Learning for Autonomy

From Policies to Operators to Distributions

Adam Thorpe

TEXAS The University of Texas at Austin

ODEN INSTITUTE FOR COMPUTATIONAL ENGINEERING & SCIENCES

CENTER FOR auTonomy

Learning for Autonomy

Learning-Based Algorithms

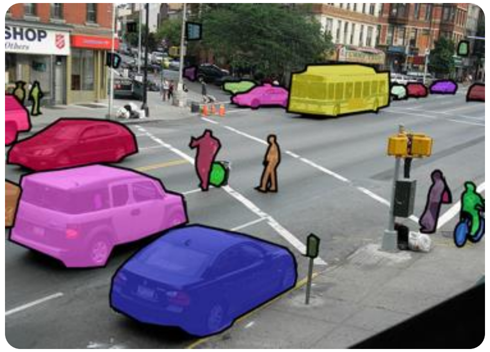

Autonomous Systems

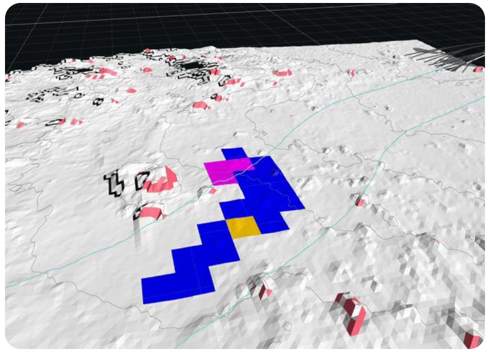

Computational Methods

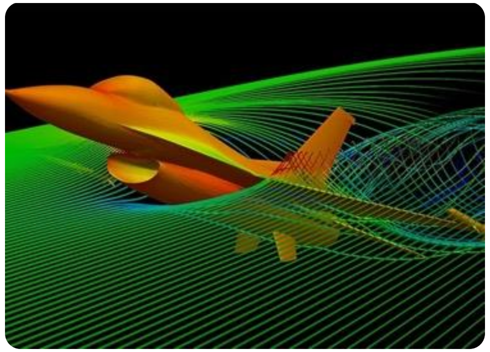

Does scale lead to generalization?

Open X-Embodiment

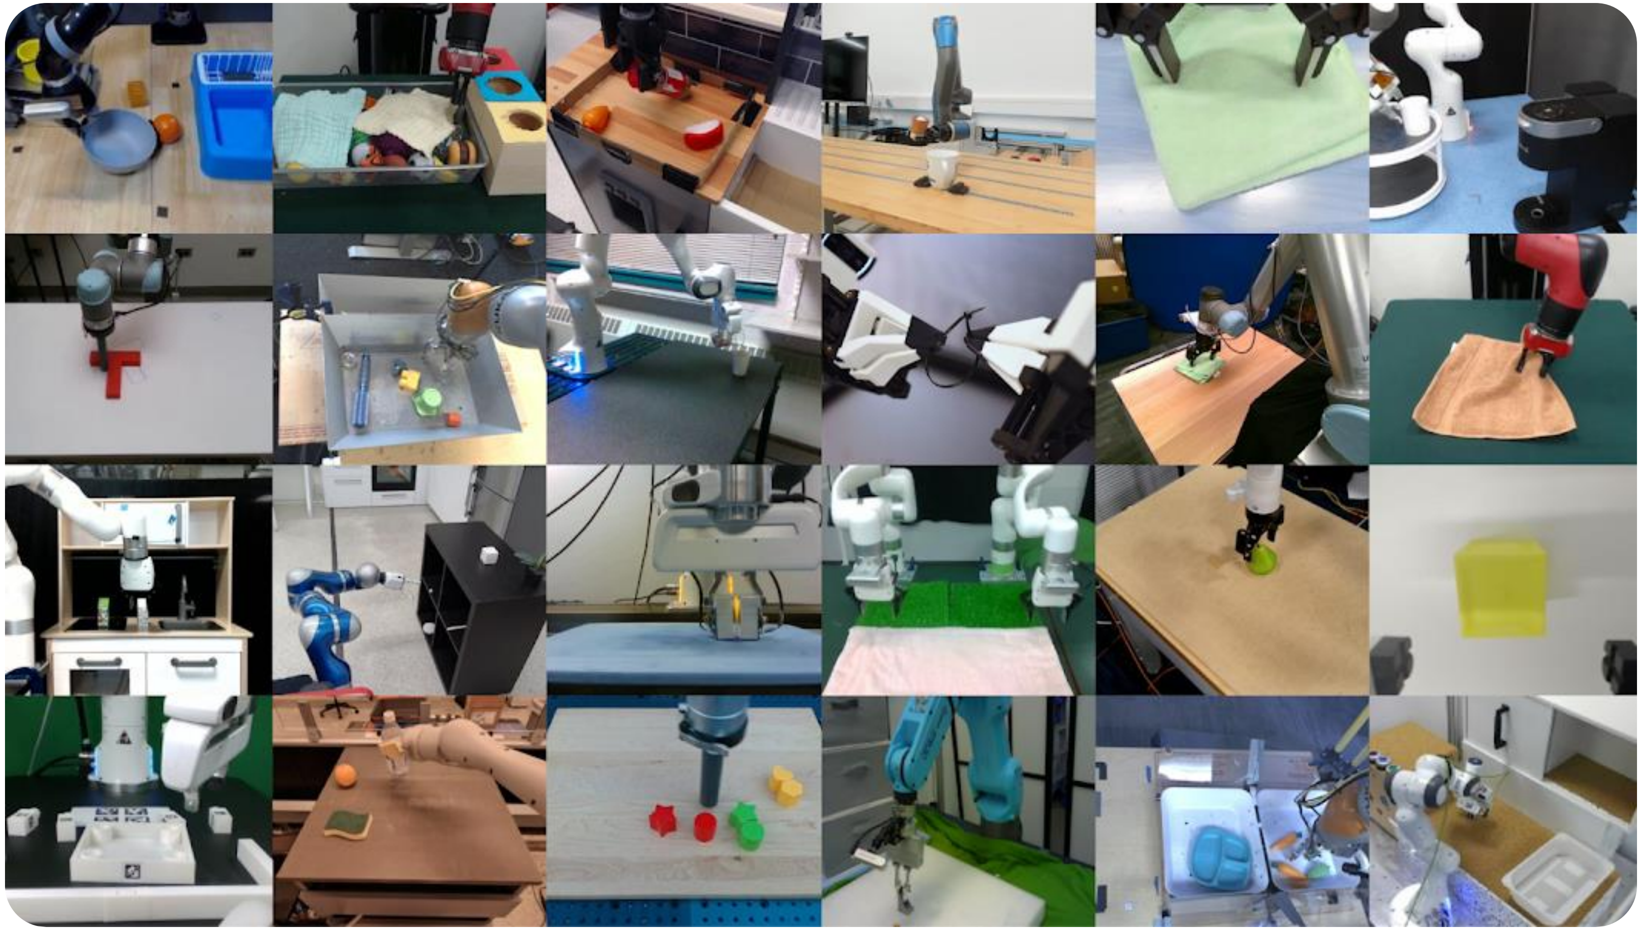

Edge cases

Autonomous systems will encounter scenarios outside the training data.

Unseen environments

Autonomy needs to be able to adapt to unseen environments.

Practical autonomy will need to operate safely with humans

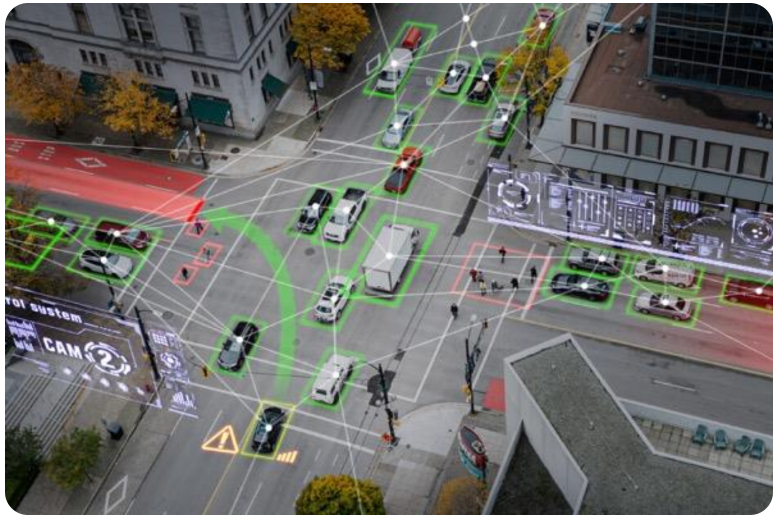

We need to design autonomous systems to avoid **misuse**, **disuse**, and **abuse**.

Parasuraman, R., & Riley, V. (1997). Humans and Automation: Use, Misuse, Disuse, Abuse. Human Factors

Rare events

Autonomy must continue to make informed decisions in the face of uncertainty.

Learning-based algorithms can be fragile

Training Data

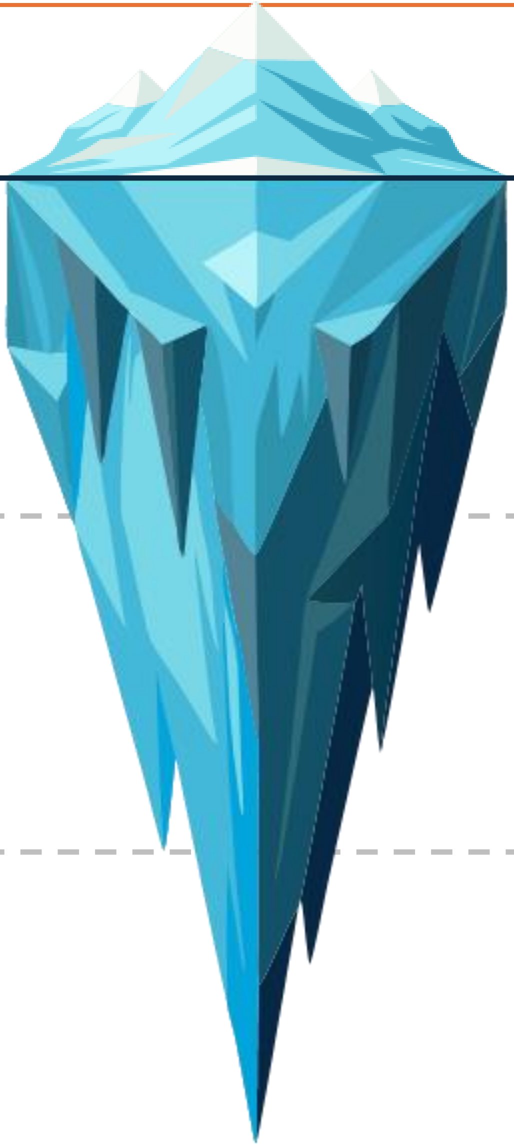

Edge Cases

Unseen Environments

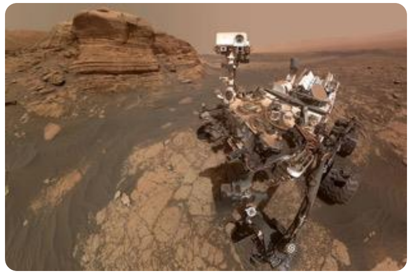

Rare Events

Unknowns

We need algorithms that can **adapt** & **transfer**…

…across domains

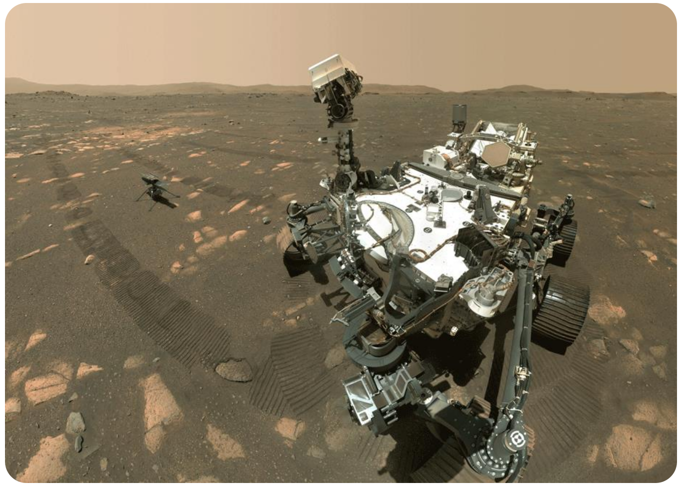

…across tasks

…across platforms

What do I mean by adaptation and transfer?

Adaptation updating or refining learned models using new data

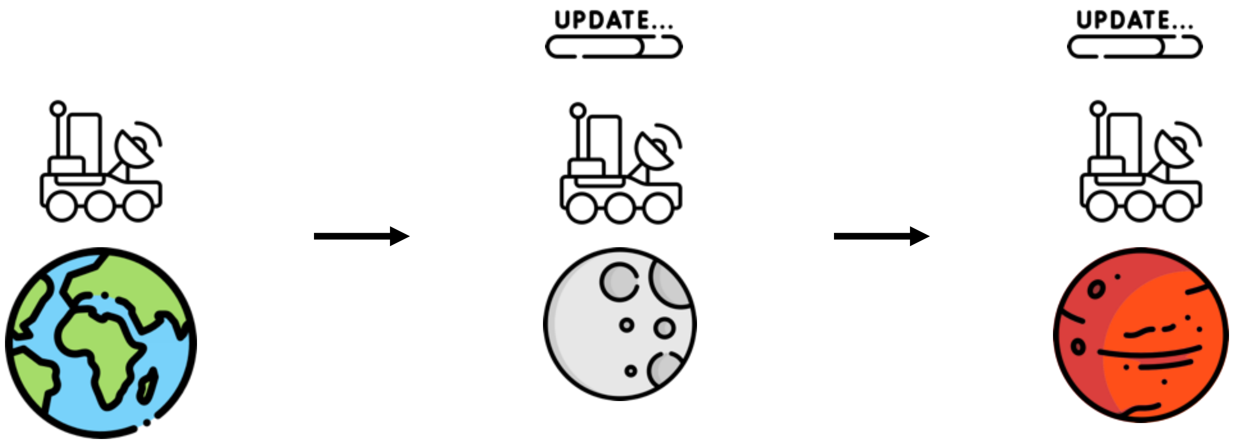

Transfer leveraging knowledge from diverse sources

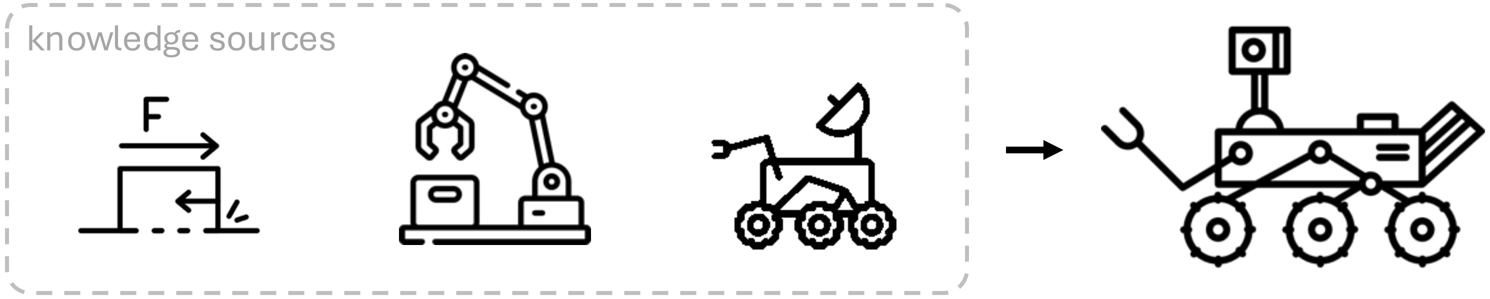

Moving beyond black-box learning

Incorporating known physics and mathematical structure

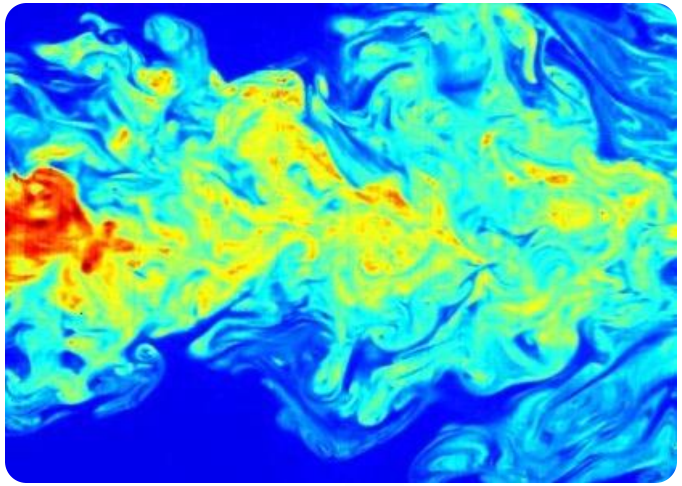

knowledge

data

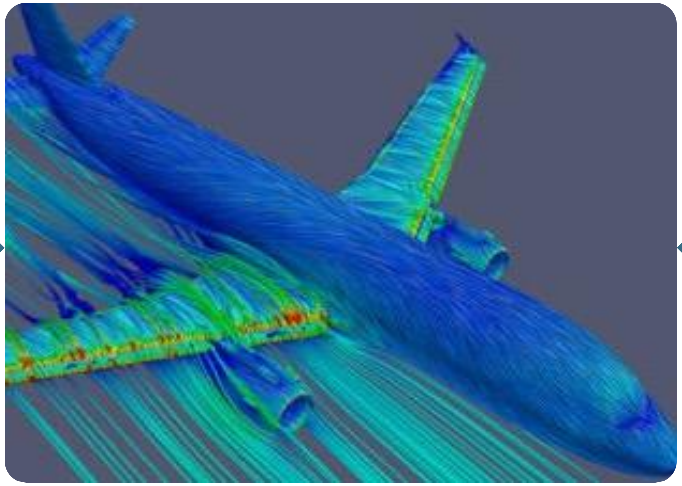

My research: learning for autonomy

Learning Structured Representations

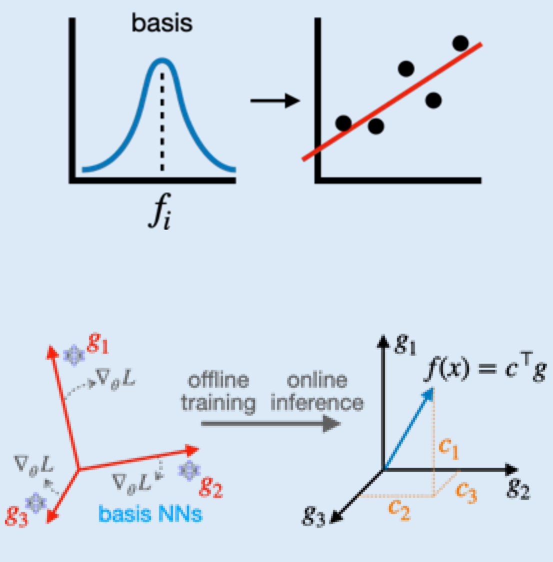

Function Encoders

Adaptable Autonomy & Inductive Transfer

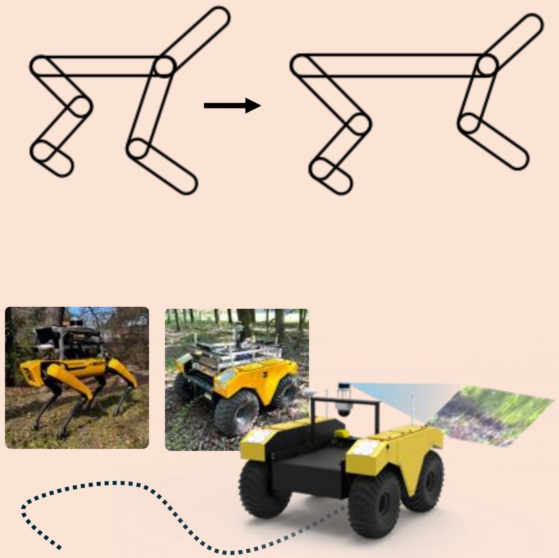

Zero-Shot Modeling & Control

Neural Operator Learning

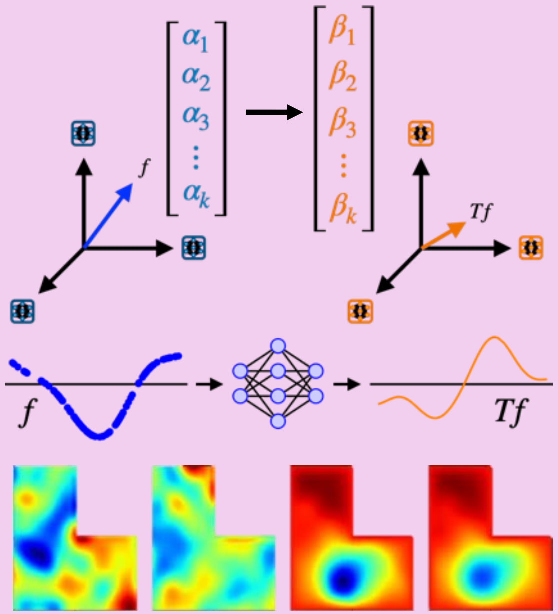

Basis to Basis Operators

Human-Autonomy Interaction

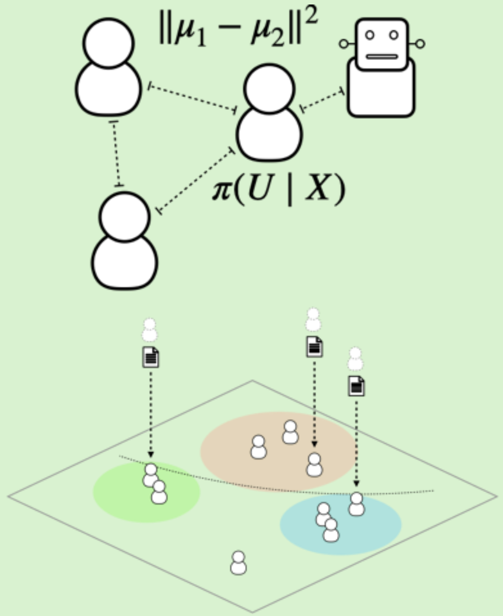

Characterizing Heterogeneity

Preview of main results

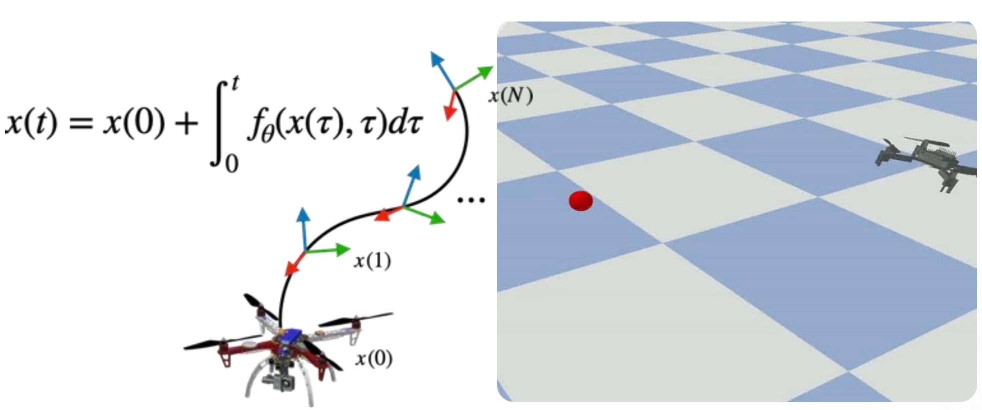

Zero-Shot System Identification & Control

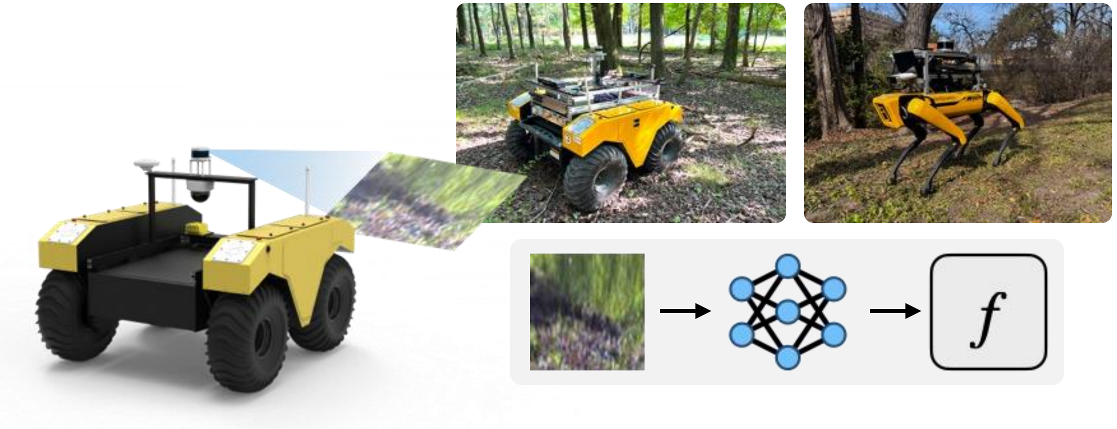

Vision to Dynamics

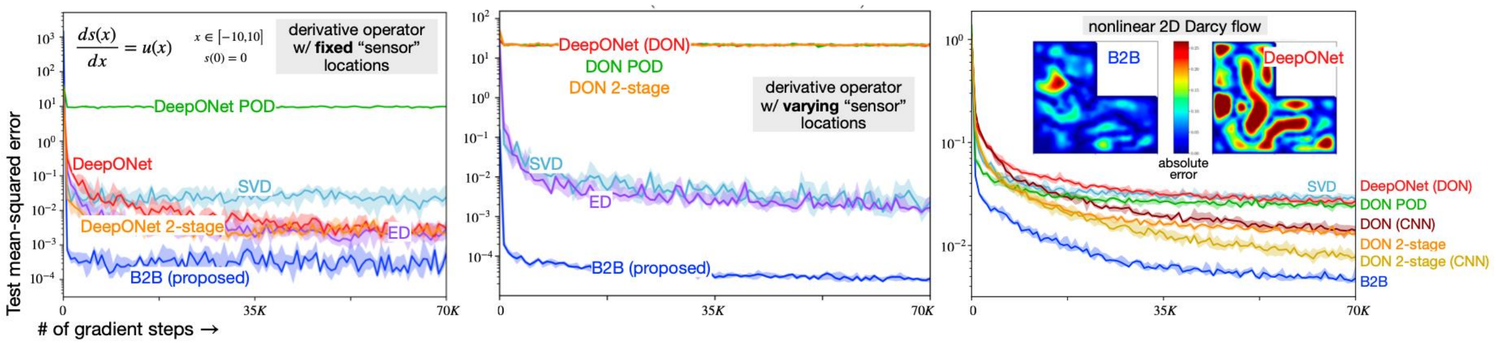

Basis to Basis Operator Learning

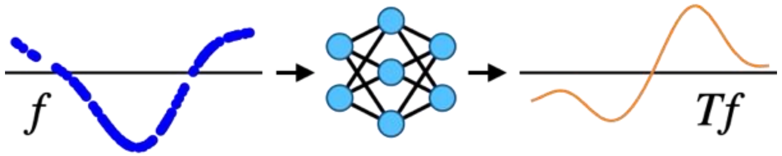

Existing transfer approaches

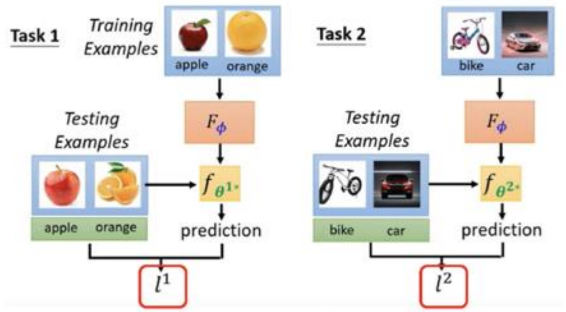

Meta-Learning

Chelsea Finn, Pieter Abbeel, Sergey Levine. (2017). Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks.

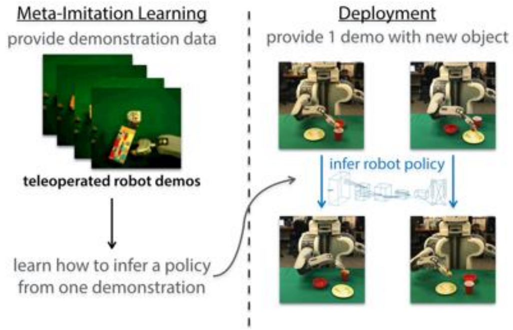

Imitation Learning

O'Neill, A., Rehman, A., Maddukuri, A., Gupta, A., Padalkar, A., Lee, A., ... & Chen, M. (2024). Open X-Embodiment: Robotic Learning Datasets and RT-X Models

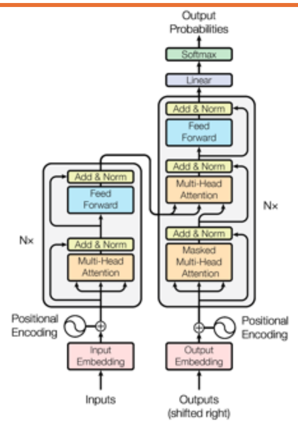

Transformers

Ashish Vaswani, et. al. (2017). Attention is All you Need.

D. Celestini, D. Gammelli, T. Guffanti, S. D'Amico, E. Capello and M. Pavone. (2024). Transformer-Based Model Predictive Control: Trajectory Optimization via Sequence Modeling

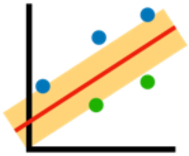
SVMs

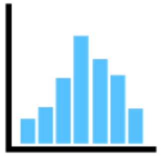
Density Estimation

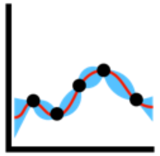
GPs


Kernel PCA


Ridge Regression

Hilbert Space Representations

✓ guarantees ✓ interpretable ✓ efficient

Prior work: kernel-based stochastic optimal control

Intractable Optimal Control Problem

(unknown dynamics/disturbance)

kernel-based reformulation

Prior Physics Knowledge

Data $\xrightarrow{\approx}$ Kernel Representation → Linear Program

approximate solution

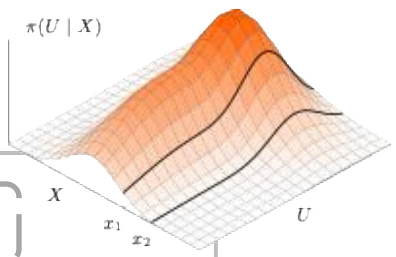

we want the best of both worlds

uncertain mass
unknown disturbance
unknown dynamics

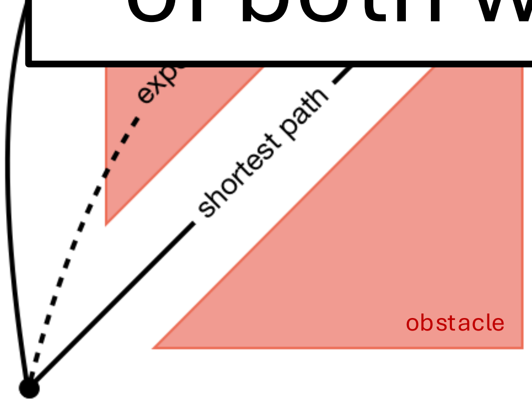

Approximate Policy:

scales with the amount of data

$$\sum_{i=1}^{m} \alpha_i(x) k(u_i, \cdot)$$

Kernel Methods

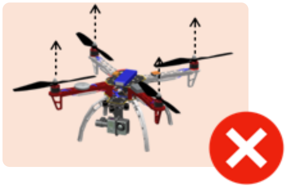

data complexity

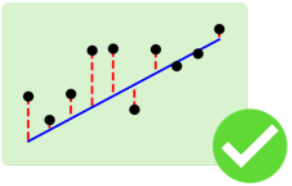

least squares

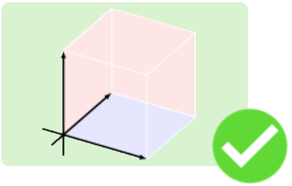

structure

Neural Networks

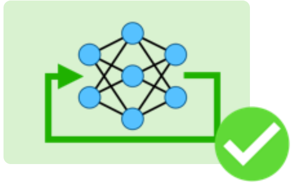

offline training

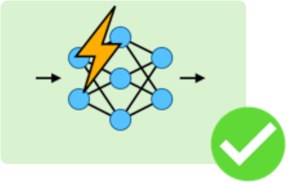

fast inference

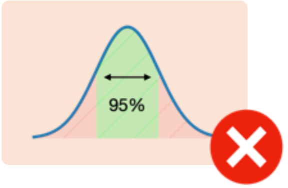

guarantees

Thorpe, A., Lew, T., Oishi, M. & Pavone, M. (2022). Data-Driven Chance Constrained Control using Kernel Distribution Embeddings. L4DC

Function encoders: combining neural networks and Hilbert spaces

Problem: How can we represent Hilbert spaces?

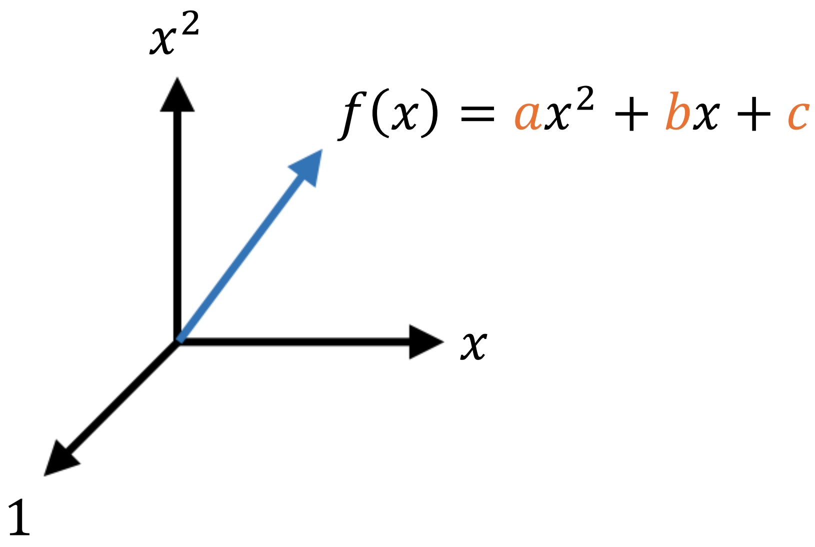

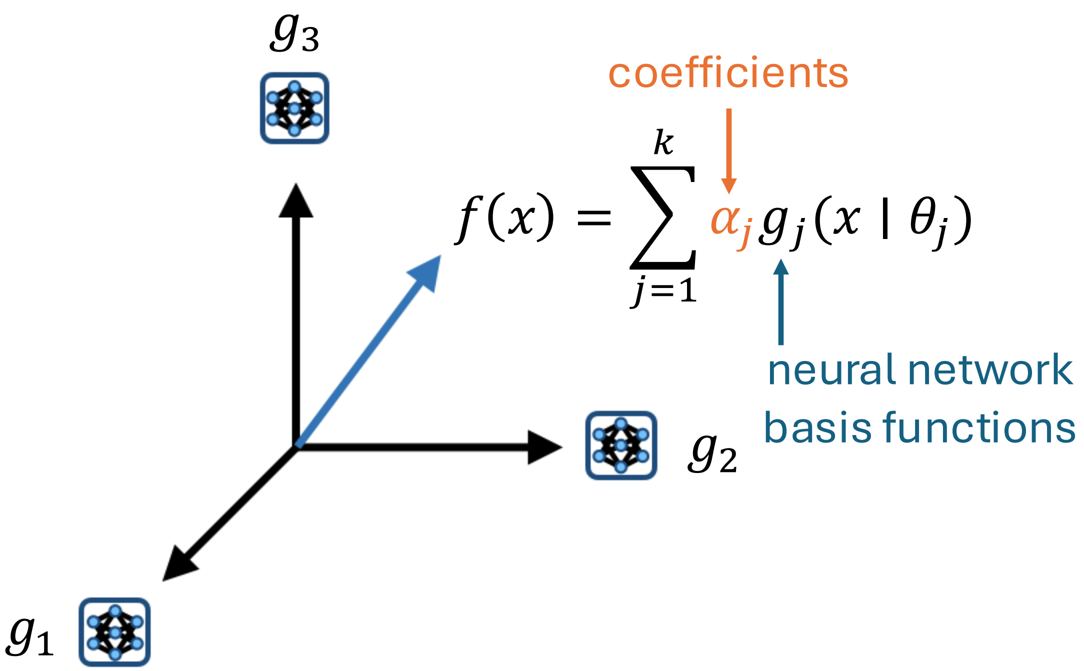

	simple polynomial example	function encoders
Basis:	$\{1 \quad x \quad x^2\}$	$\{g_1 \quad g_2 \quad g_3 \quad \cdots \quad g_k\}$
Representation:	$[a \quad b \quad c]$	$[\alpha_1 \quad \alpha_2 \quad \alpha_3 \quad \cdots \quad \alpha_k]$

Breaking function encoders down: **offline training, online inference**

Offline Training

learn the basis functions

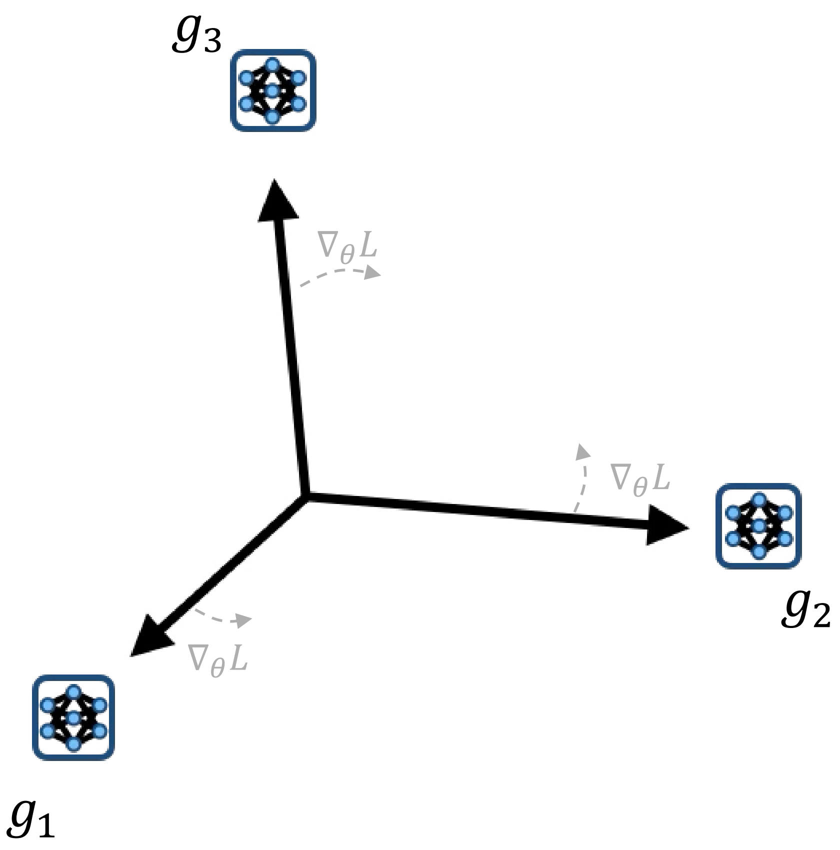

Online Inference

compute the coefficients α

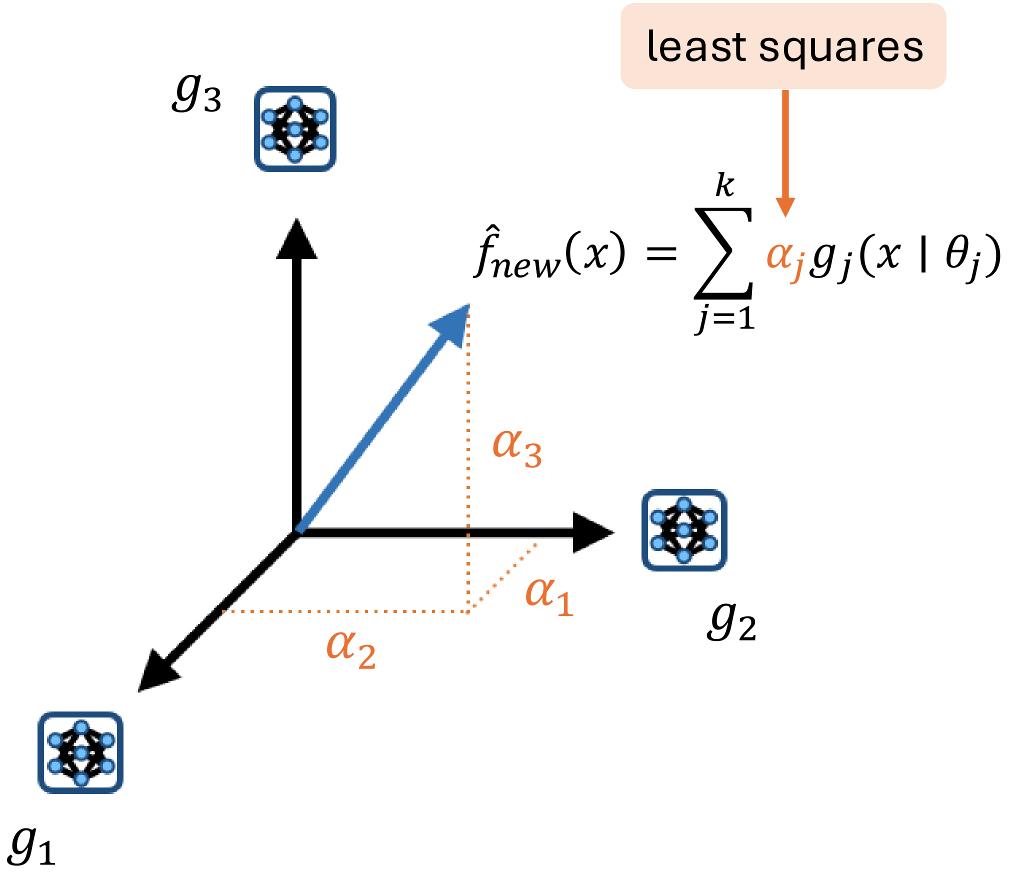

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2025). Function Encoders: A Principled Approach to Transfer Learning in Hilbert Spaces.

Offline Training: Training neural network basis functions

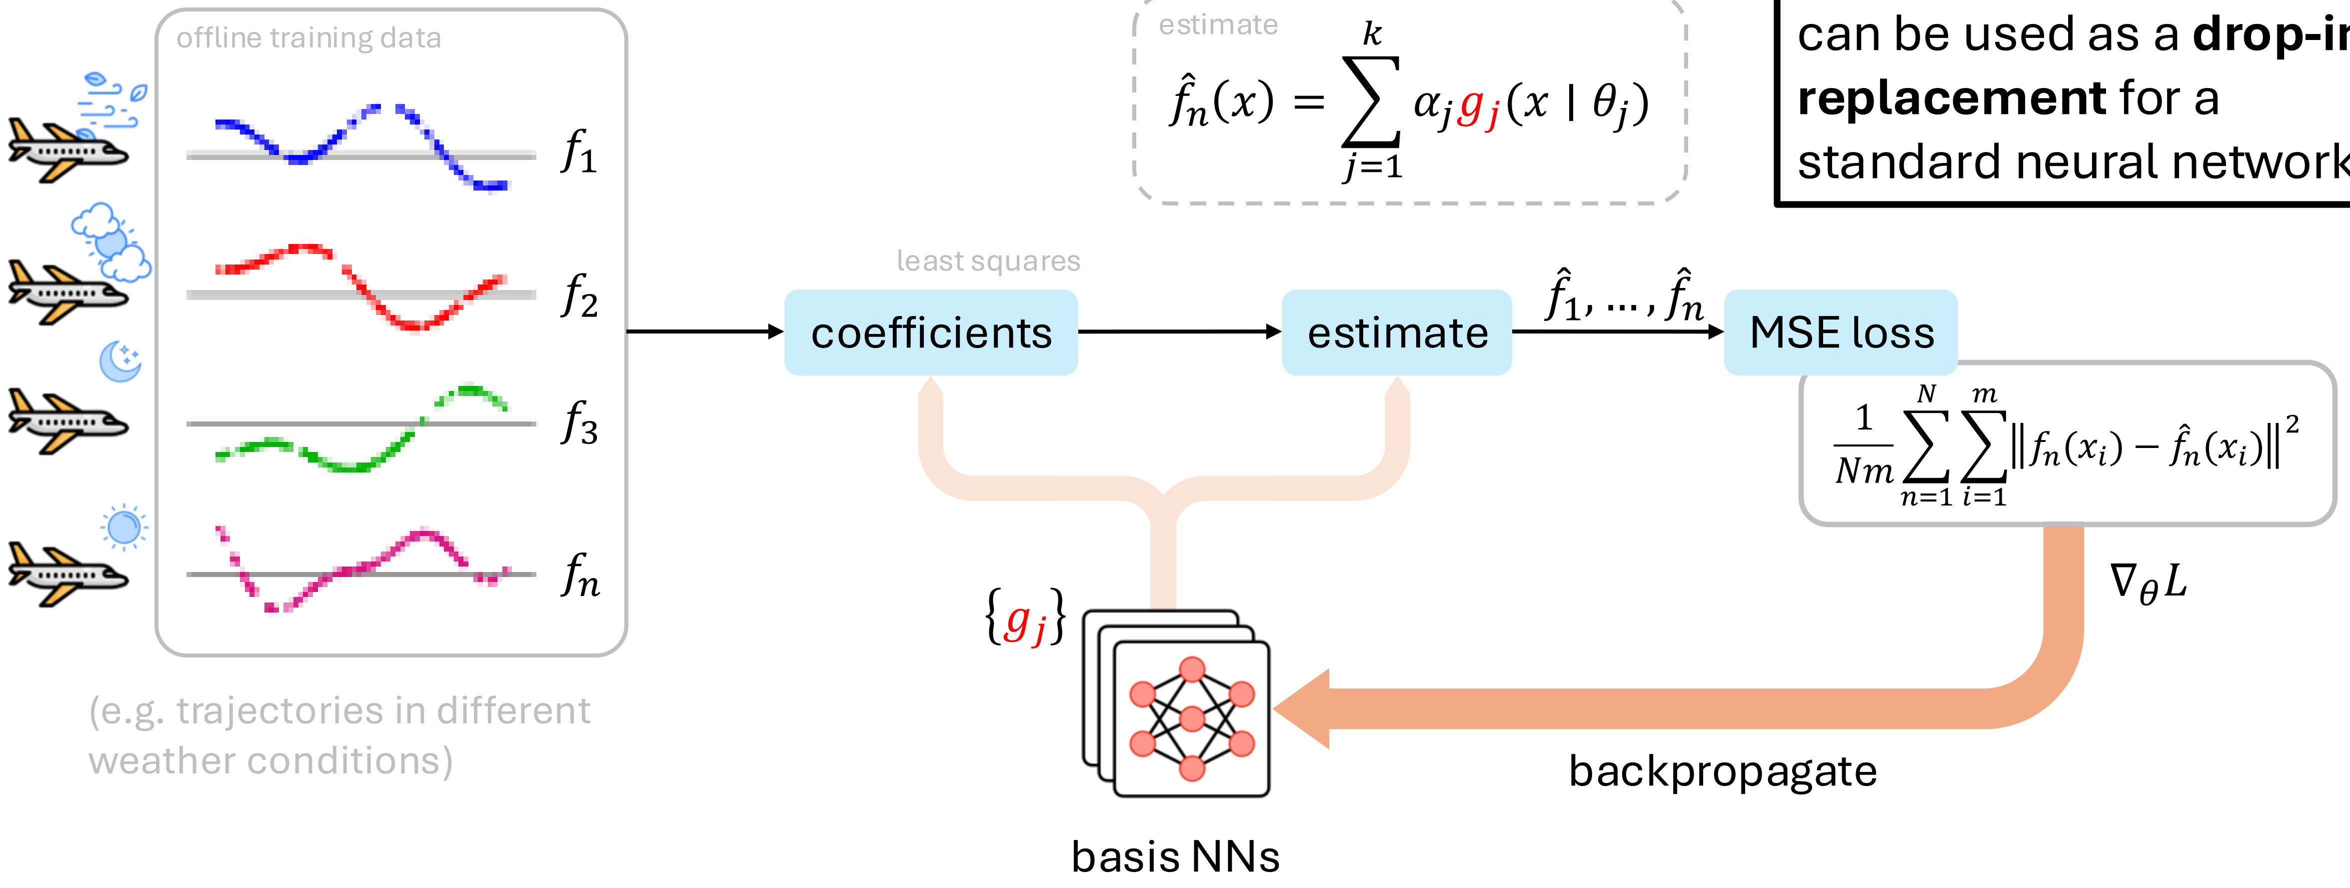

(e.g. trajectories in different weather conditions)

estimate

$$\hat{f}_n(x) = \sum_{j=1}^{k} \alpha_j g_j(x \mid \theta_j)$$

can be used as a **drop-in replacement** for a standard neural network

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2025). Function Encoders: A Principled Approach to Transfer Learning in Hilbert Spaces.

Online Inference: Computing coefficients for a new function

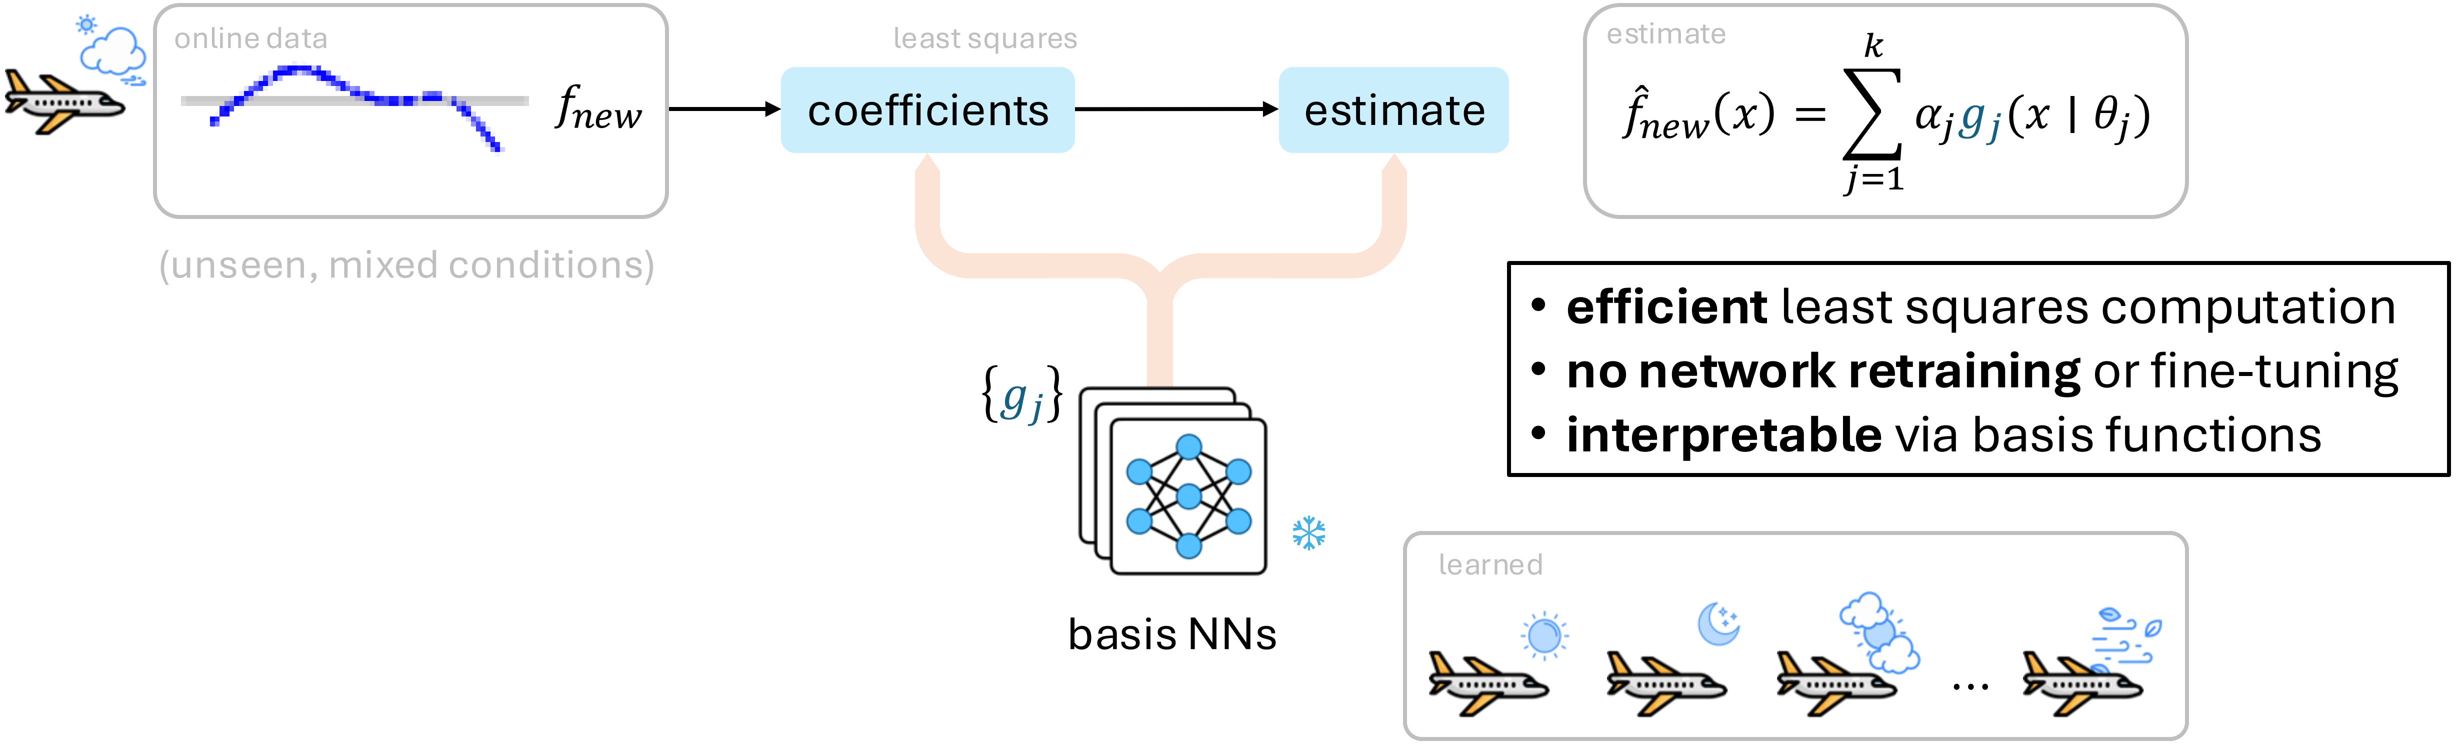

- **efficient** least squares computation
- **no network retraining** or fine-tuning
- **interpretable** via basis functions

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2025). Function Encoders: A Principled Approach to Transfer Learning in Hilbert Spaces.

Preview of main results

$$x(t) = x(0) + \int_0^t f_\theta(x(\tau), \tau) d\tau$$

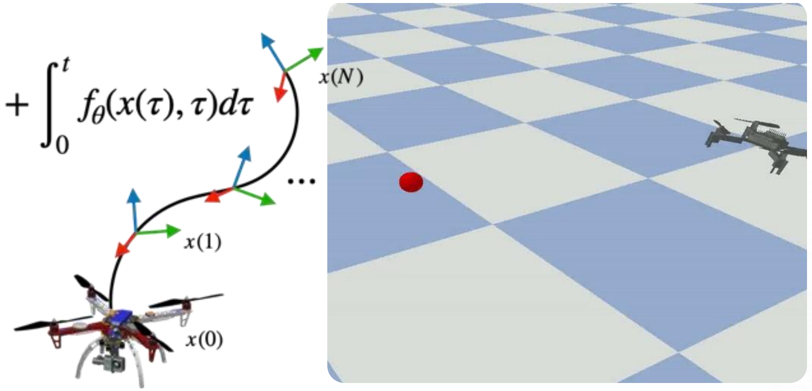

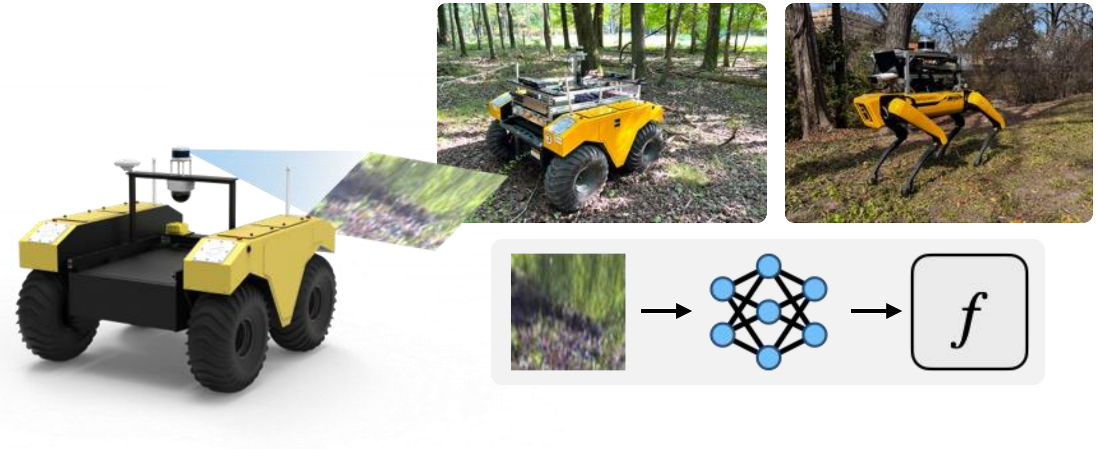

Zero-Shot System Identification & Control

Vision to Dynamics

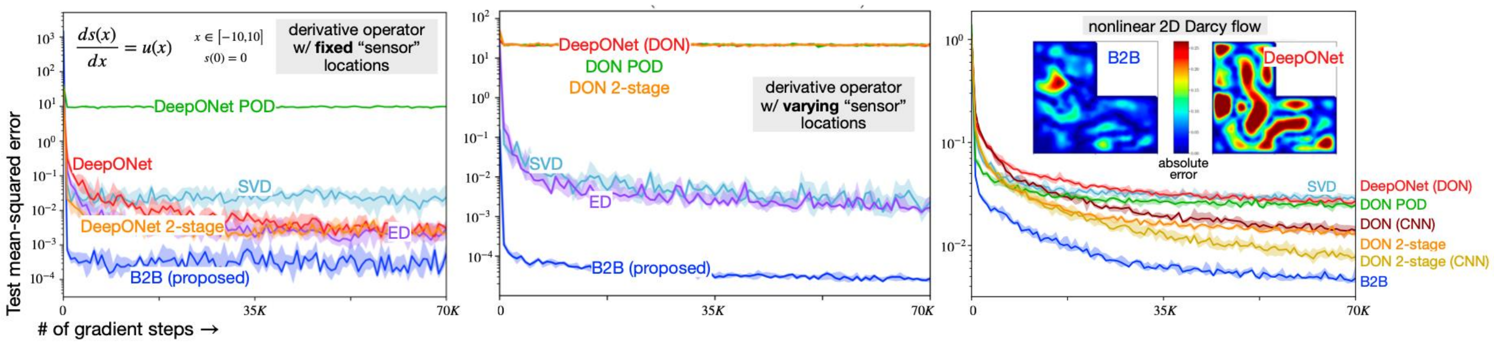

Basis to Basis Operator Learning

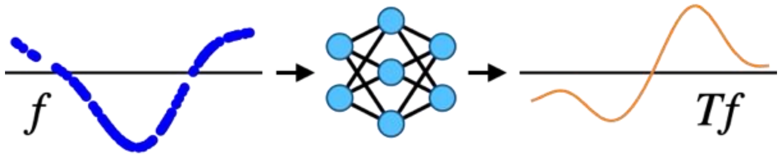

Zero-shot transfer of neural ODEs

Problem: how can we identify system dynamics at runtime using limited data?

Scenario: delivery drone with unknown mass/inertia

Given:
- historical trajectories with varying masses
- a small amount of online data

Goal: estimate dynamics

$$\dot{x}(t) = f(x(t), t)$$

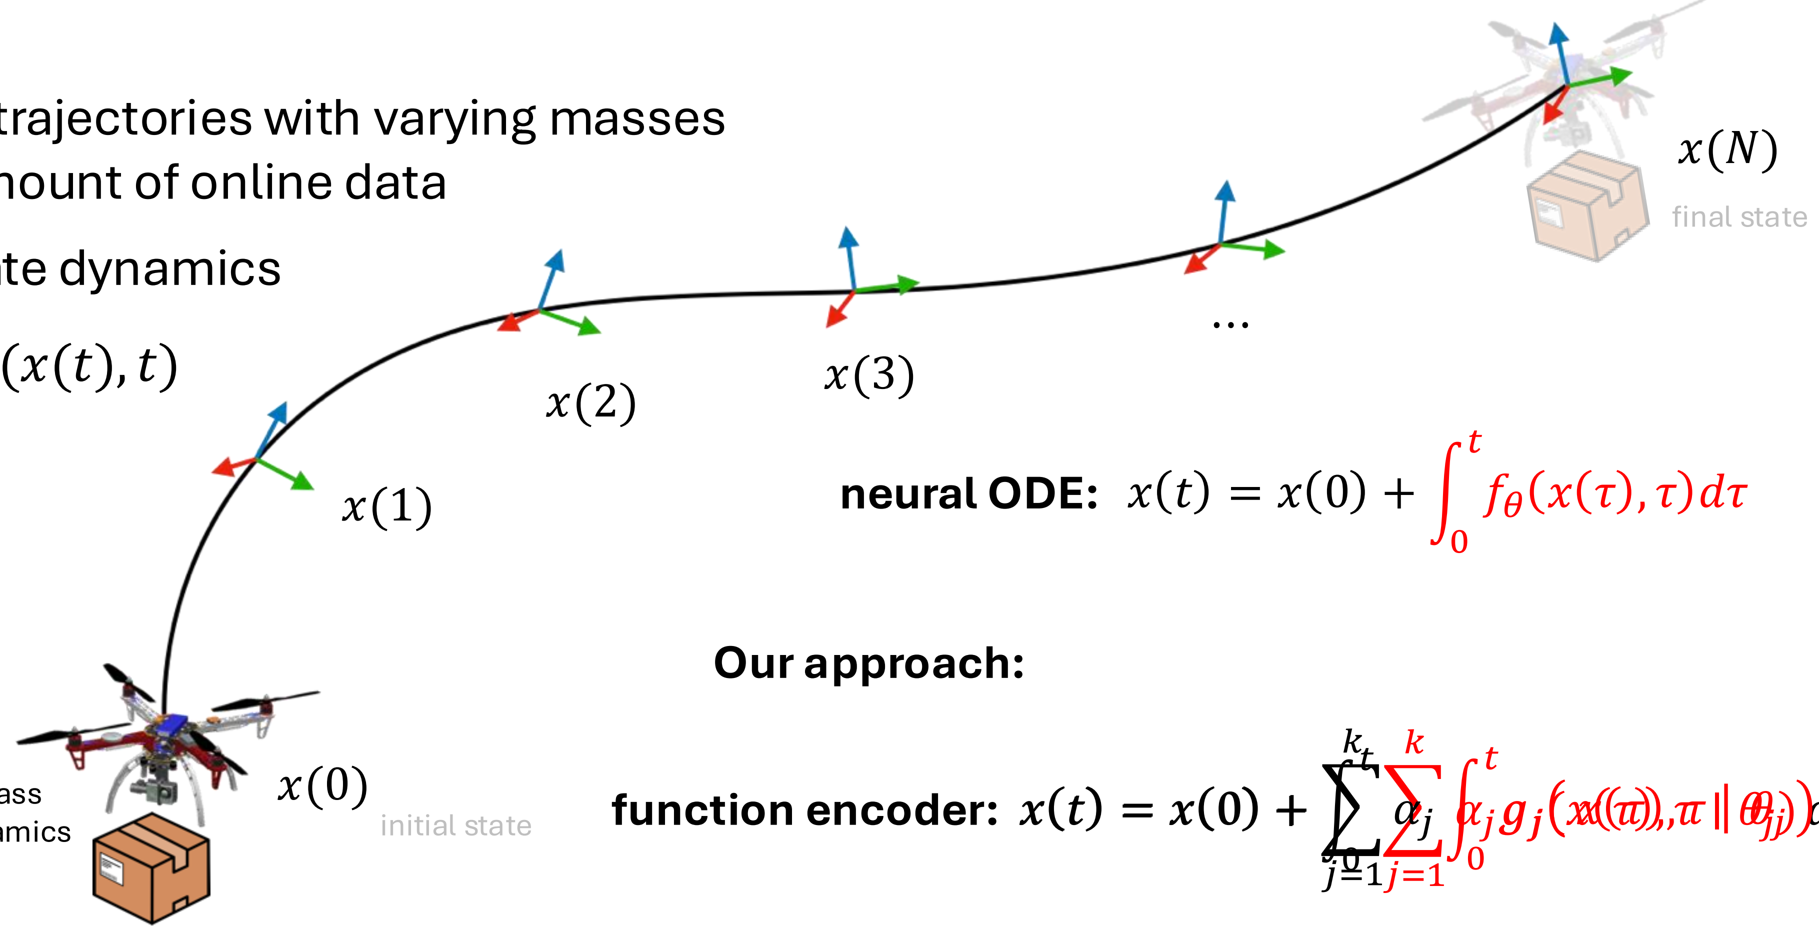

neural ODE: $x(t) = x(0) + \int_0^t f_\theta(x(\tau), \tau) d\tau$

Our approach:

function encoder: $x(t) = x(0) + \sum_{j=1}^{k} \int_0^t \alpha_j g_j(x(\tau), \tau \| \theta_j) d\tau$

Function encoders enable accurate zero-shot prediction

Zero shot: zero gradient updates, using only online data

long-horizon prediction accuracy

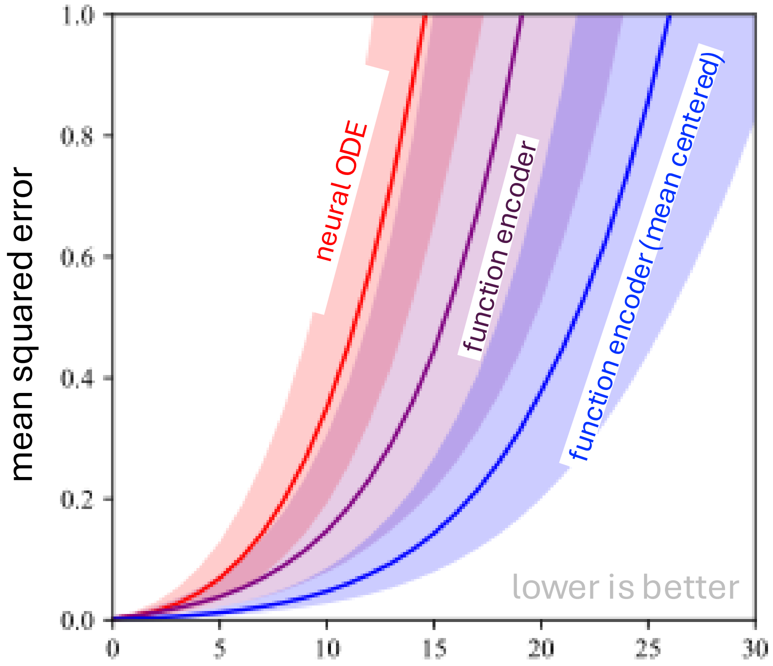

our approach makes accurate predictions from only a few **seconds** of online data

10-step prediction accuracy

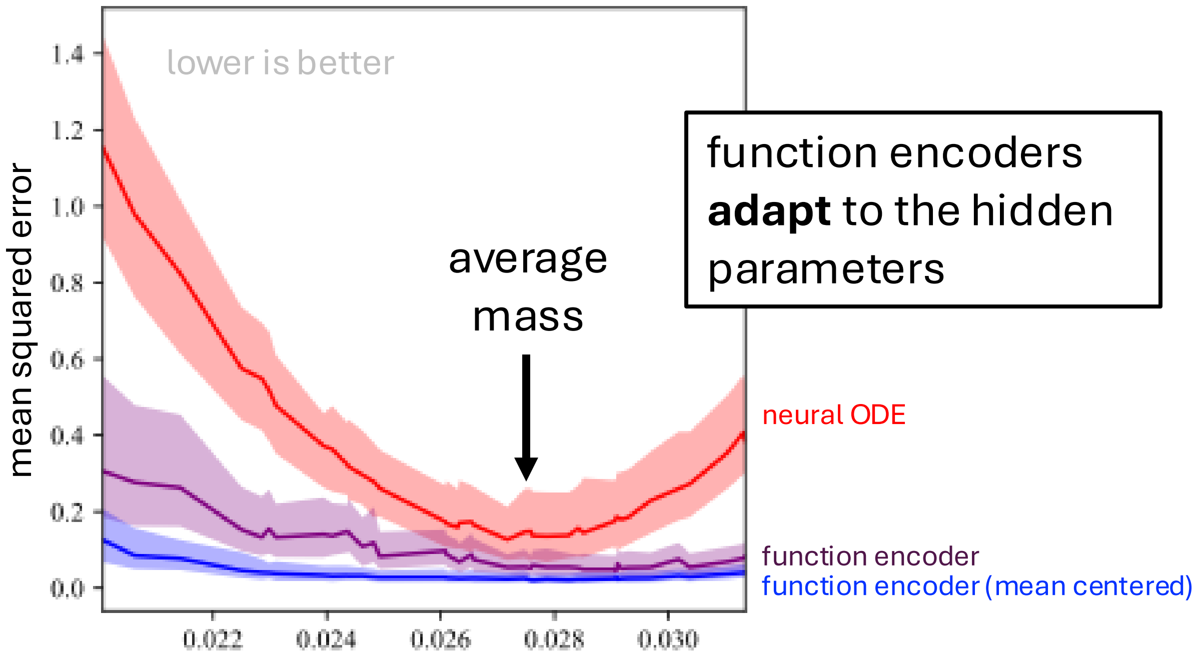

function encoders **adapt** to the hidden parameters

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2024). Zero-shot transfer of neural ODEs. NeurIPS

Function encoders enable accurate downstream control

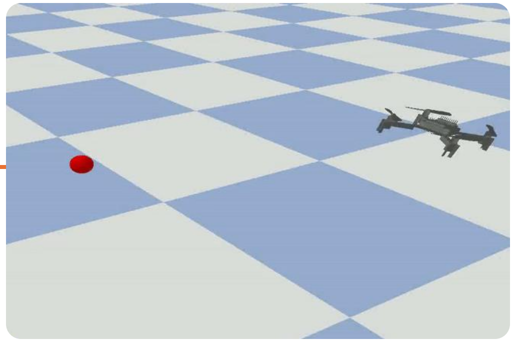

Accurate models are critical for **safe** control

neural ODE model has too much or too little thrust

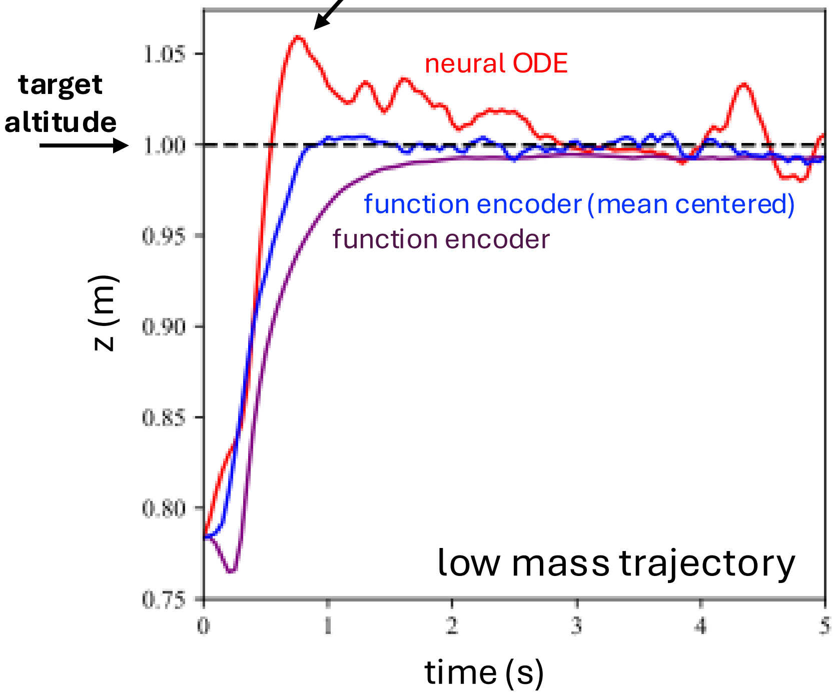

more accurate models lead to better control performance

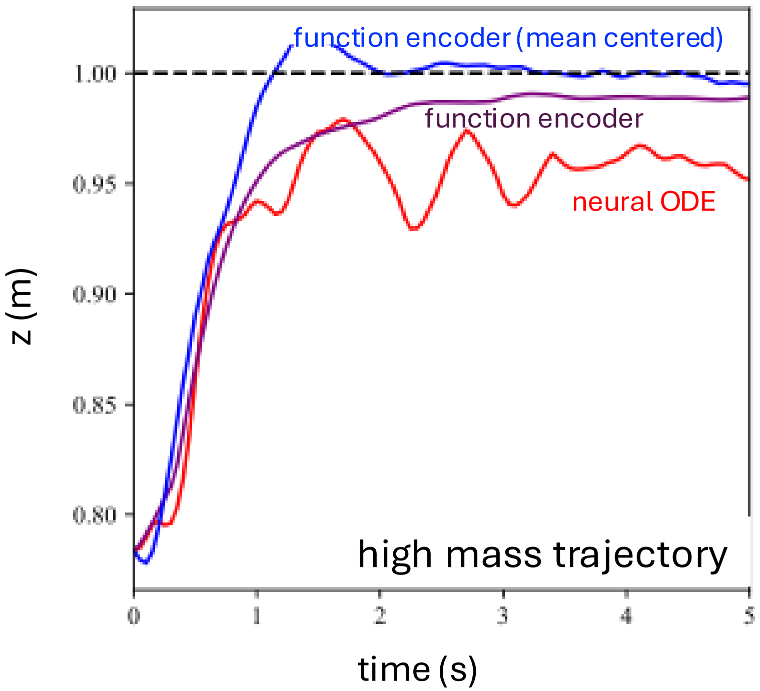

neural ODE baseline oscillates, **fails** to reach target altitude

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2024). Zero-shot transfer of neural ODEs. NeurIPS

Function encoders enable transfer **beyond** the training data

Using the properties of the Hilbert space for transfer

Van der Pol oscillator:

$$\dot{x}_1 = x_2$$
$$\dot{x}_2 = \mu(1 - x_2^2)x_2 - x_1$$

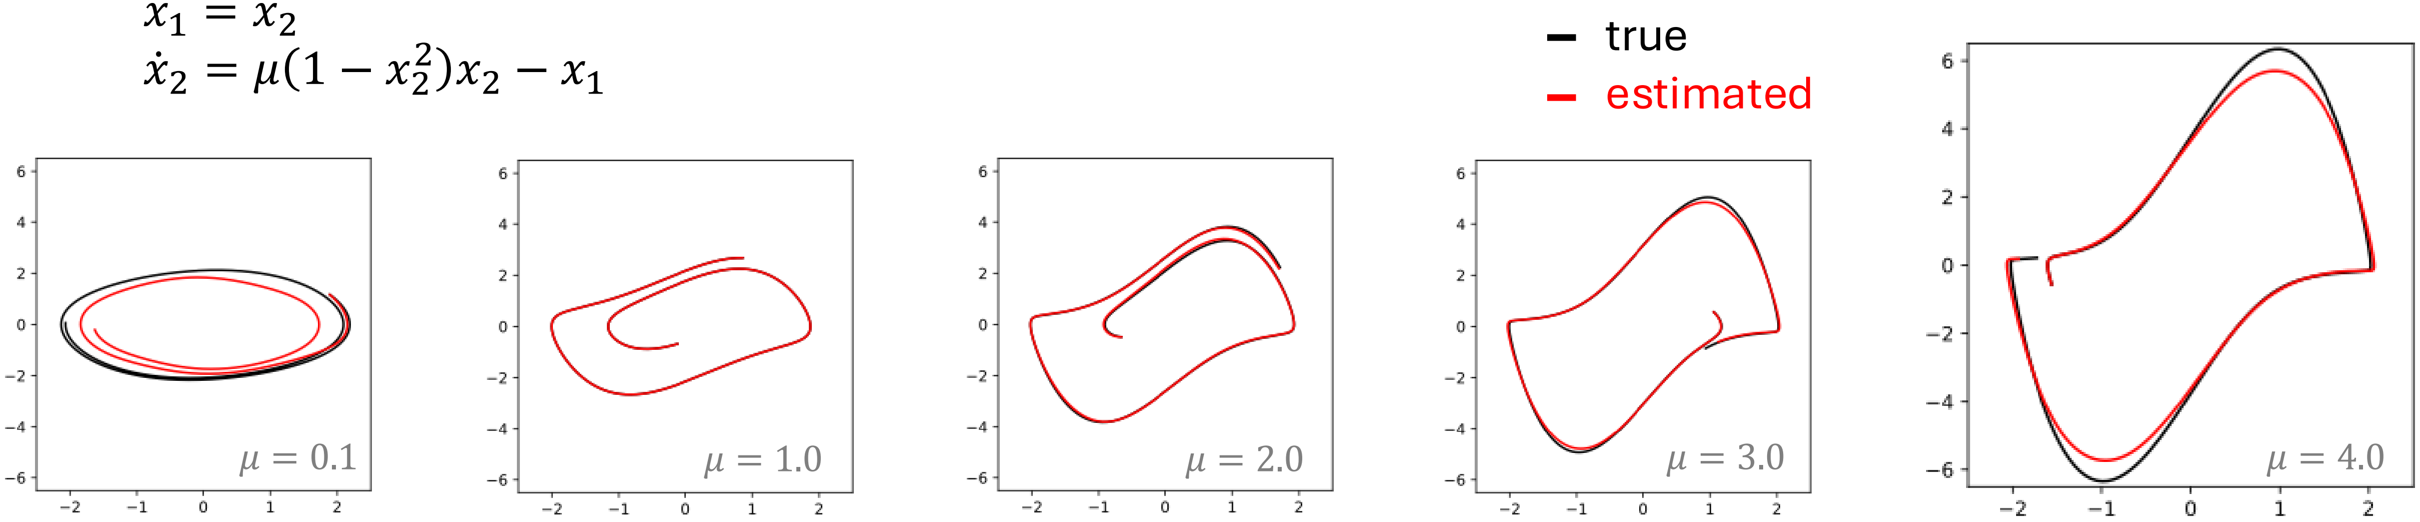

training data α_1

interpolation in the convex hull of α_1 and α_2

training data α_2

extrapolation outside the convex hull

inside the training data

outside the training data

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2024). Zero-shot transfer of neural ODEs. NeurIPS

A geometric characterization of transfer

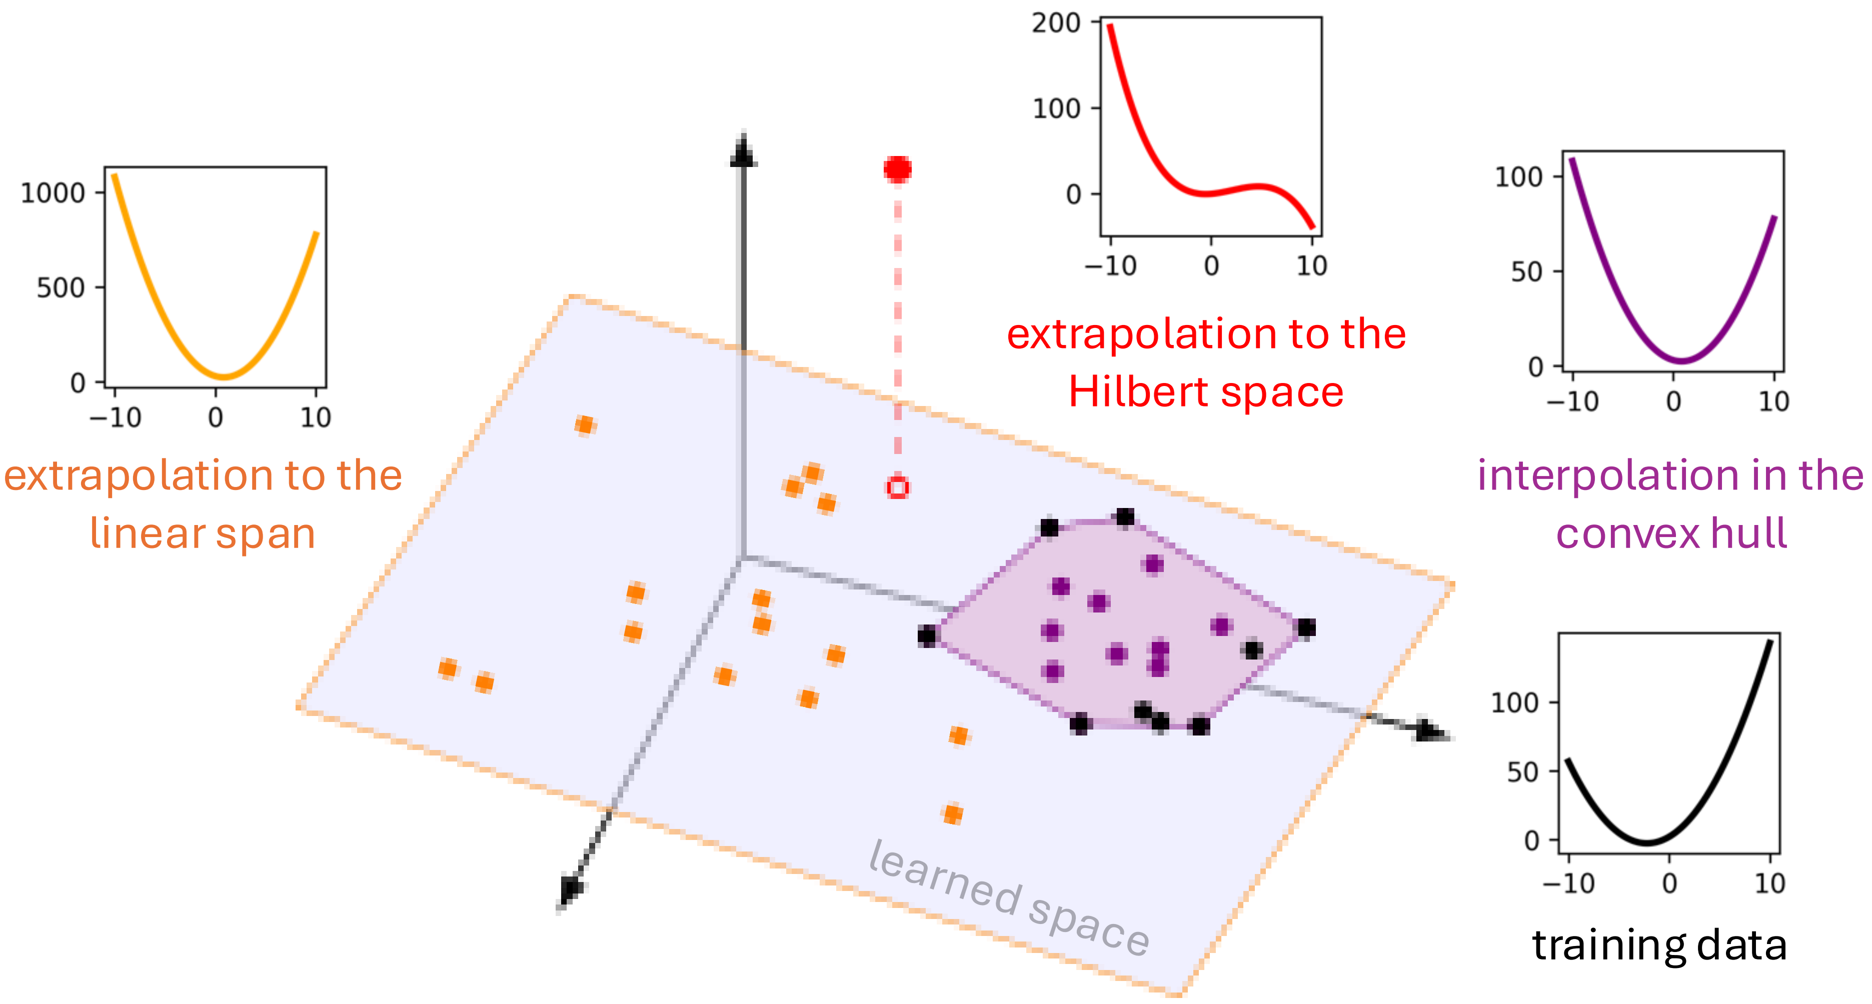

A very simple transfer test

How well do existing approaches transfer?

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2025). Function Encoders: A Principled Approach to Transfer Learning in Hilbert Spaces. ICML

More transfer applications

Image Classification

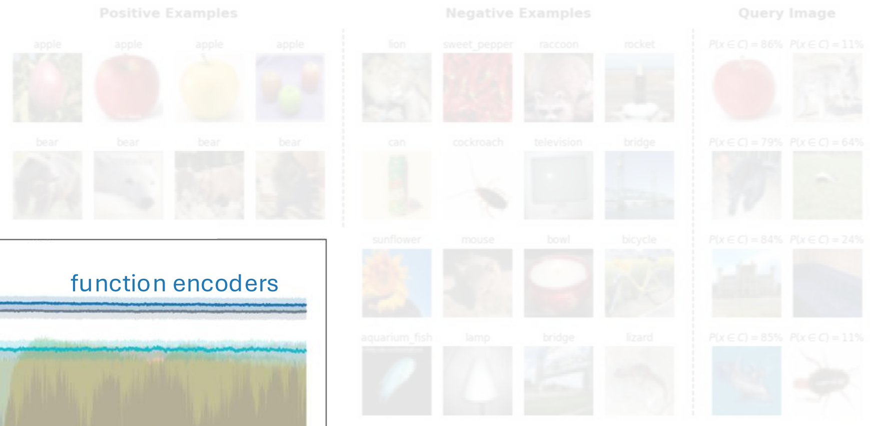

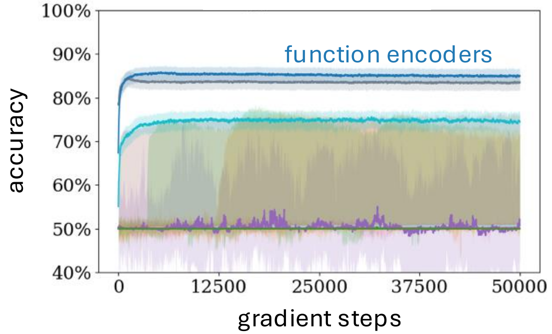

higher is better

Pose Estimation

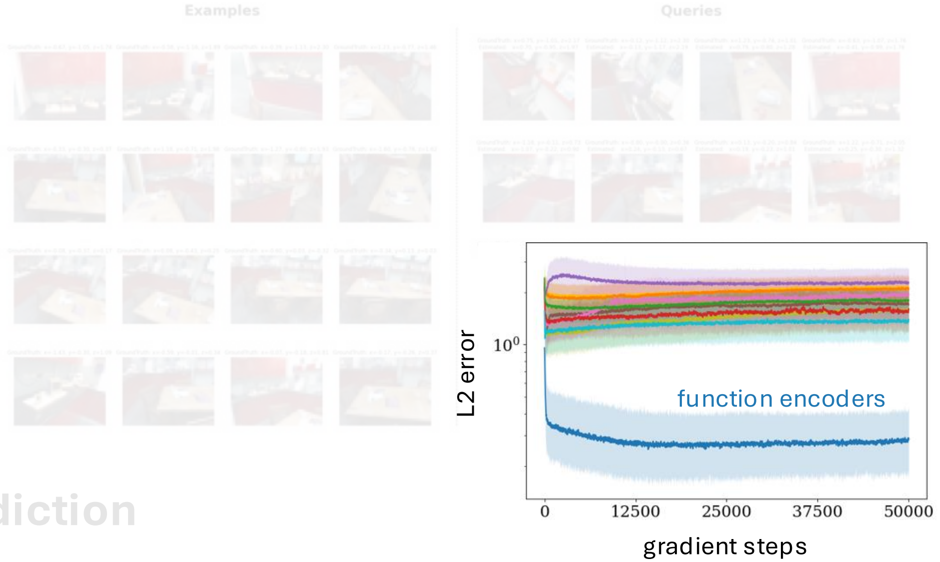

lower is better

Dynamics Prediction

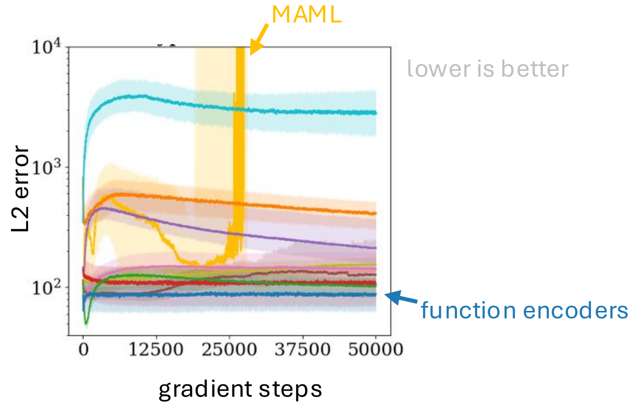

lower is better

Ingebrand, T., Thorpe, A. J., & Topcu, U. (2025). Function Encoders: A Principled Approach to Transfer Learning in Hilbert Spaces. ICML

Outline

Zero-Shot System Identification & Control

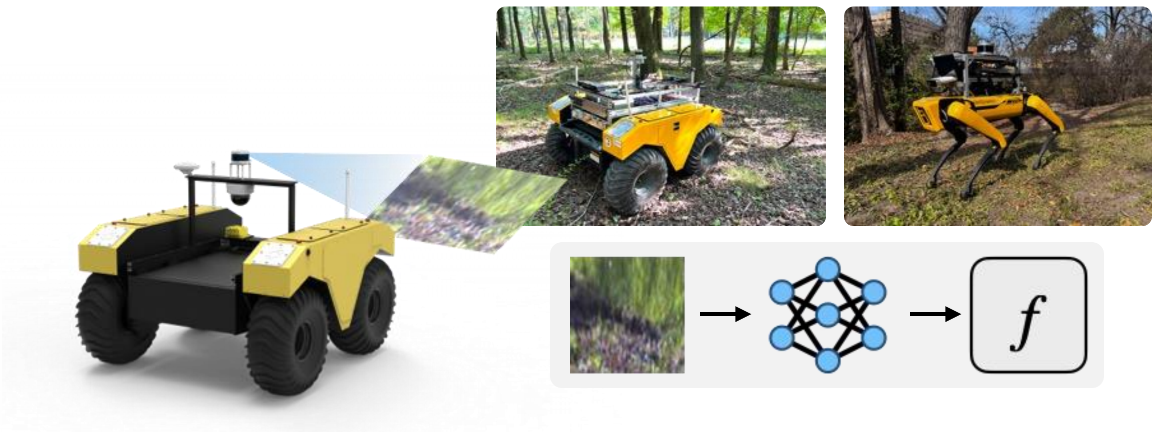

Vision to Dynamics

Basis to Basis Operator Learning

Adaptation and transfer for robotics

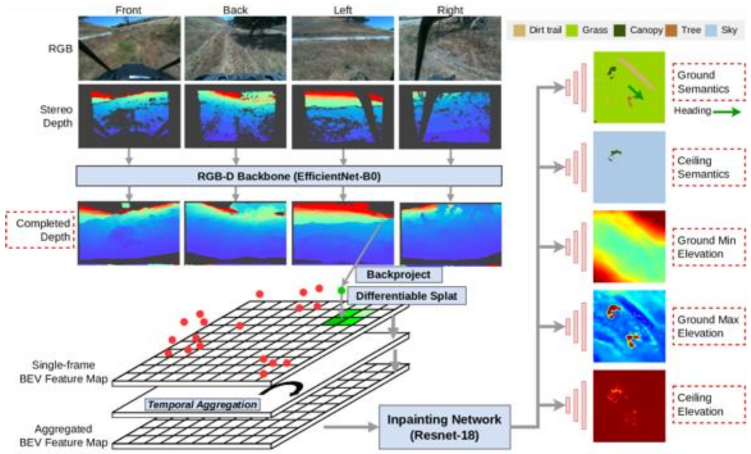

Mapping

X. Meng, N. Hatch, A. Lambert, A. Li, N. Wagener, M. Schmittle, J. Lee, W. Yuan, Z. Chen, S. Deng, G. Okopal, D. Fox, B. Boots, A. Shaban (2023). Terrainnet: Visual modeling of complex terrain for high-speed, off-road navigation.

Navigation

Han, T., Liu, A., Li, A., Spitzer, A., Shi, G., & Boots, B. (2023). Model predictive control for aggressive driving over uneven terrain.

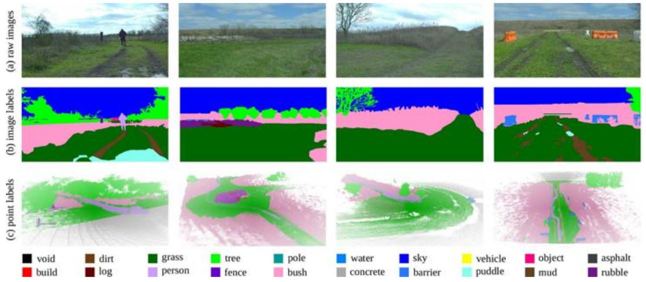

Semantic Segmentation

Jiang, P., Osteen, P., Wigness, M., & Saripalli, S. (2021). Rellis-3d dataset: Data, benchmarks and analysis

What's missing?

the ability to **adapt** & **transfer** to unseen environments or terrains

Adapting robots to new conditions at runtime

Problem: how can we adapt to different terrains?

Using camera images to adapt to the terrain

Problem: how can we adapt to different terrains?
Given: historical data on various terrains
Goal: estimate the robot dynamics from camera images

The Great Outdoors Dataset

DEVCOM
ARMY RESEARCH
LABORATORY

V2D: Vision to Dynamics

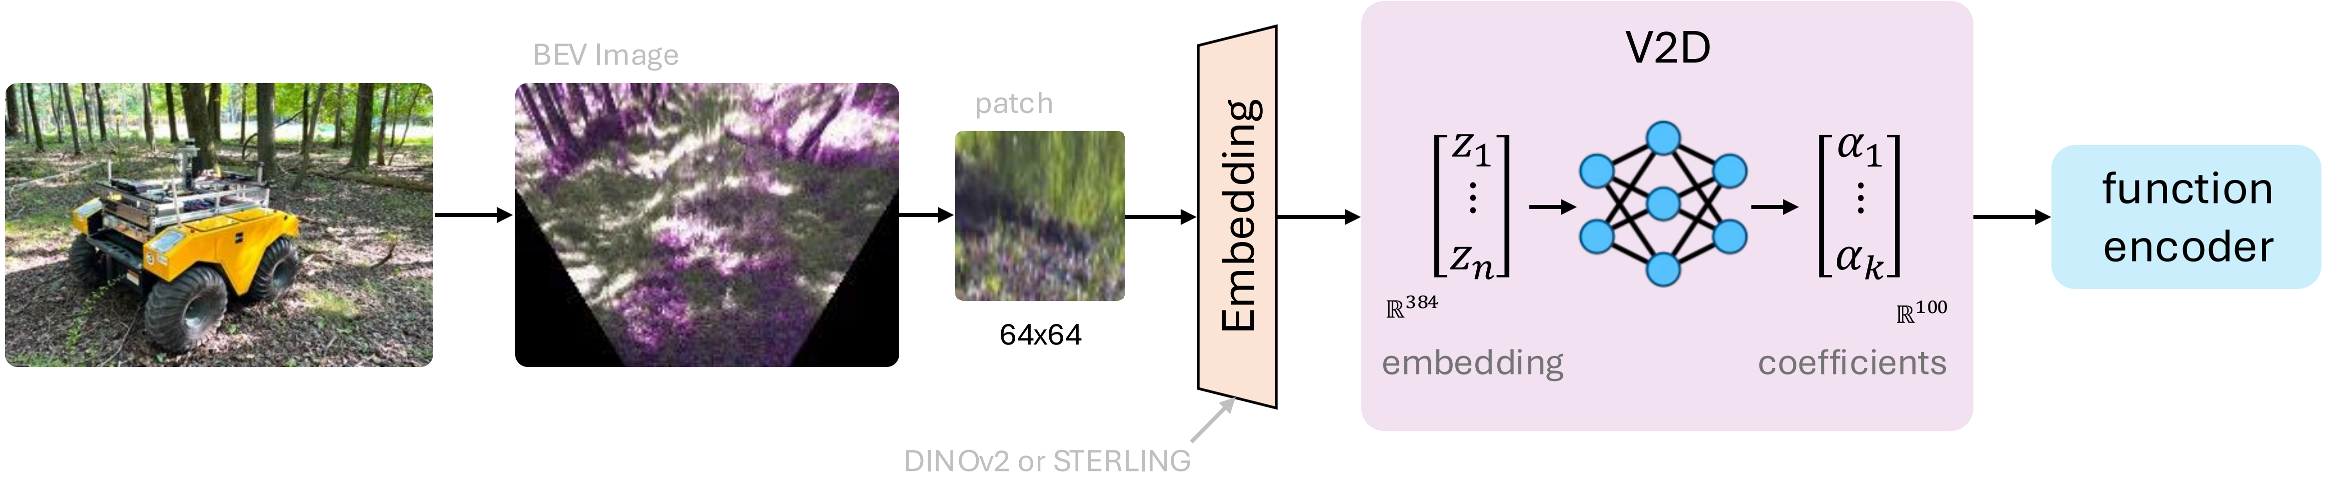

Results

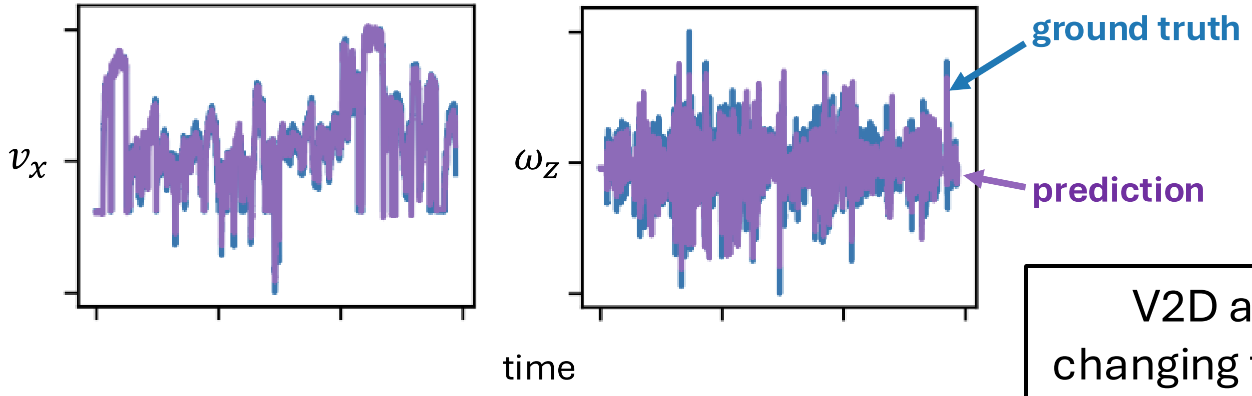

V2D adapts to changing terrain using **only** camera images

Driving on Mars:

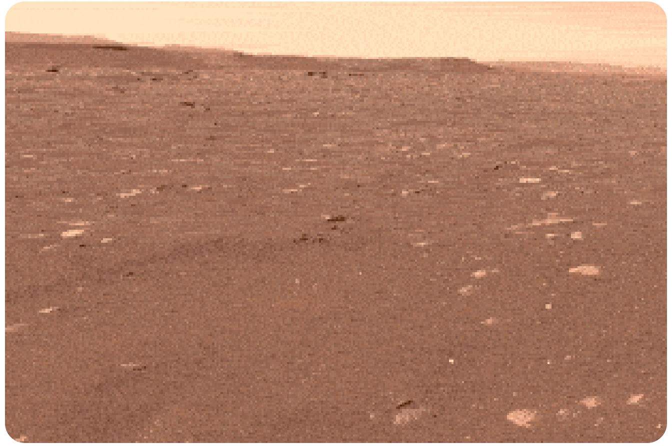

Outline

Zero-Shot System Identification & Control

Vision to Dynamics

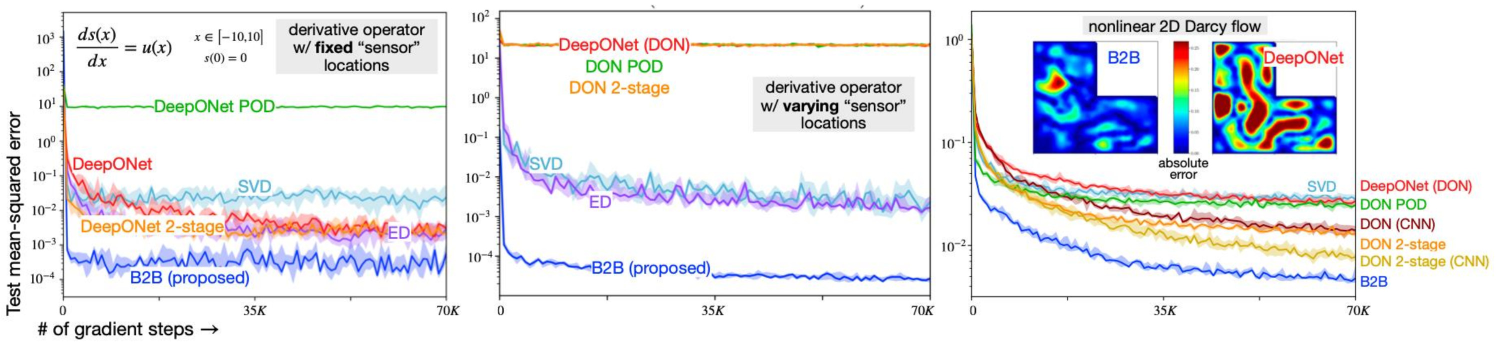

Basis to Basis Operator Learning

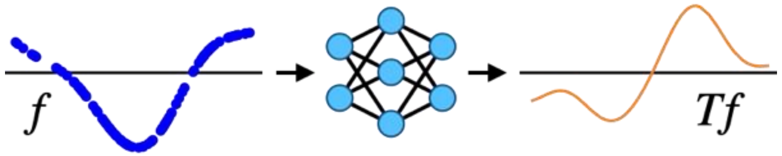

Neural operator learning: function to function maps

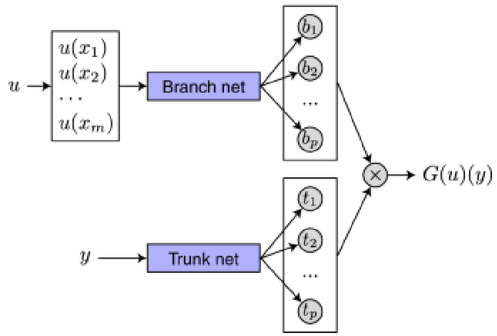

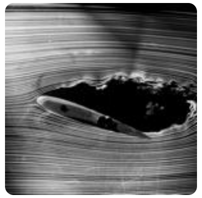

DeepONet

Lu, L., Jin, P., Pang, G., Zhang, Z., & Karniadakis, G. E. (2021). Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators.

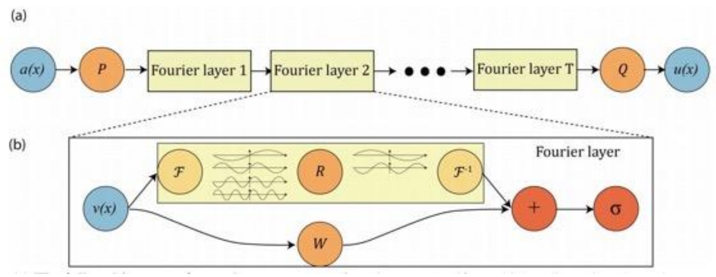

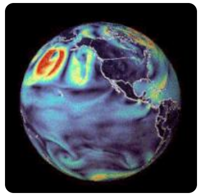

Fourier Neural Operators

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2020). Fourier neural operator for parametric partial differential equations.

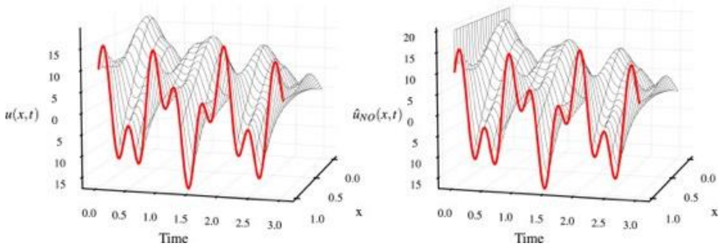

PDE Control

Bhan, L., Shi, Y., & Krstic, M. (2023). Neural operators for bypassing gain and control computations in PDE backstepping

Main challenge:

DeepONet & FNO require the input data to be on a **fixed grid** or **mesh**

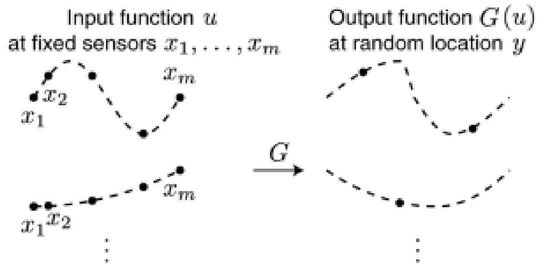

Basis to Basis Operator Learning (B2B)

Given: input-output pairs of transformations (f, Tf)
Goal: approximate $T: \mathcal{F} \to \mathcal{H}$

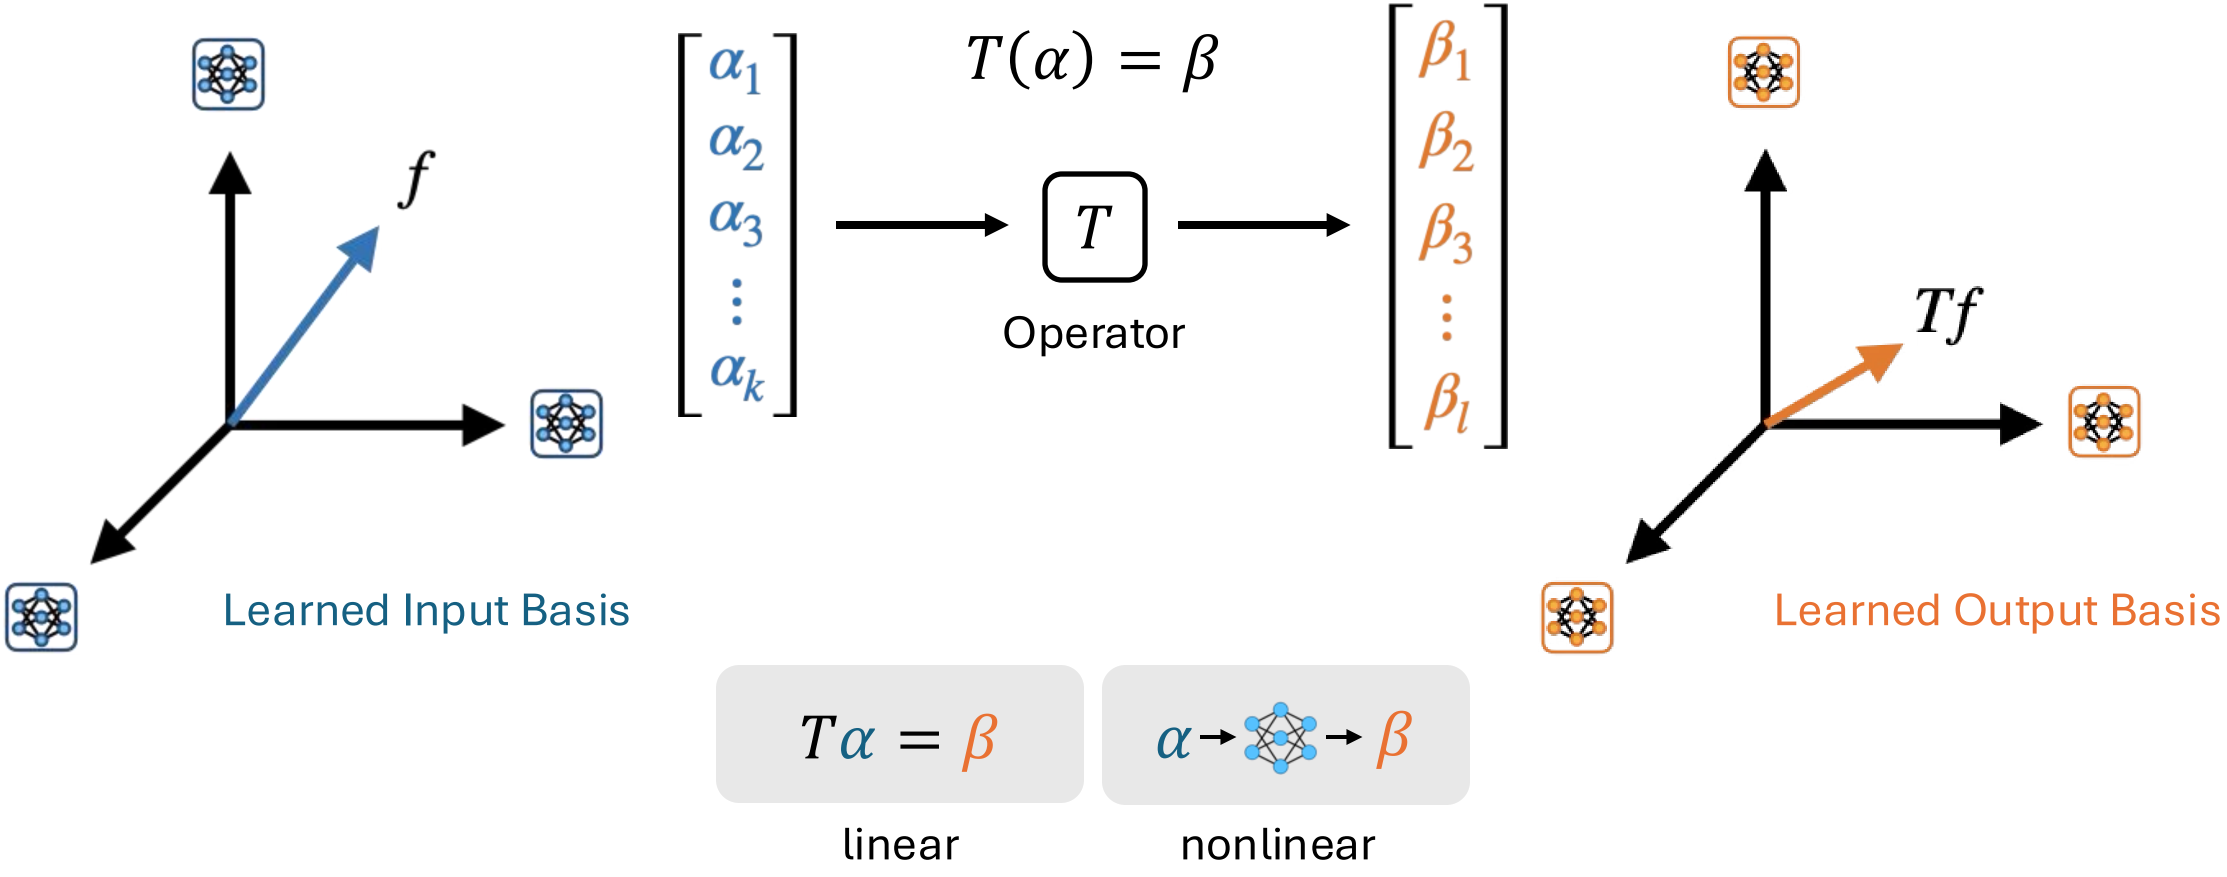

Ingebrand, T., Thorpe, A. J., Goswami, S., Kumar, K., & Topcu, U. (2025). Basis-to-basis operator learning using function encoders. CMAME

An illustrative linear example: derivative & antiderivative

$$\frac{ds(x)}{dx} = u(x), \qquad s(0) = 0, \qquad Tu(x) = s(x = 0) + \int_0^x u(t)\,dt$$

derivative operator

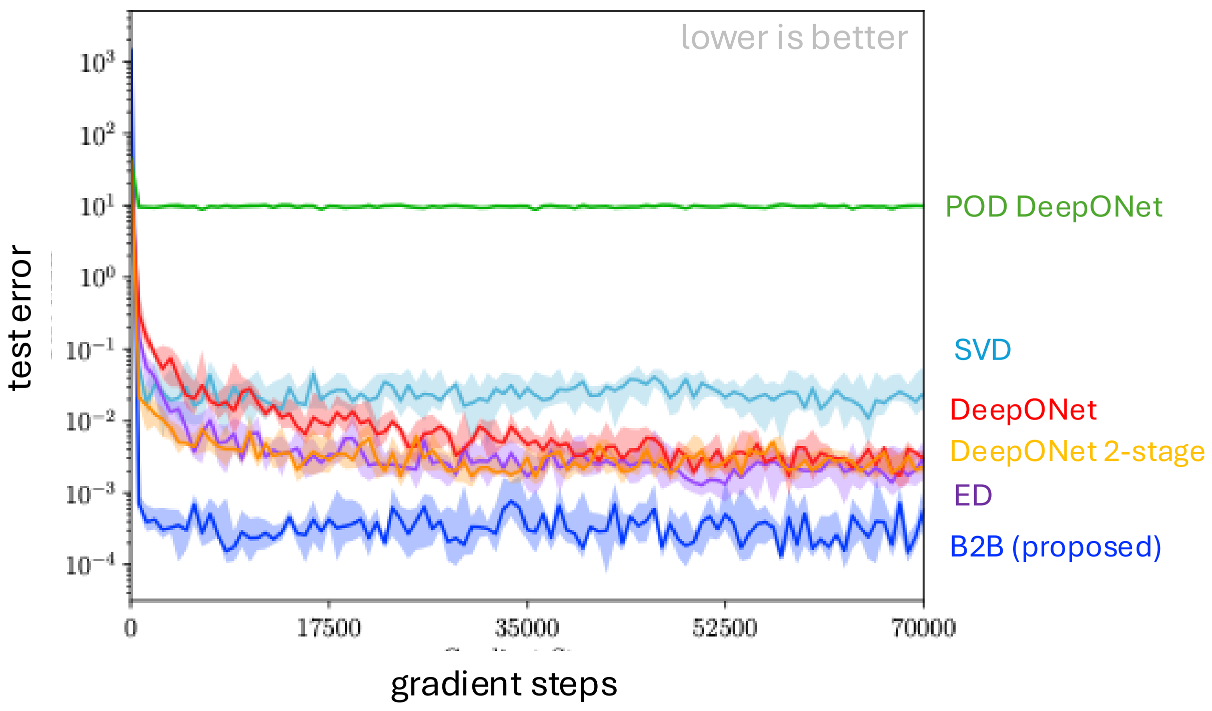

antiderivative operator

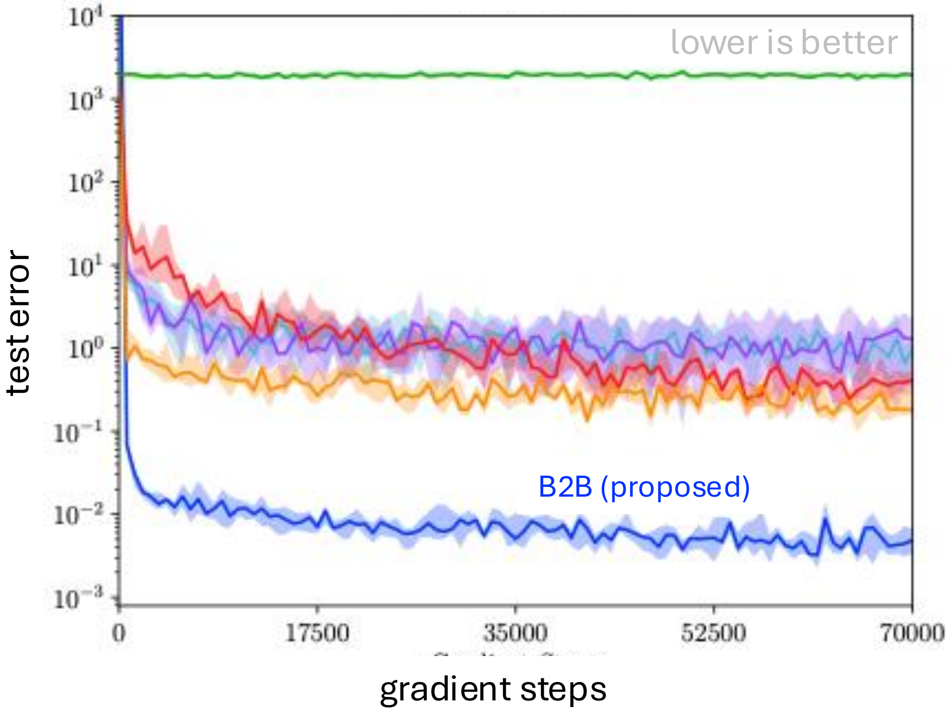

varying sensor locations:

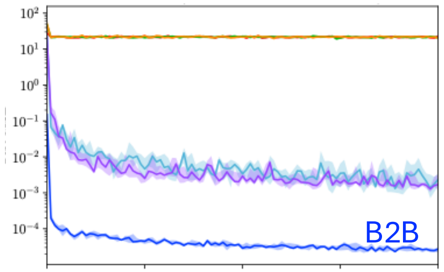

- **extrapolates** to the linear span
- **maintains accuracy,** even when the measurement locations **change**

A nonlinear example of basis to basis for PDE modeling

Modeling the solution of a partial differential equation

input function

$$\nabla \cdot \sigma + f(\boldsymbol{x}) = 0, \boldsymbol{x} = (x, y)$$
$$(u, v) = 0, \forall x = 0$$

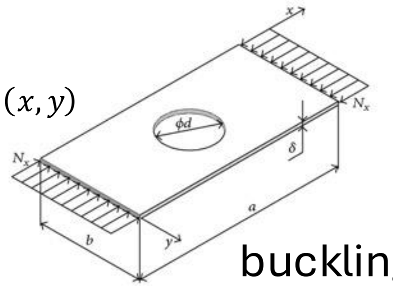

buckling

loading function (input function) | displacement (output function) | ground truth | error

B2B

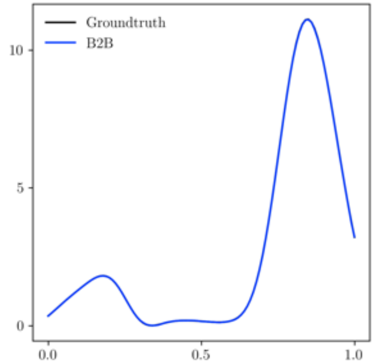

→

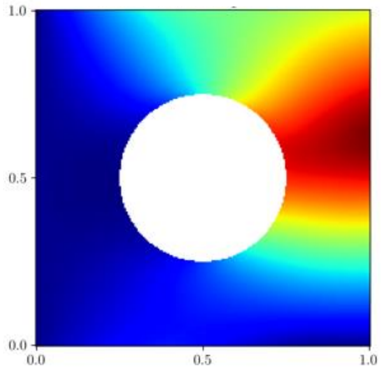

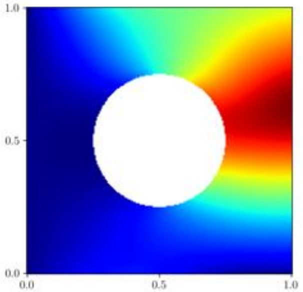

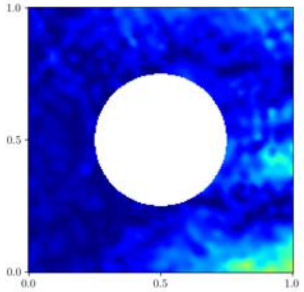

DeepONet

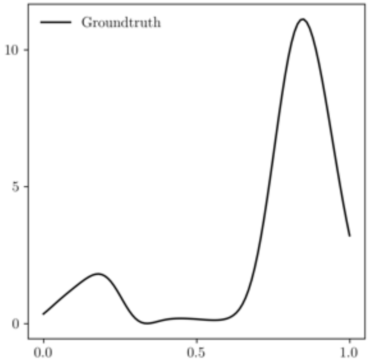

→

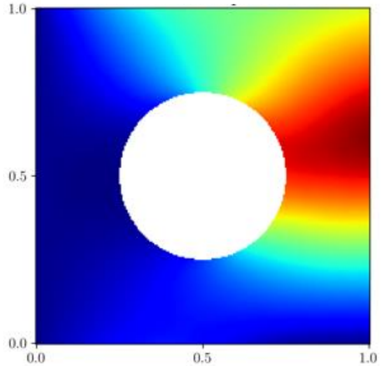

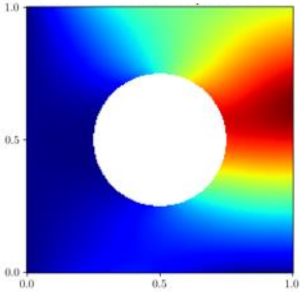

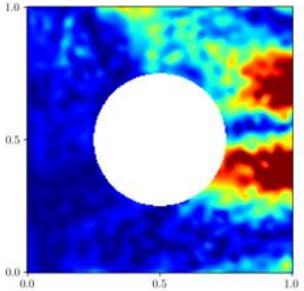

B2B has lower error, and **doesn't rely on a fixed grid or mesh**.

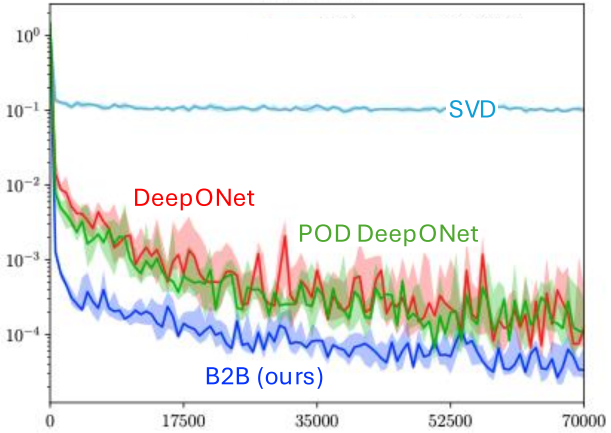

Neural operators model the **entire** solution, not just one instance!

My research: learning for autonomy

Learning Structured Representations

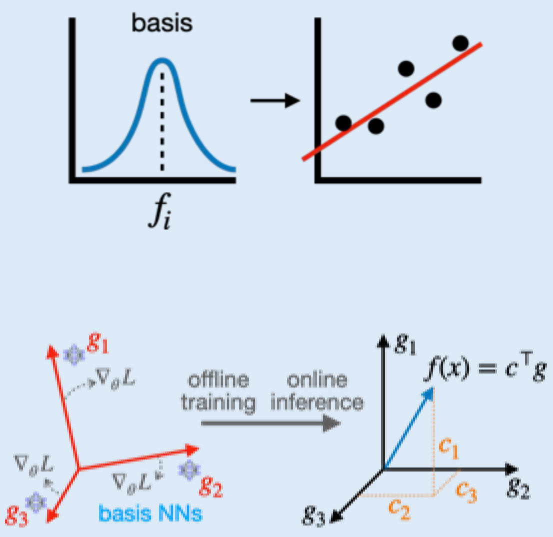

Function Encoders

Adaptable Autonomy & Inductive Transfer

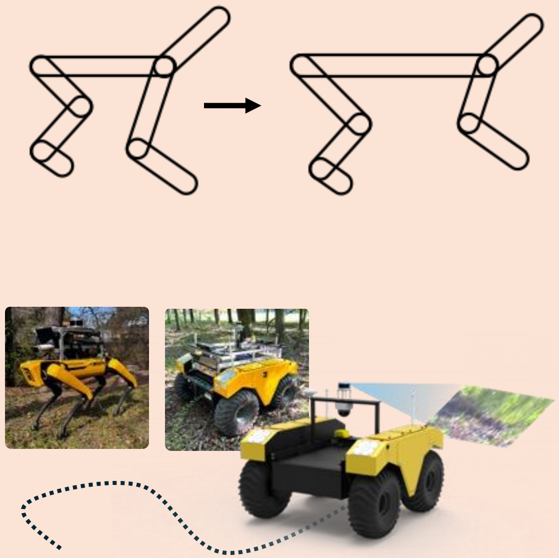

Zero-Shot Modeling & Control

Neural Operator Learning

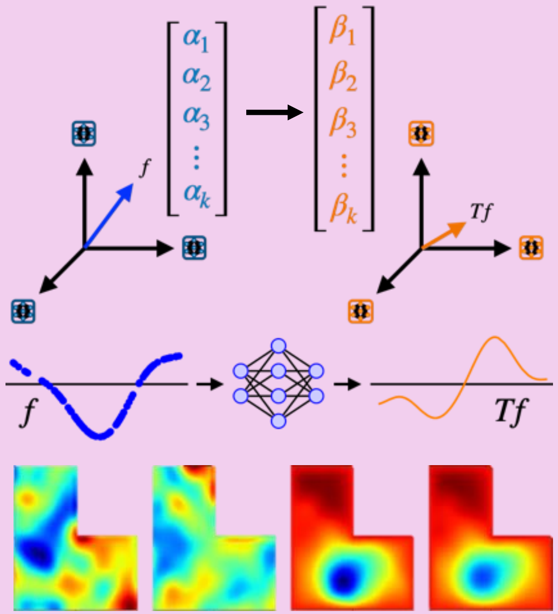

Basis to Basis Operators

Human-Autonomy Interaction

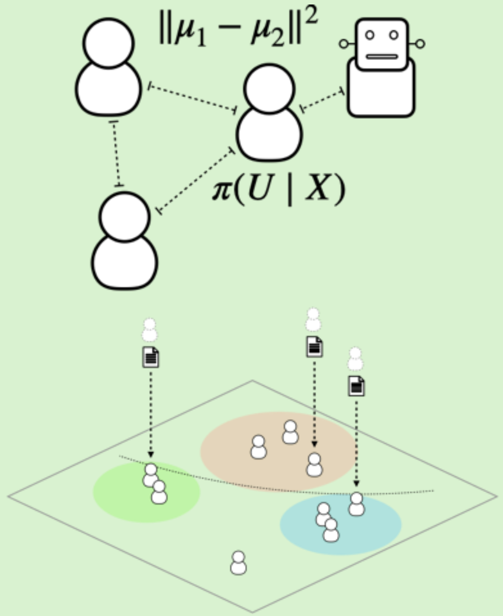

Characterizing Heterogeneity

Operator learning & inverse problems

Neural operator learning represents a **new frontier in learning and autonomy**.

We need fundamentally new algorithms to handle adaptation & transfer that offer **global** solutions.

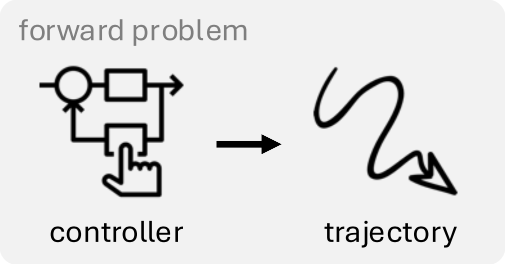

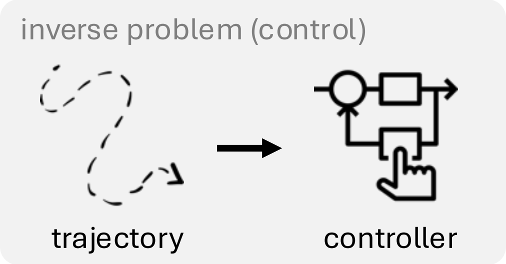

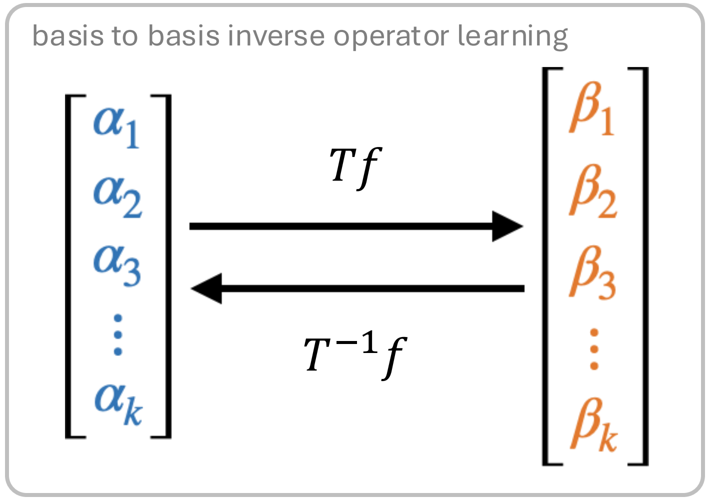

Bellman operator: $Tf(x,u) = R(x,u) + \gamma\langle \mathbb{P}(\cdot \mid x,u), V_f\rangle$

"inferring causes from effects"

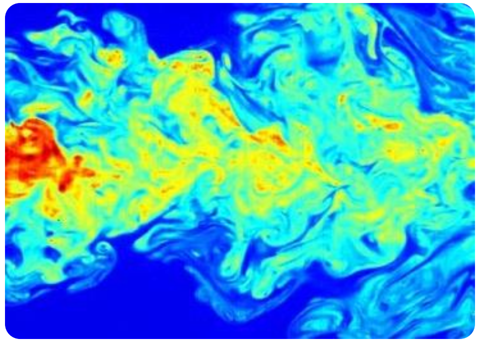

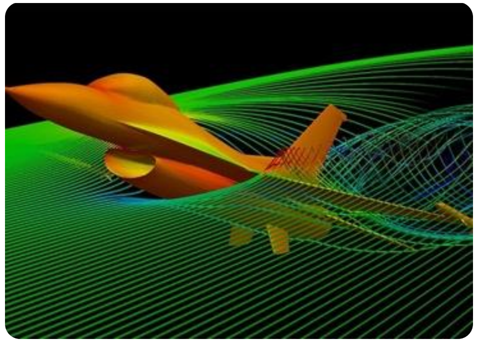

Beyond scale: adaptation and transfer through structure

To advance autonomy, we need to explore **new strategies** for adaptation and transfer that **incorporate structure**.

same-day autonomy

historical data
physics knowledge
surrogate models
simulated data
deployment
analytic models
low fidelity simulation
high fidelity simulation
adaptation

0 hours — 12 hours — 24 hours

multi-fidelity transfer

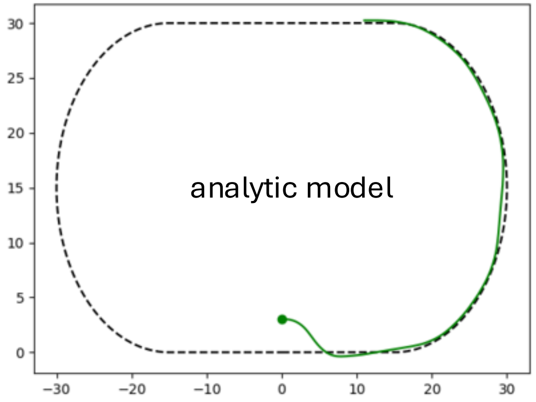

Putting it together: digging on the Moon (...or wherever we go to next)

The Problem: we won't know **until we get there...**

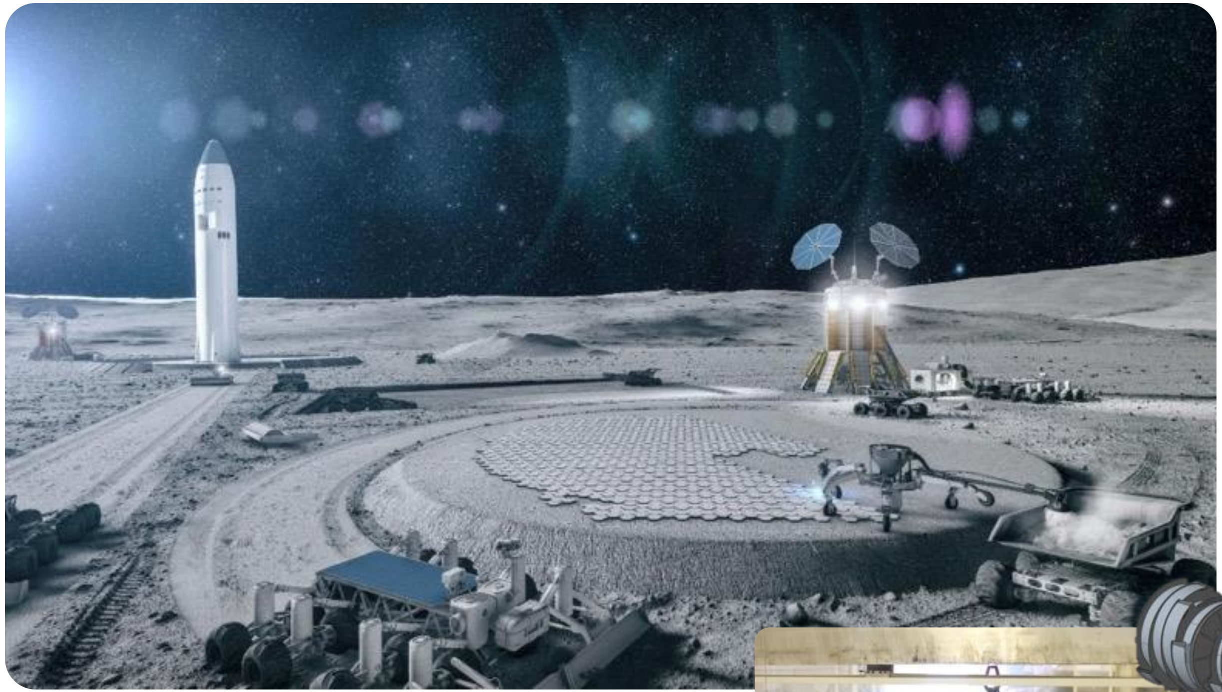

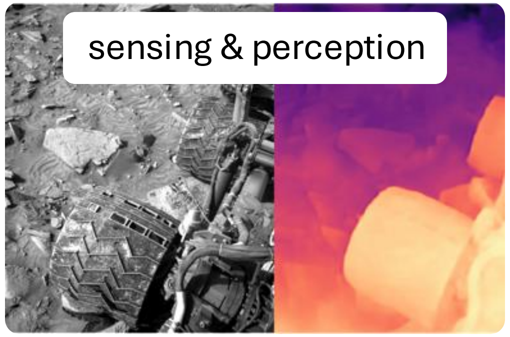

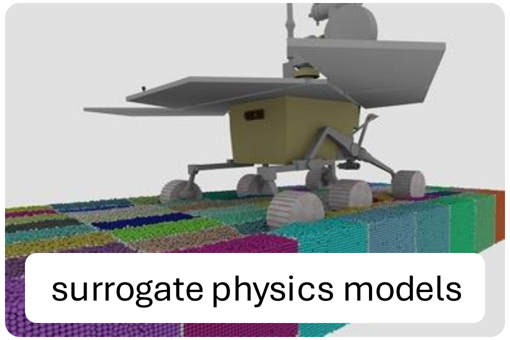

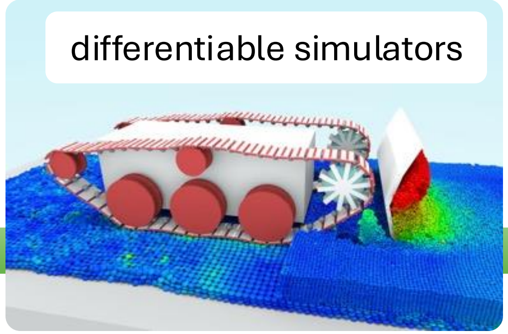

$$\frac{ds(x)}{dx}$$

Putting it together: applications in aerospace autonomy

urban air mobility

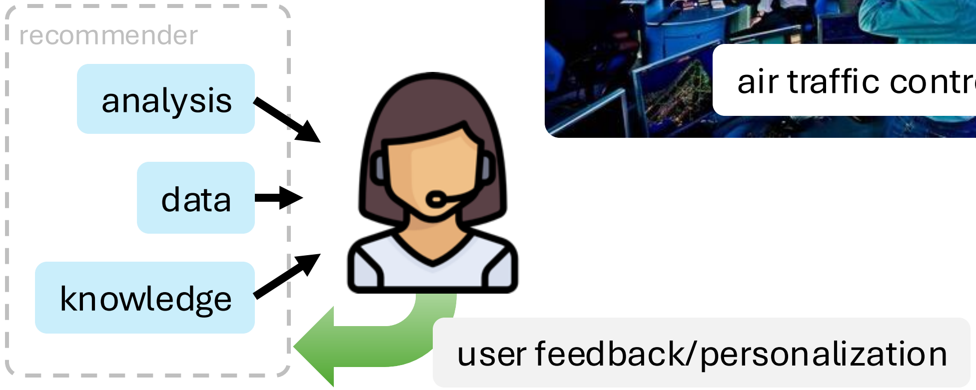

disaster response

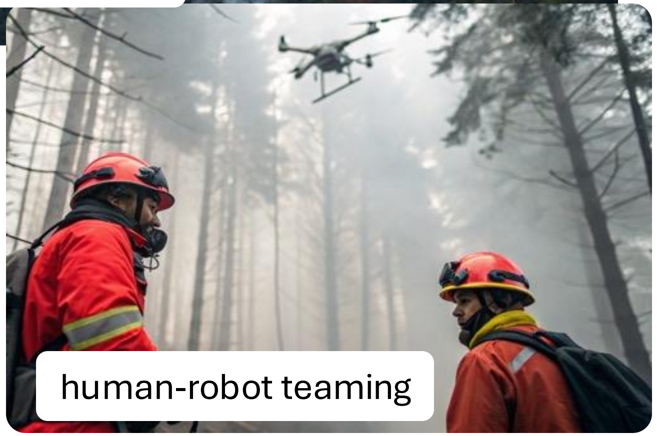

Cognitively-Aware Autonomy

Demo: design a cognitively-aware intelligent tutoring system

intelligently adjusts tutor to the learner

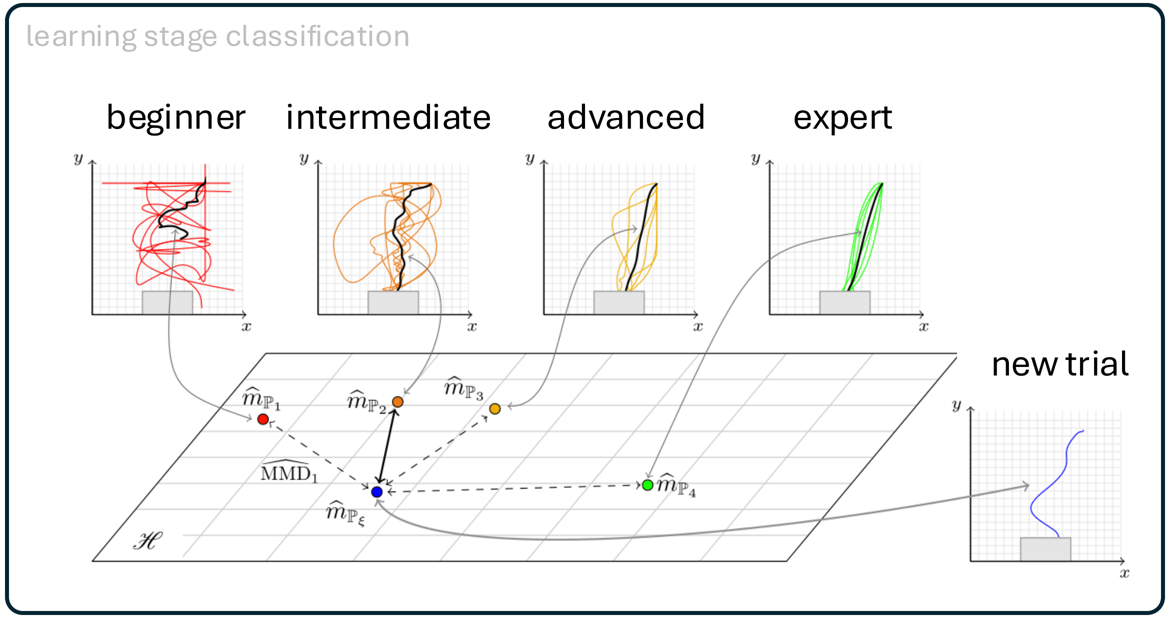

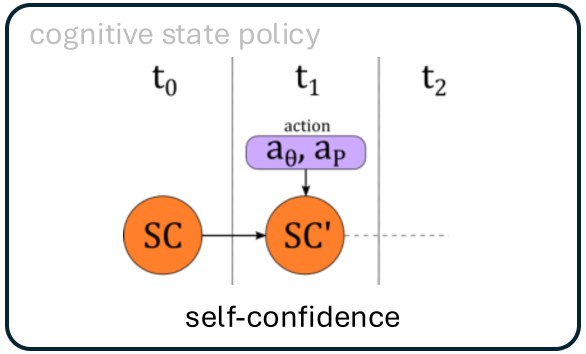

THE UNIVERSITY OF NEW MEXICO · PURDUE UNIVERSITY · TEXAS The University of Texas at Austin · University of Colorado Boulder · Raytheon Technologies

M. S. Yuh, E. Rabb, A. Thorpe and N. Jain. (2024). Using Reward Shaping to Train Cognitive-Based Control Policies for Intelligent Tutoring Systems
Ortiz, K. R., Hunter, J. G., Thorpe, A. J., Yuh, M., Reid, T., Jain, N., & Oishi, M. (2024). Assessing the Relationship Between Learning Stages and Prefrontal Cortex Activation in a Psychomotor Task

Teaching & Mentoring

- meeting student needs through **personalized support** and frequent, individualized feedback
- **focusing on problems and projects** to ground theoretical concepts in practical problems
- designing **inclusive classroom environments**

enhance existing courses with modern theory (e.g. connections between controls & learning)

learning-based control

computational

Outreach & Service

explicit goals, measurable impacts, long-term initiatives

- early research opportunities, e.g. undergraduate internships, REU programs
- improve access through outreach, e.g. Code2College, Texas Advanced Computing Center REU, NSF Summer Intensive Research Initiative

- commitment to service
- cultivating a culture of belonging
- building mentorship networks, e.g. STEM Muse Mentorship Program

Summer Intensive Research Initiative


STEM Muse Kickoff

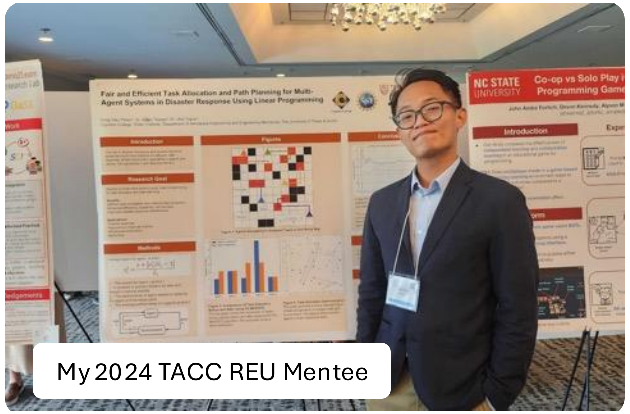

My 2024 TACC REU Mentee

My research: learning for autonomy

Learning Structured Representations

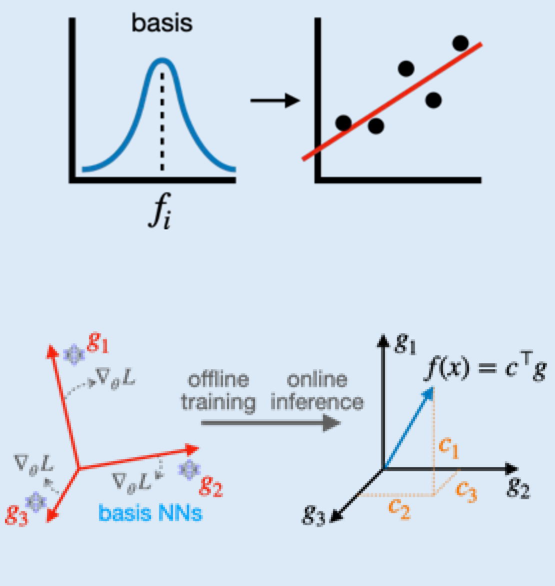

Function Encoders

Adaptable Autonomy & Inductive Transfer

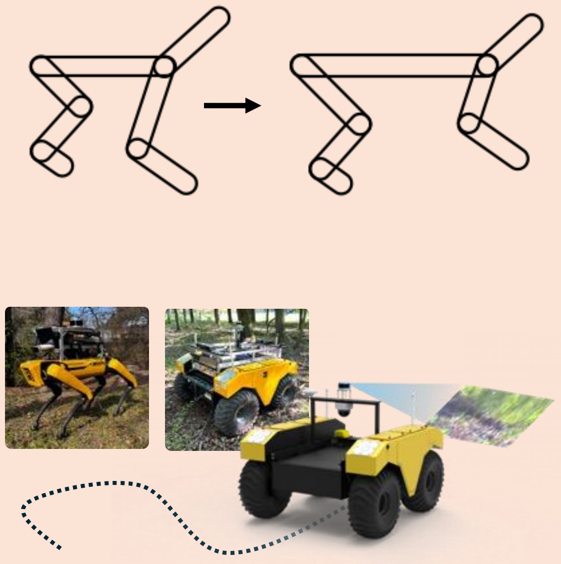

Zero-Shot Modeling & Control

Neural Operator Learning

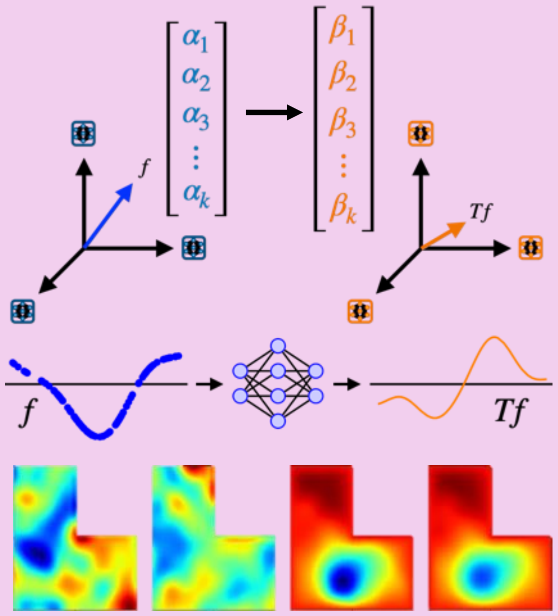

Basis to Basis Operators

Human-Autonomy Interaction

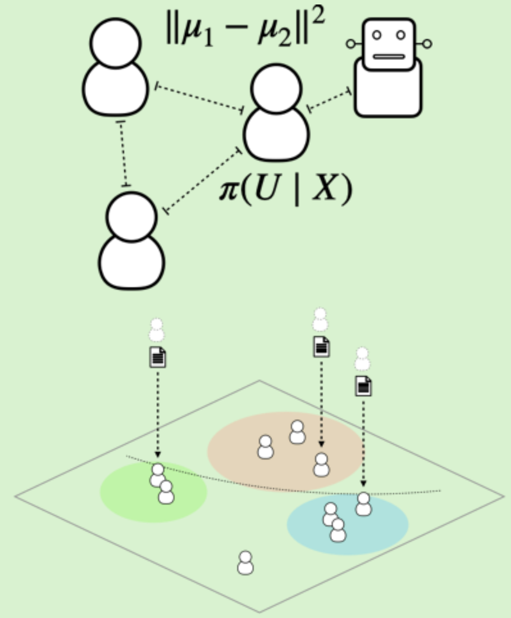

Characterizing Heterogeneity

Contact: adam.thorpe@austin.utexas.edu