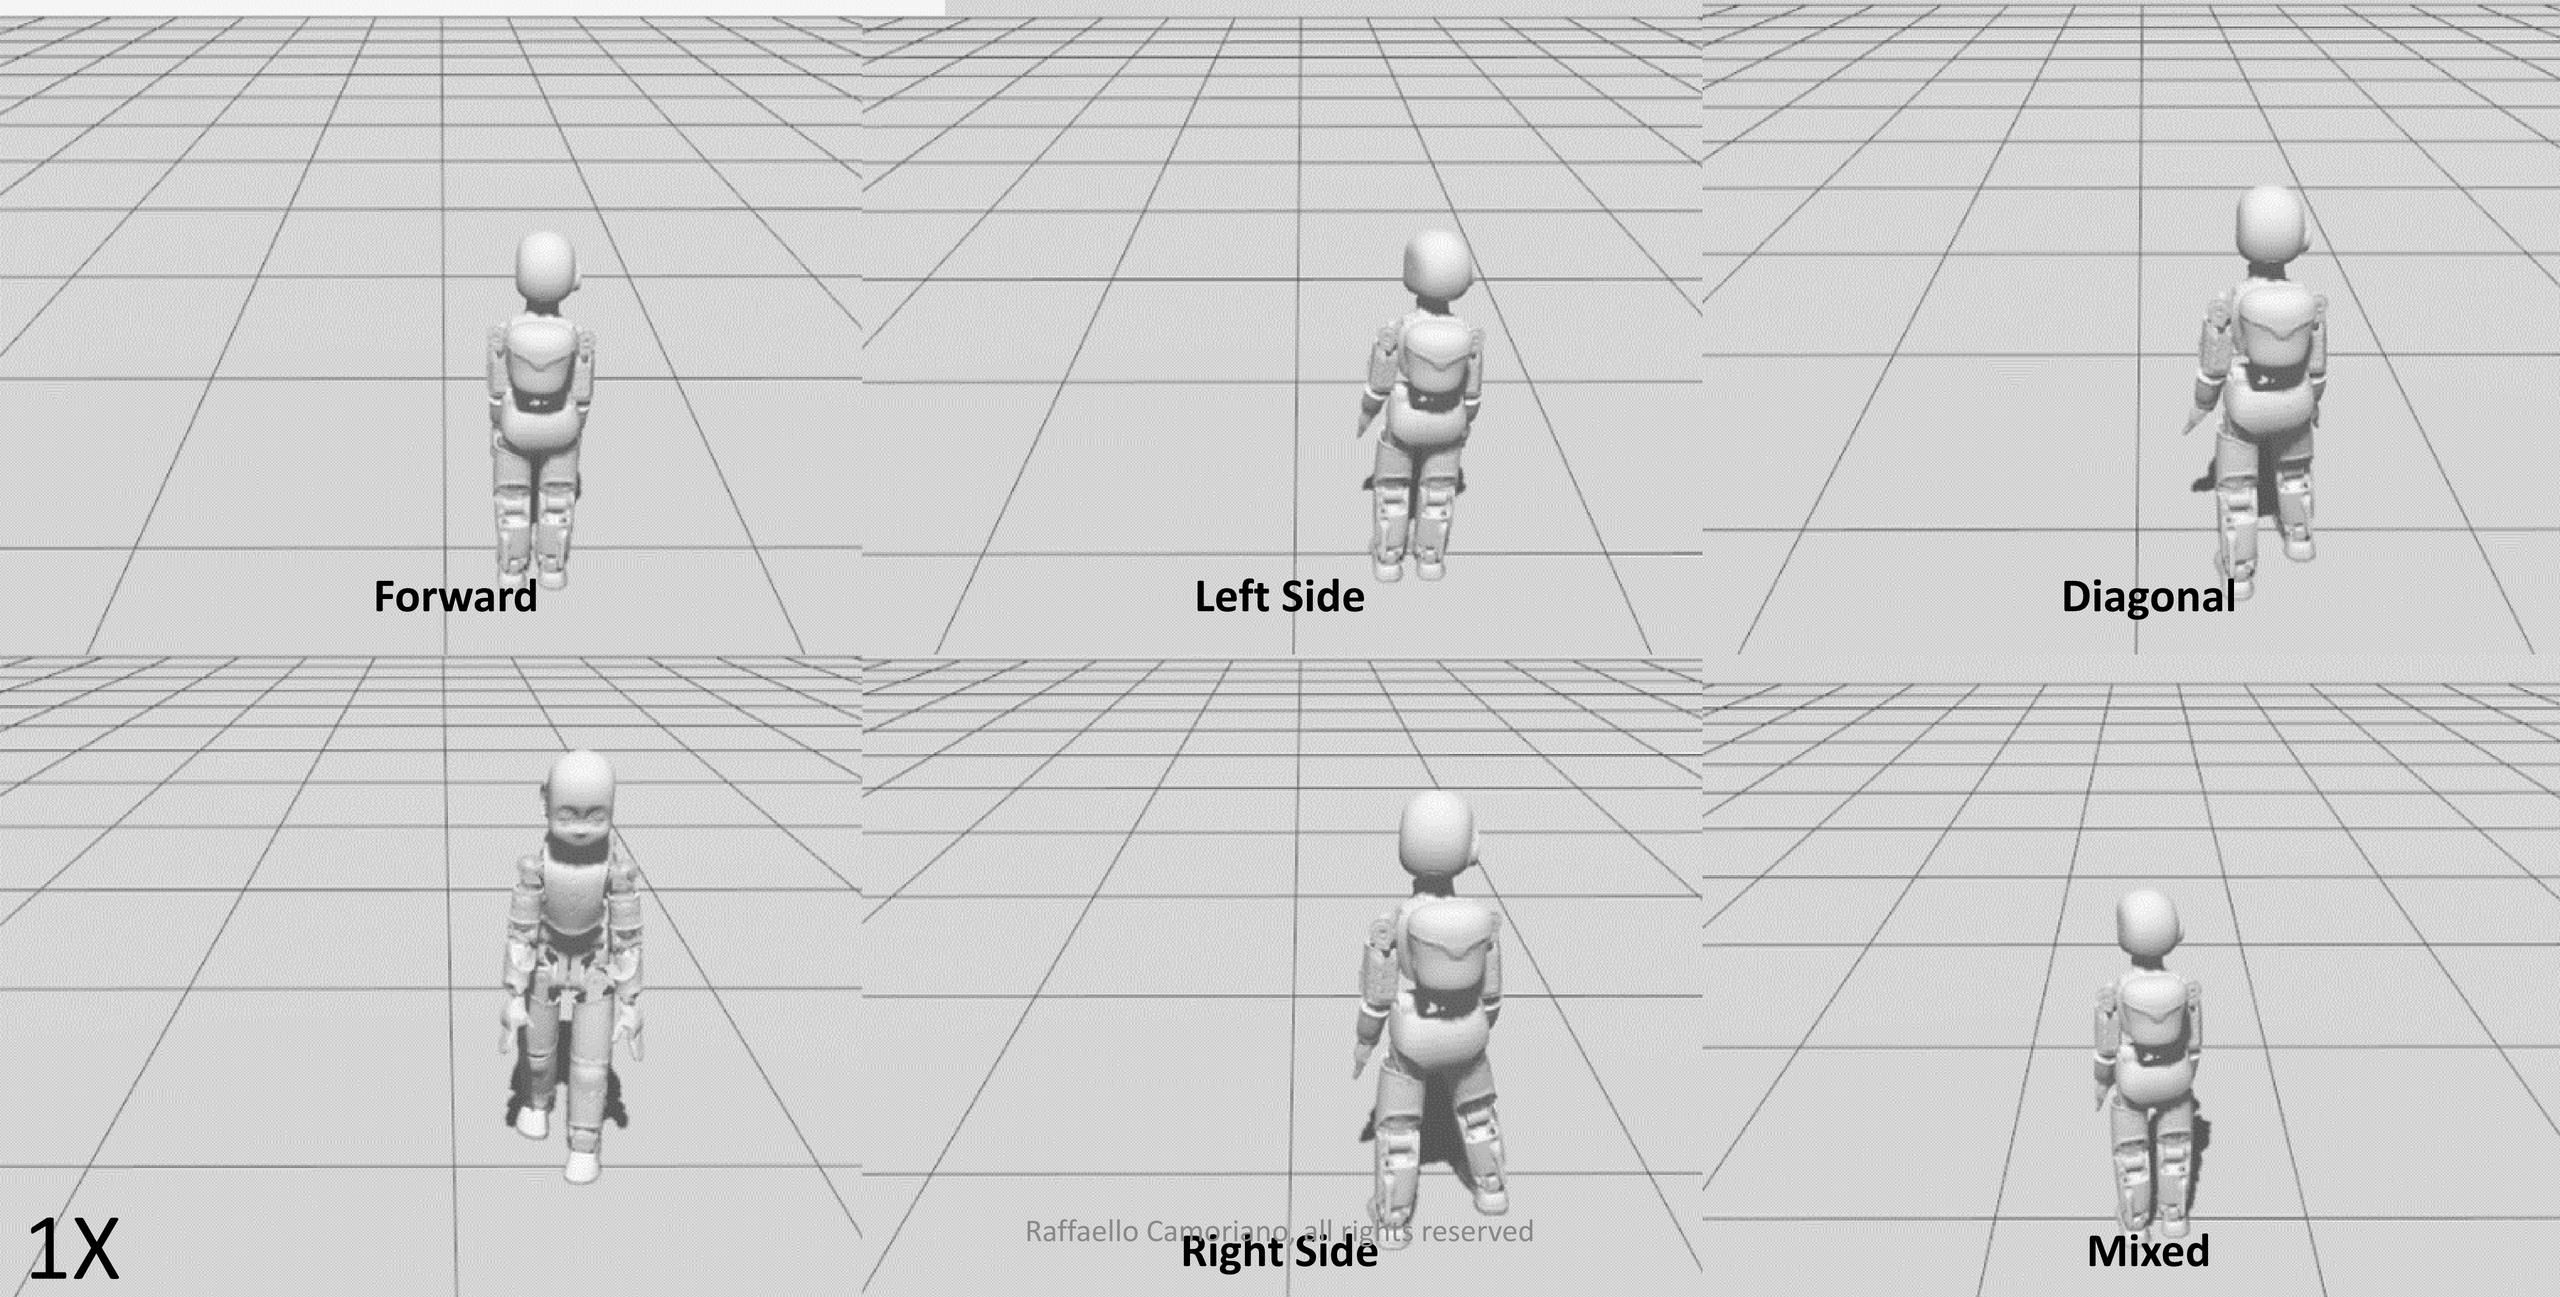

Trajectory Generation

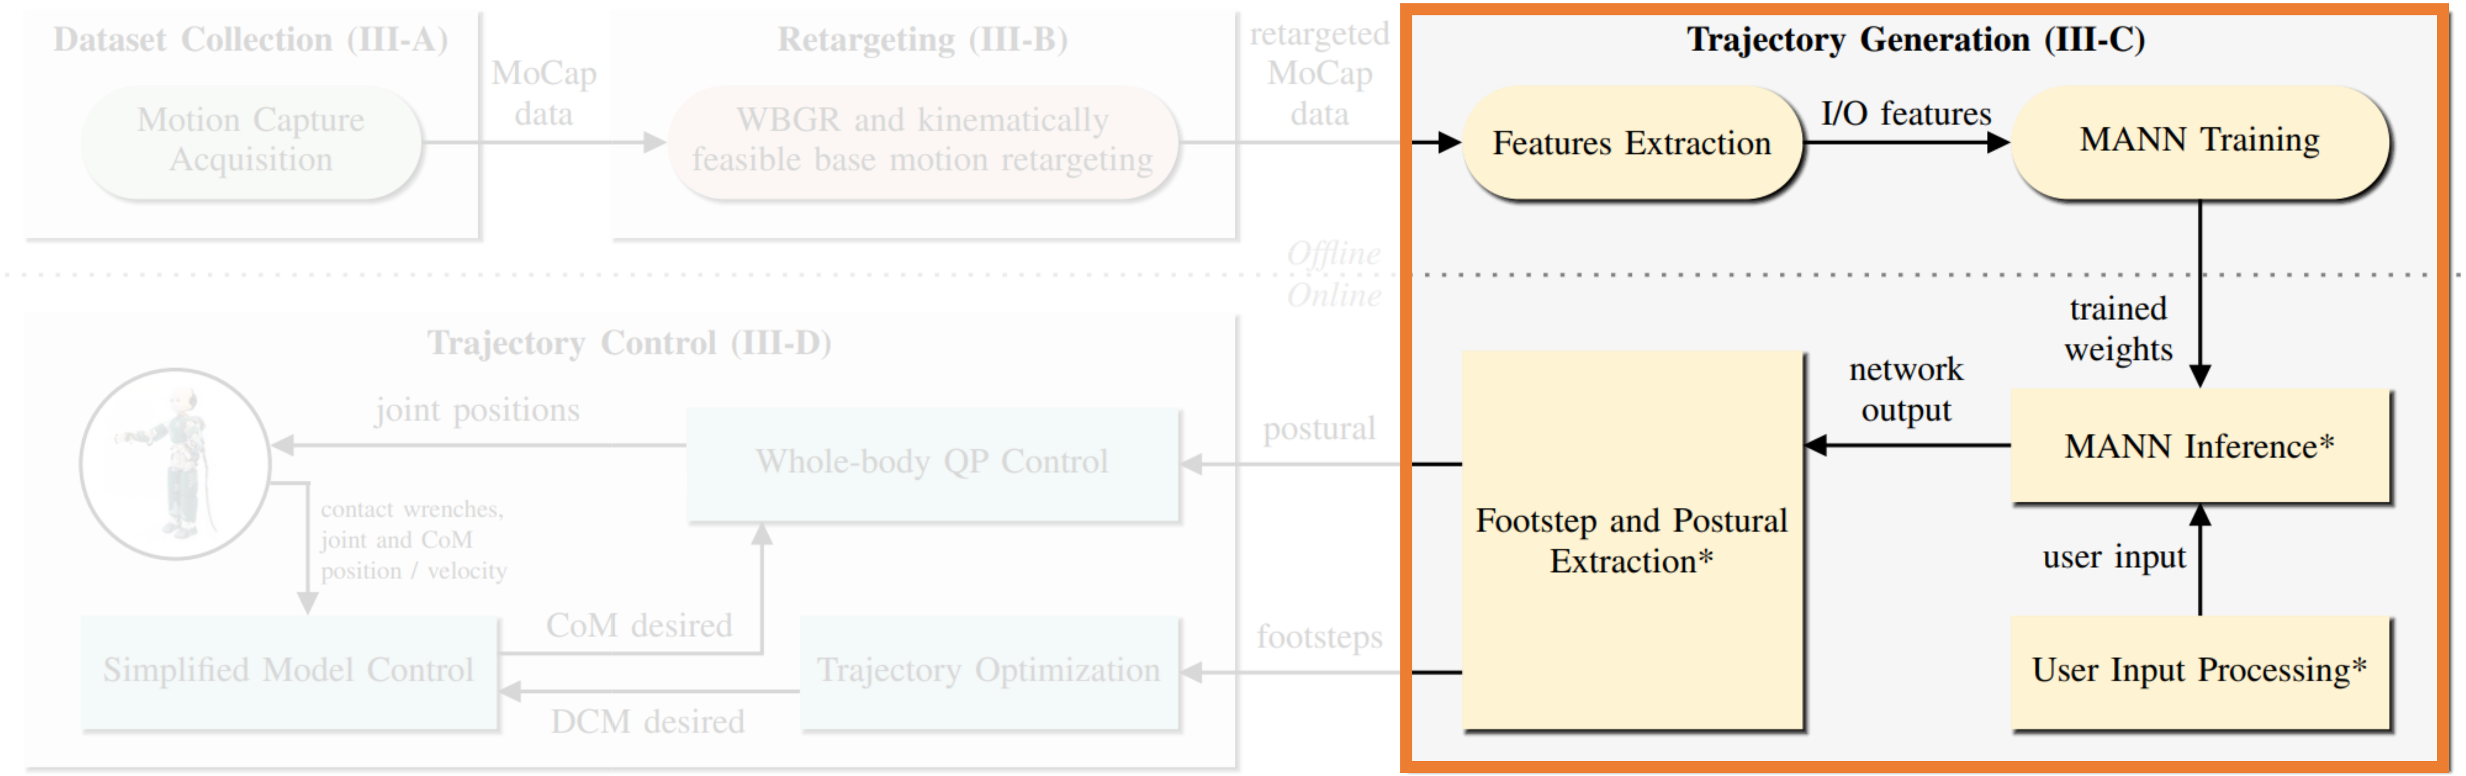

Network Structure

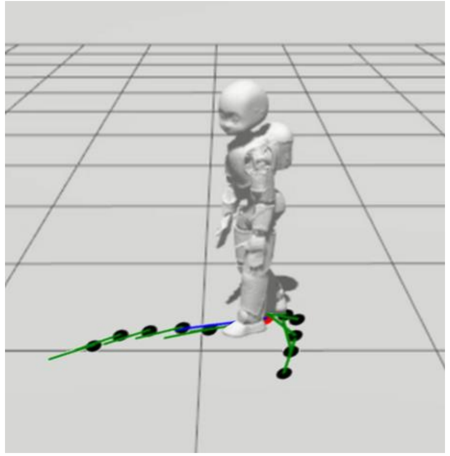

$x_i \in \mathbb{R}^{137}$

- ground trajectory data at t_i
- state of the character at t_{i-1}

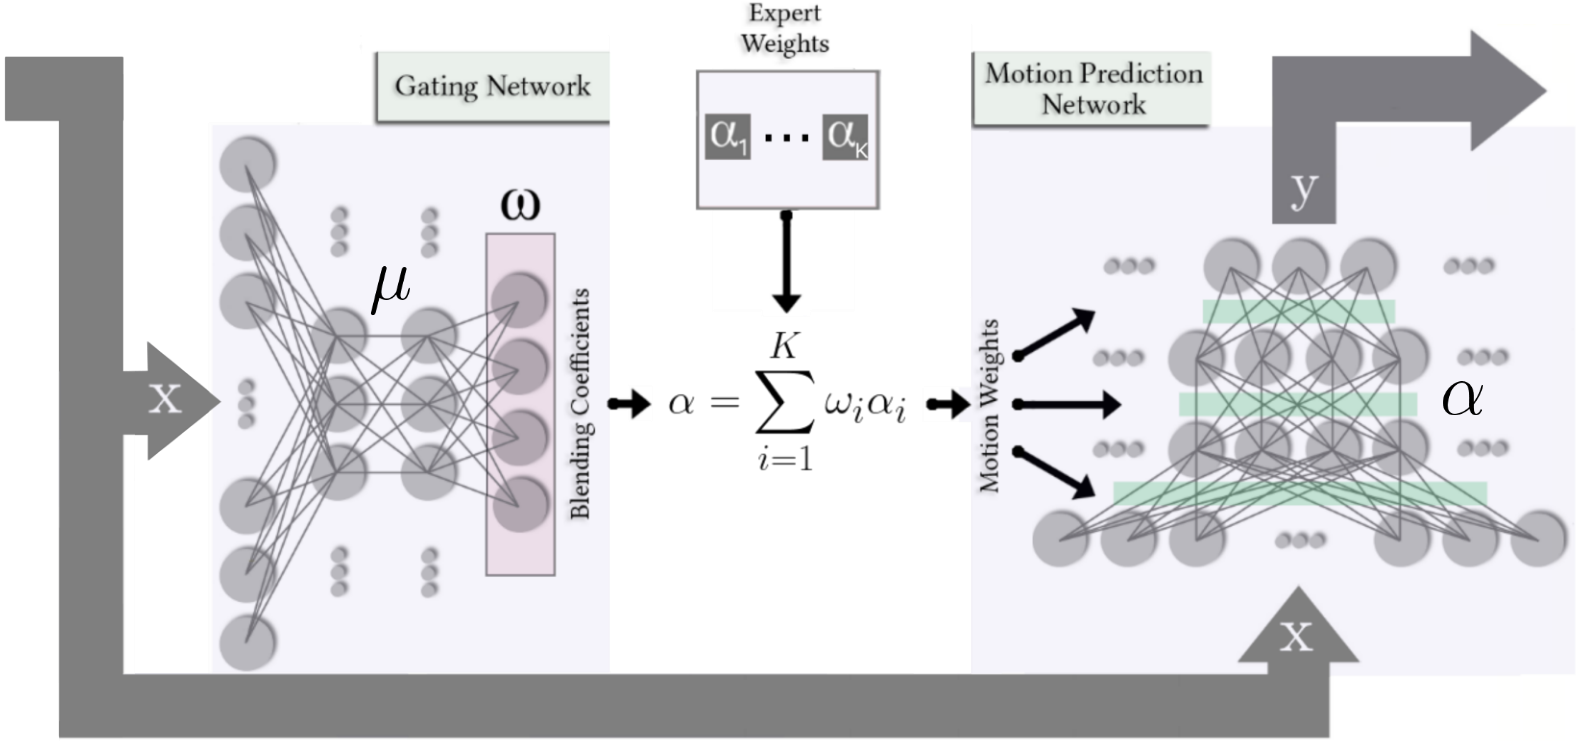

Adapted from Zhang et al., “Mode-Adaptive Neural Networks for Quadruped Motion Control”, *ACM Transactions on Graphics*, 2018

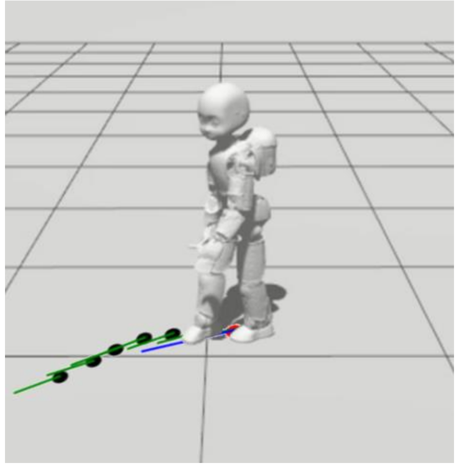

$y_i \in \mathbb{R}^{103}$

- future ground trajectory data at t_{i+1}
- state of the character at t_i
- root transformation from t_{i-1} to t_i

User Input Processing and MANN Inference

Desired Motion Direction

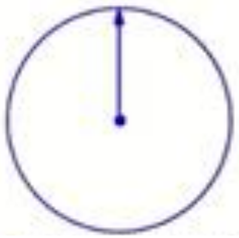

Desired Facing Direction

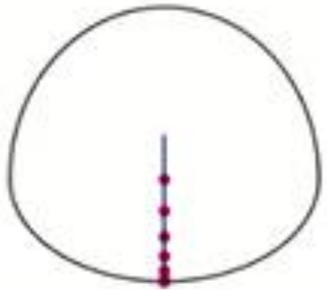

Desired future ground base trajectory

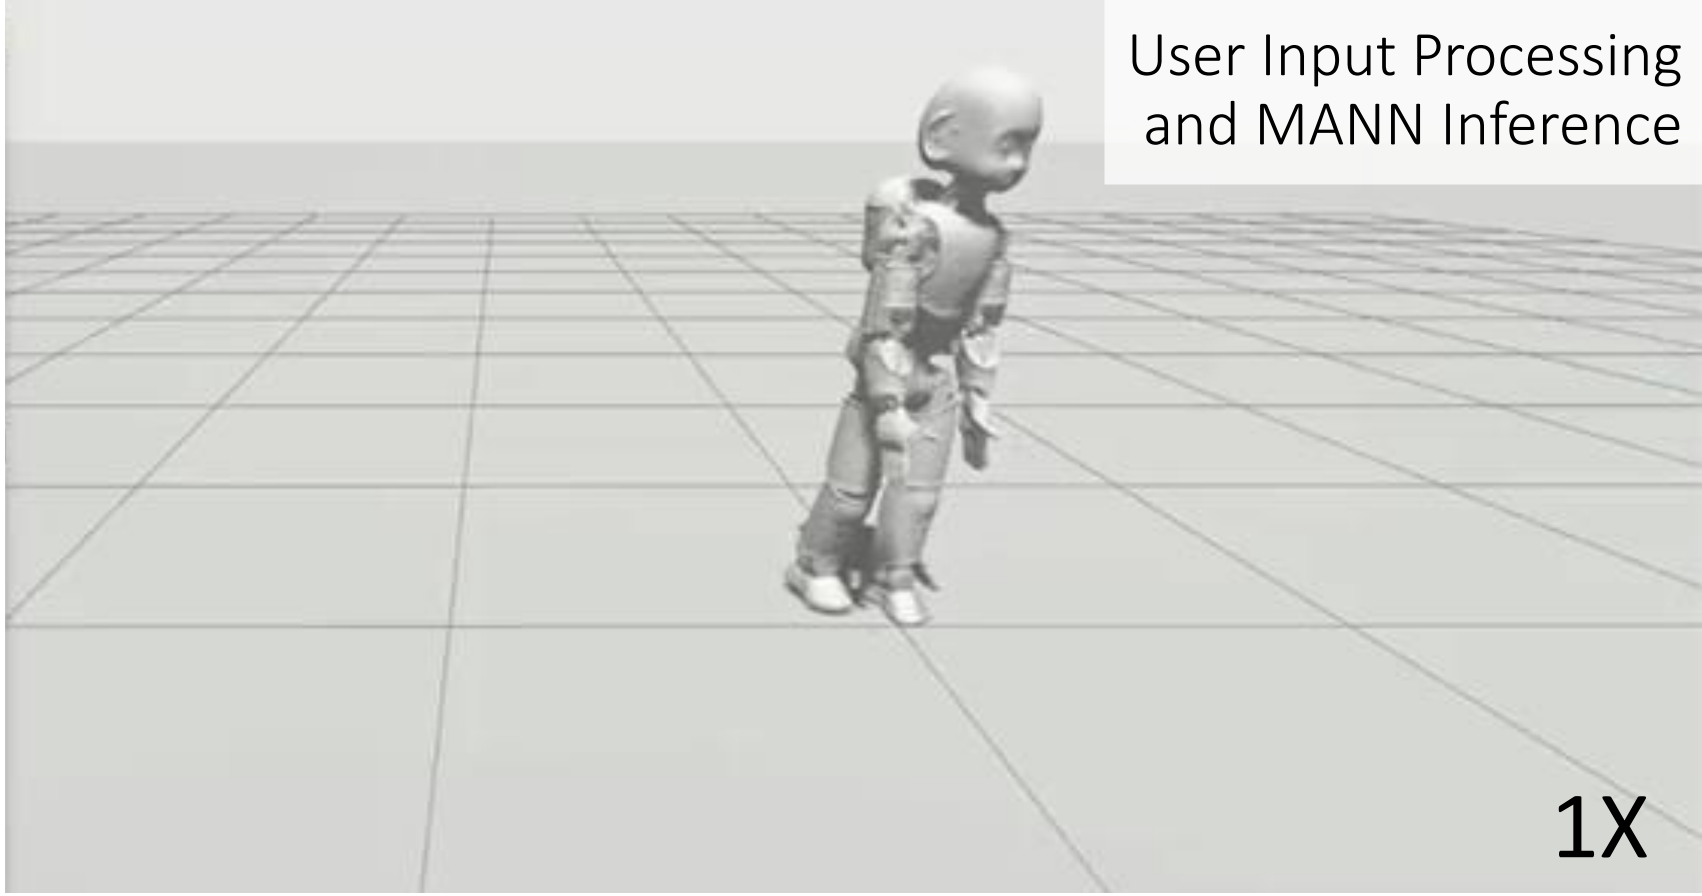

MANN Blending Coefficients

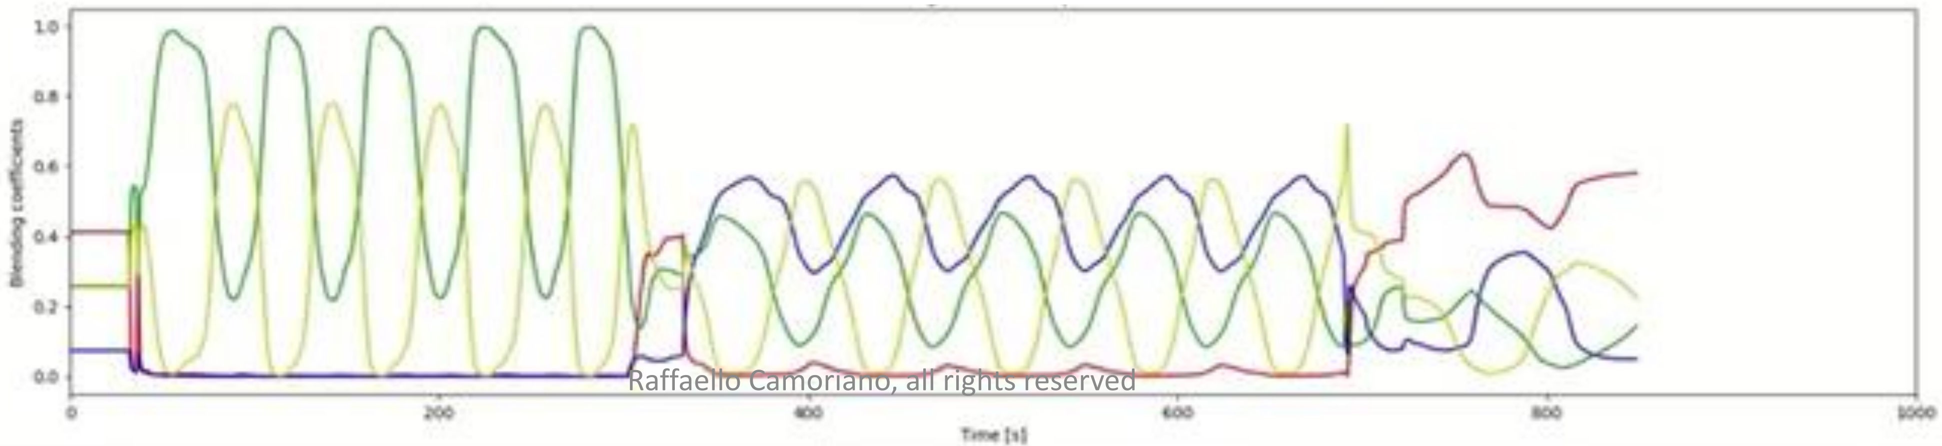

Trajectory Control

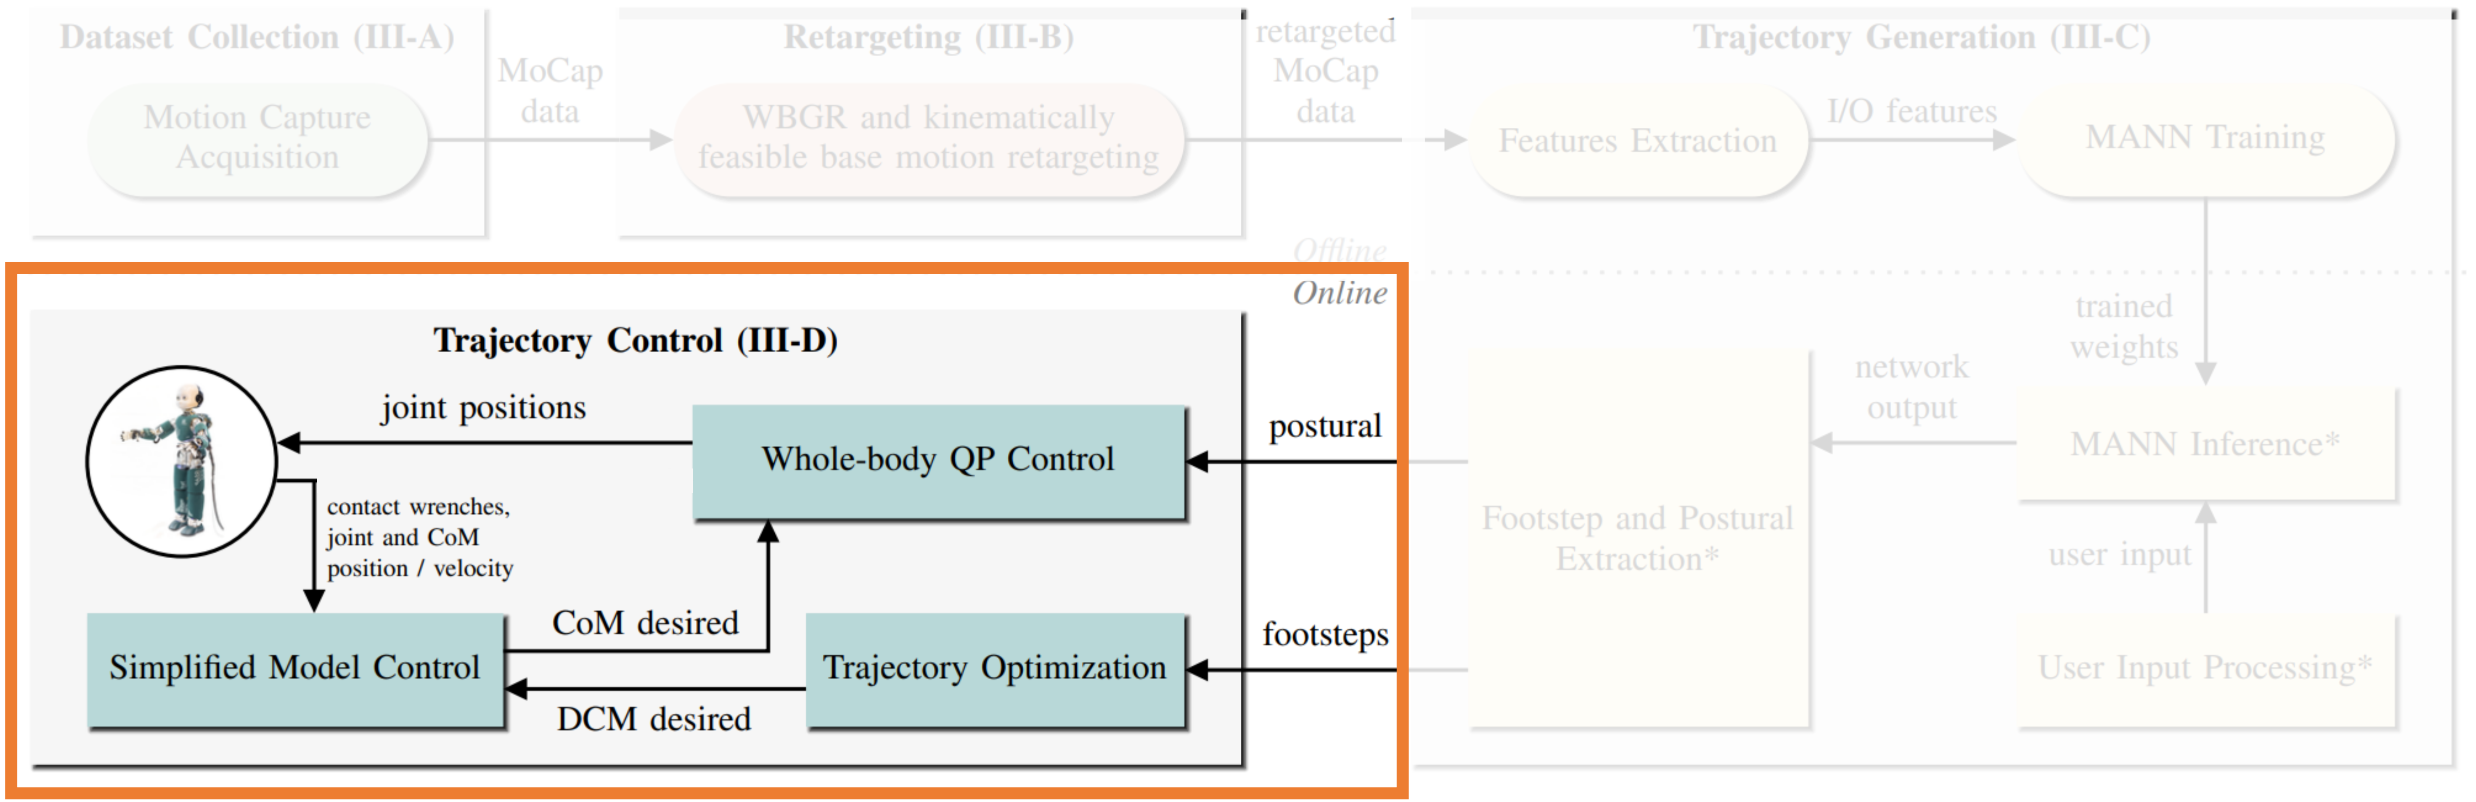

Trajectory Control

Forward walking

Left-side walking

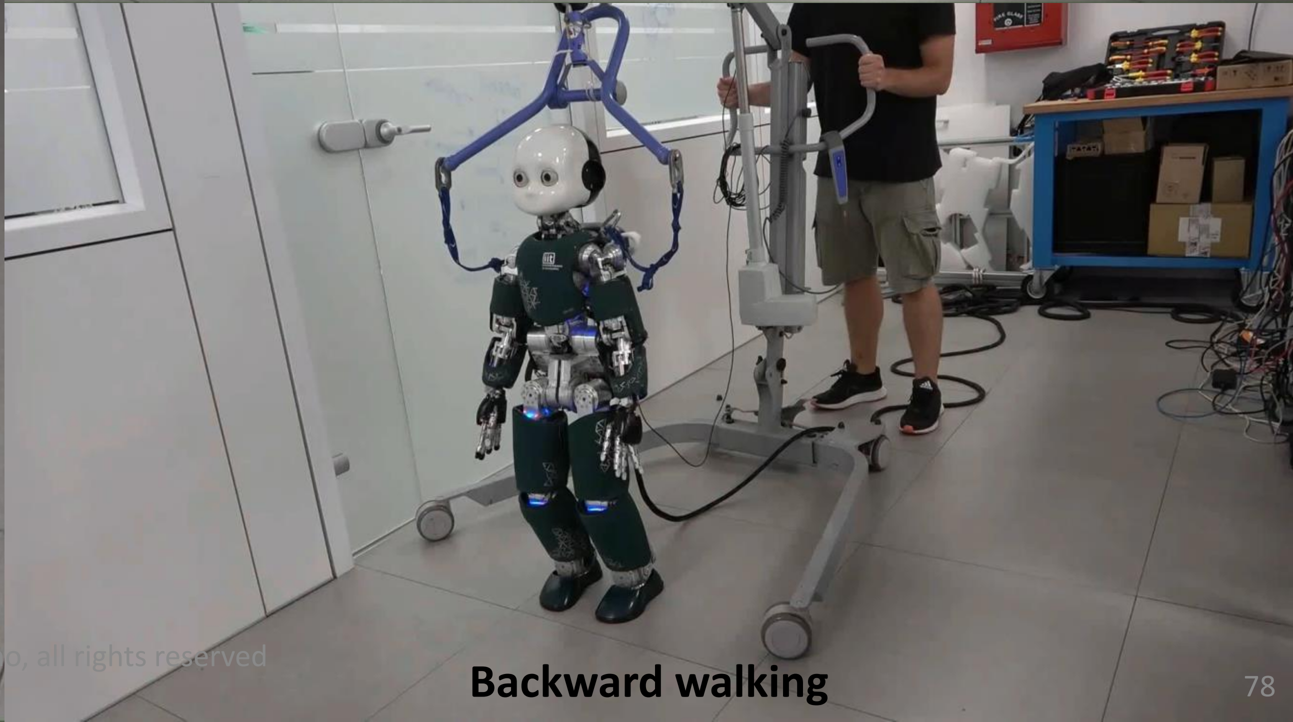

Backward walking

Real-robot Trajectory Control: Mixed Walking Sequence

1X

Extension: Physics-Informed Motion Generation

- Main ADHERENT limitations: drift and feet stability
- **Idea:** Generate whole-body motion of a humanoid incorporating **kinematically feasible contacts** via a physics-informed loss
- **Control-informed blending of user input** with network predictions to reduce drift

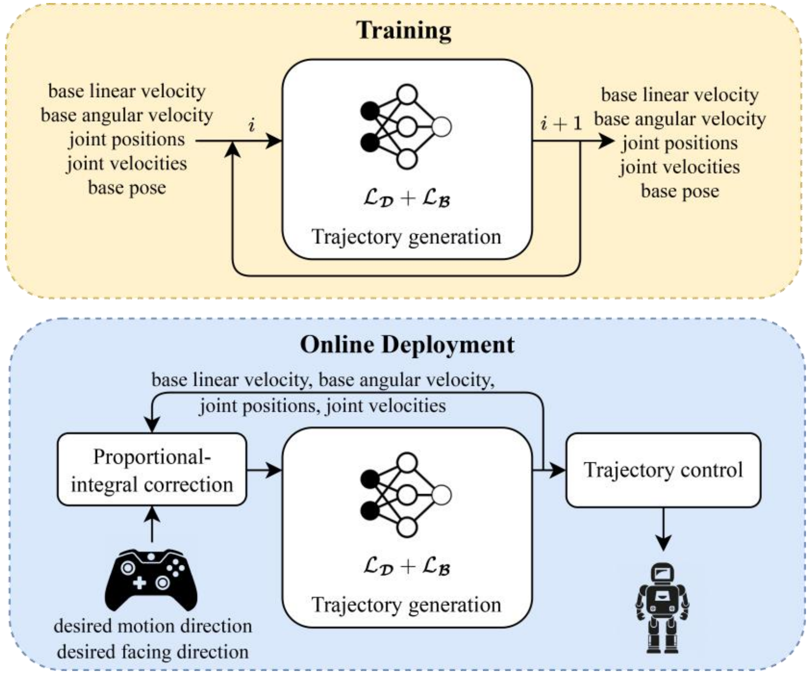

Physics-informed loss

We use known physics to design a loss component to prevent foot sliding.

Physics law

$$\mathbf{v}_{SF} = \mathbf{J}_{SF}\boldsymbol{\nu} = \begin{bmatrix}\mathbf{J}^{\mathcal{B}}_{SF} & \mathbf{J}^{\dot{\mathbf{s}}}_{SF}\end{bmatrix}\begin{bmatrix}\mathbf{v}_{\mathcal{B}} \\ \dot{\mathbf{s}}\end{bmatrix} = 0$$

Physics-informed loss component

$$\mathcal{L}_{\mathcal{B}}(\mathbf{x};\theta) = \left\| \begin{bmatrix}{}^{\mathcal{B}}\hat{\mathbf{v}}_{\mathcal{I},\mathcal{B}} \\ {}^{\mathcal{B}}\hat{\boldsymbol{\omega}}_{\mathcal{I},\mathcal{B}}\end{bmatrix}(\mathbf{x};\theta) + \left[\alpha(\mathbf{J}^{\mathcal{B}}_{LF})^{-1}\mathbf{J}^{\dot{\mathbf{s}}}_{LF} + (1-\alpha)(\mathbf{J}^{\mathcal{B}}_{RF})^{-1}\mathbf{J}^{\dot{\mathbf{s}}}_{RF}\right]\hat{\dot{\mathbf{s}}}(\mathbf{x};\theta) \right\|^2$$

Overall loss

$$\mathcal{L}(\mathbf{x};\theta) = \mathcal{L}_{\mathcal{D}}(\mathbf{x};\theta) + w\mathcal{L}_{\mathcal{B}}(\mathbf{x};\theta)$$

D'Elia, Evelyn, et al. "Stabilizing Humanoid Robot Trajectory Generation via Physics-Informed Learning and Control-Informed Steering." *IEEE/RSJ IROS* 2025

Extension: Physics-Informed Motion Generation

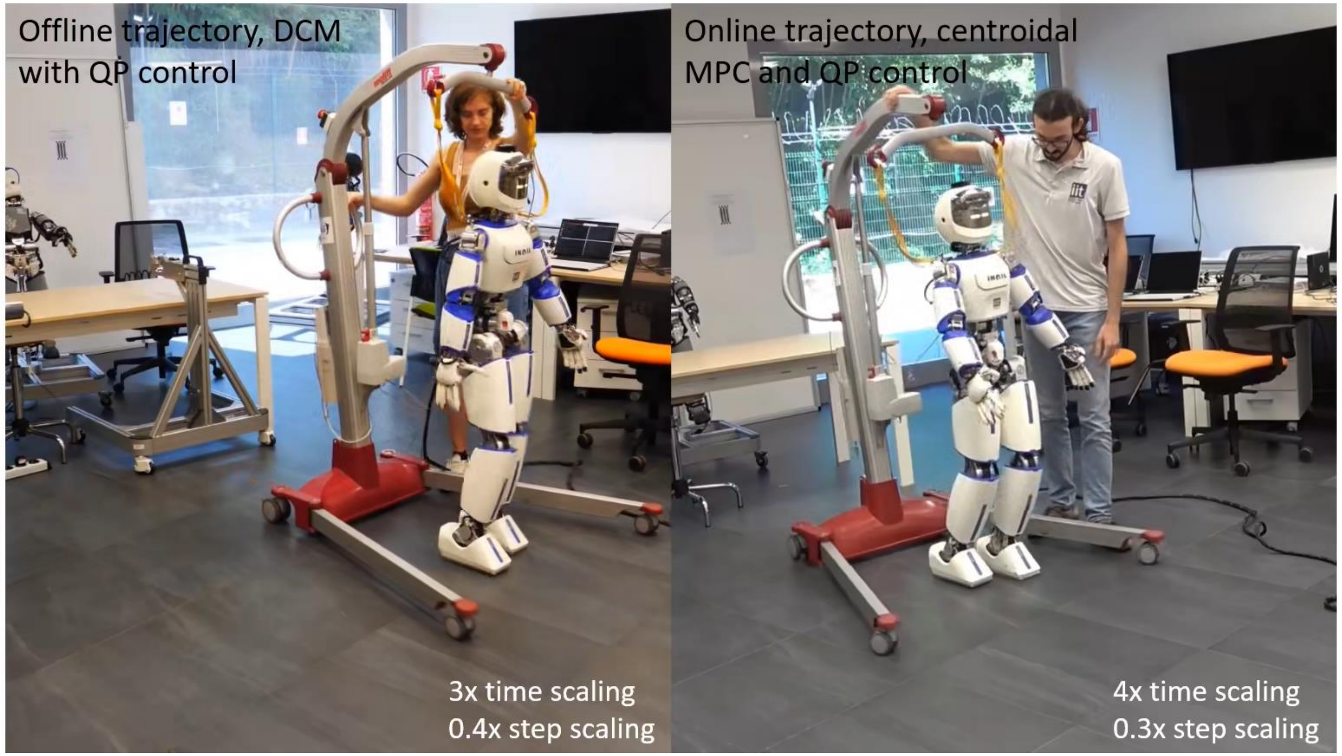

Real-world deployment & modulatity demonstration

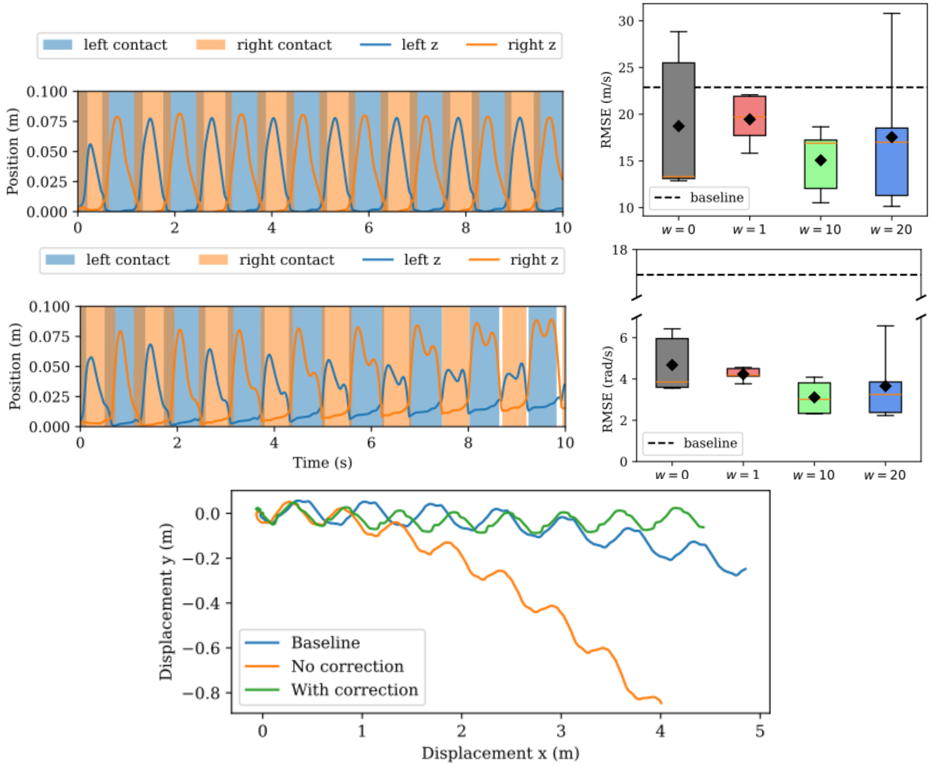

Ablations

Conclusions

- Structured Learning for robotics
  - Inverse Kinematics learning
  - Imitation learning on manifolds
  - 3D Learning for Object-centric Planning
- Incremental model learning and prosthetic myocontrol
- Learning humanoid locomotion
  - RL-based
  - Supervised imitation learning from MoCap

Contacts

Email: raffaello.camoriano@polito.it
Web: https://www.polito.it/en/staff?p=raffaello.camoriano
BlueSky: @raf_camo

Future Work

- Structured learning with dynamically changing environment and embodiment structures

- Learning structural properties from data and interaction

- 3D Learning for Grasping and Manipulation
    - Efficient task-driven point-cloud sampling for grasp proposal
    - Articulated object manipulation
    - Transferring across different grippers and embodiments

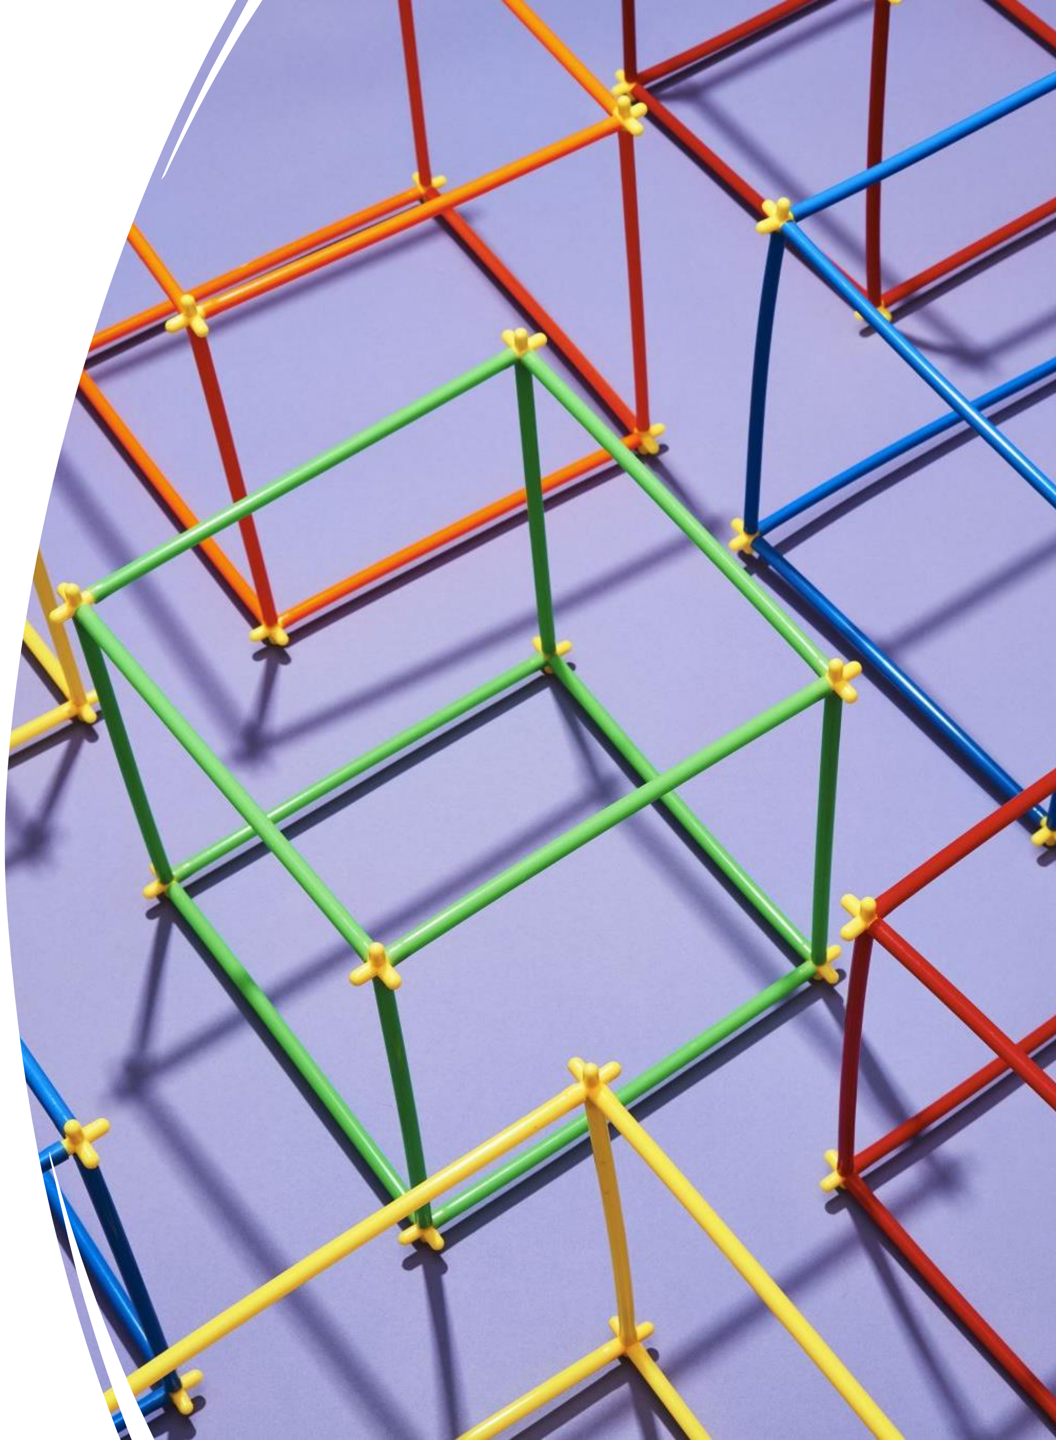

Acknowledgments

This study was carried out within the FAIR - Future Artificial Intelligence Research and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1555 11/10/2022, PE00000013). This manuscript reflects only the authors’ views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

Additional References

- A. Rudi, R. C., and L. Rosasco. ``Less is More: Nyström Computational Regularization''. NIPS, 2015.
- R. C.*, T. Angles*, A. Rudi, and L. Rosasco, "NYTRO: When Subsampling Meets Early Stopping," AISTATS, 2016.
- A. Rudi, R. C., and L. Rosasco. ``Generalization Properties of Learning with Random Features''. NIPS Workshop on Adaptive and Scalable Nonparametric Methods in ML, 2016.
- R. C.*, G. Pasquale*, C. Ciliberto, L. Natale, L. Rosasco, and G. Metta. ``Incremental robot learning of new objects with fixed update time''. ICRA, 2017
- J. Lin, R. C., and L. Rosasco. ``Generalization Properties and Implicit Regularization for Multiple Passes SGM''. ICML, 2016.
- R. C., S. Traversaro, L. Rosasco, G. Metta, and F. Nori. ``Incremental Semiparametric Inverse Dynamics Learning''. ICRA, 2016
- D. Romeres, M. Zorzi, R. C., A. Chiuso. ``Online semi-parametric learning for inverse dynamics modeling''. IEEE CDC, 2016
- Romeres, Diego, et al. "Derivative-free online learning of inverse dynamics models." IEEE Transactions on Control Systems Technology 28.3 (2019): 816-830
- Milios, Dimitrios, R. C. et al. "Dirichlet-based Gaussian processes for large-scale calibrated classification." *NeurIPS,* 2018
- Duan, A, R. C. , Ferigo, D et al. (2019) Learning to Sequence Multiple Tasks with Competing Constraints. IROS 2019
- Duan, A., R. C., et al. "Learning to Avoid Obstacles With Minimal Intervention Control." Frontiers in Robotics and AI 7 (2020)
- Duan, A., R. C., et al. "Constrained DMPs for feasible skill learning on humanoid robots." IEEE Humanoids, 2018.