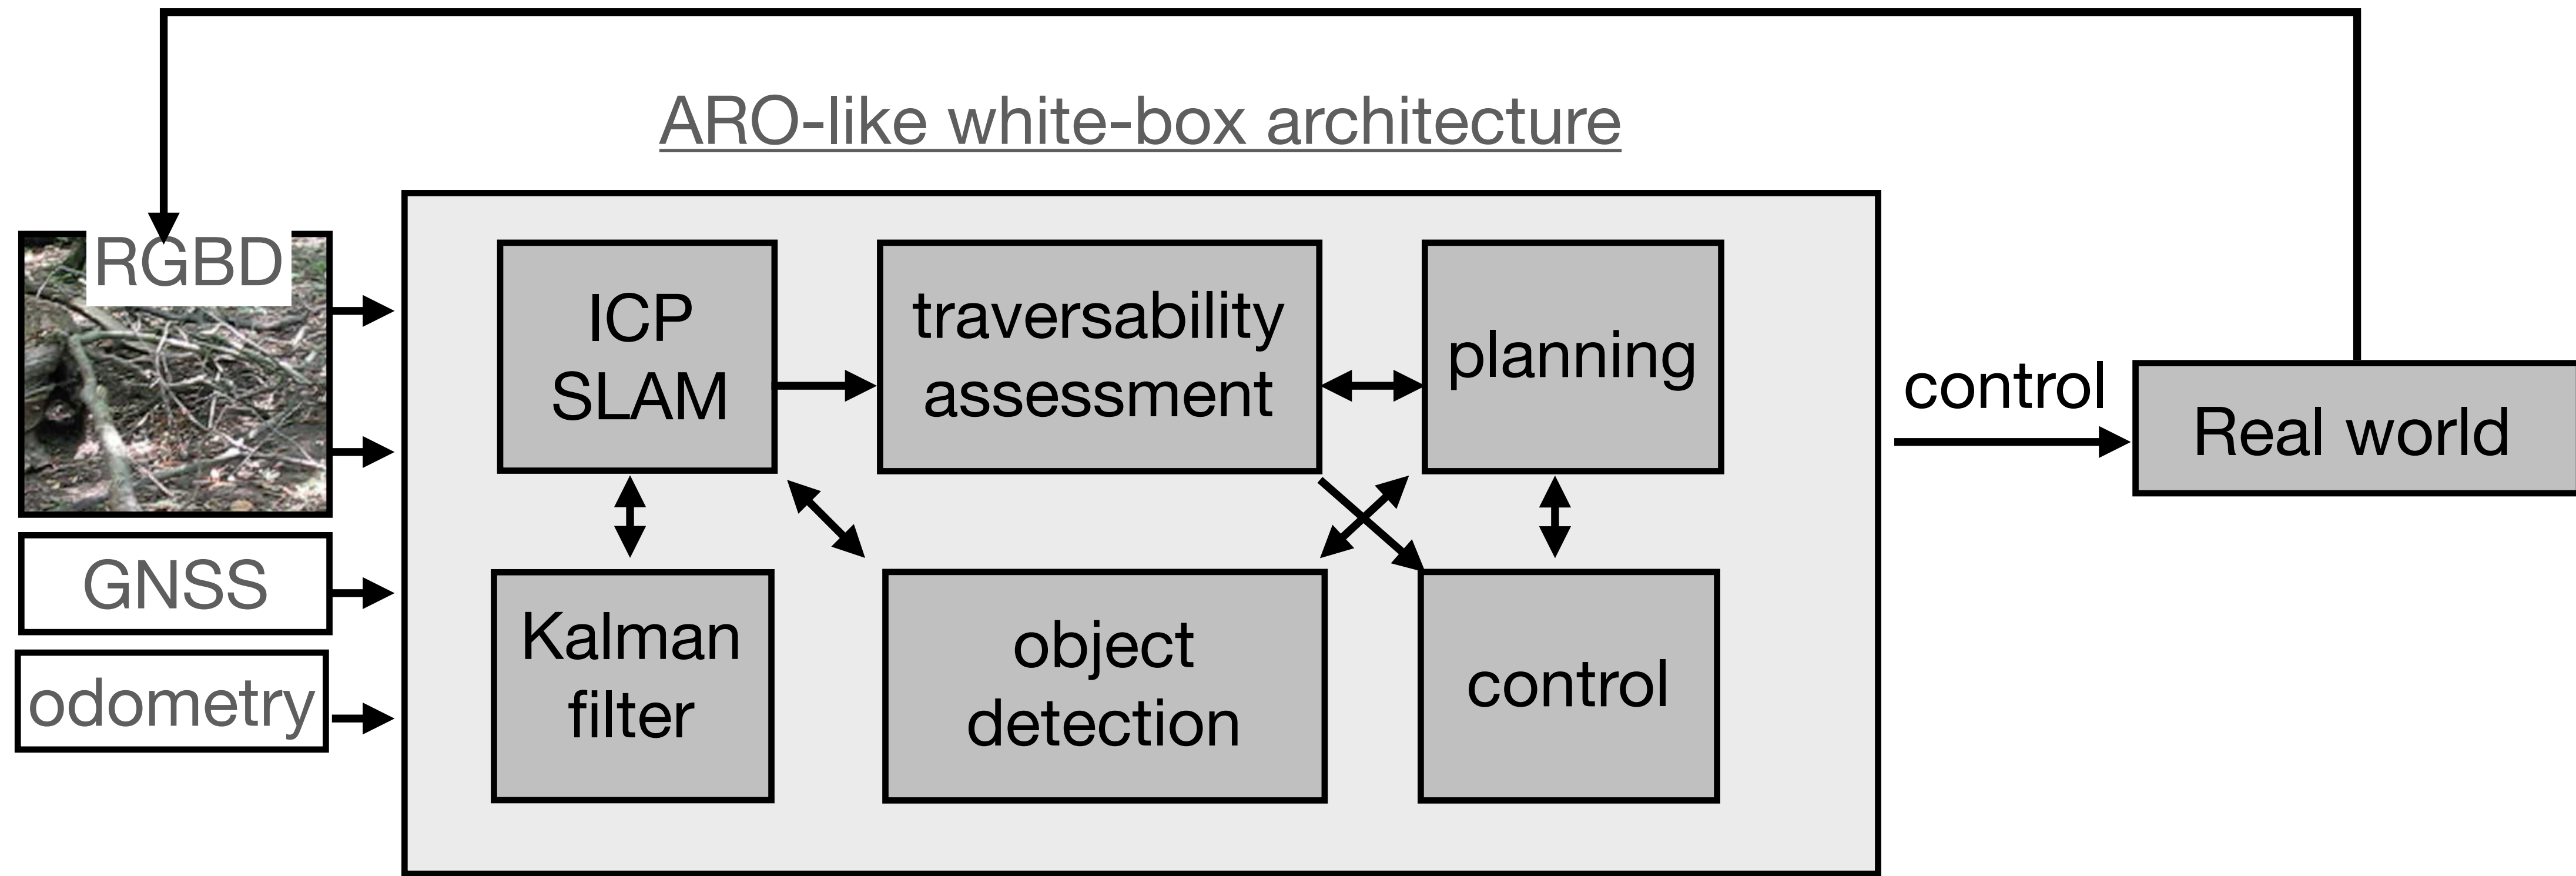



# **Take home message ARO**

**Karel Zimmermann**

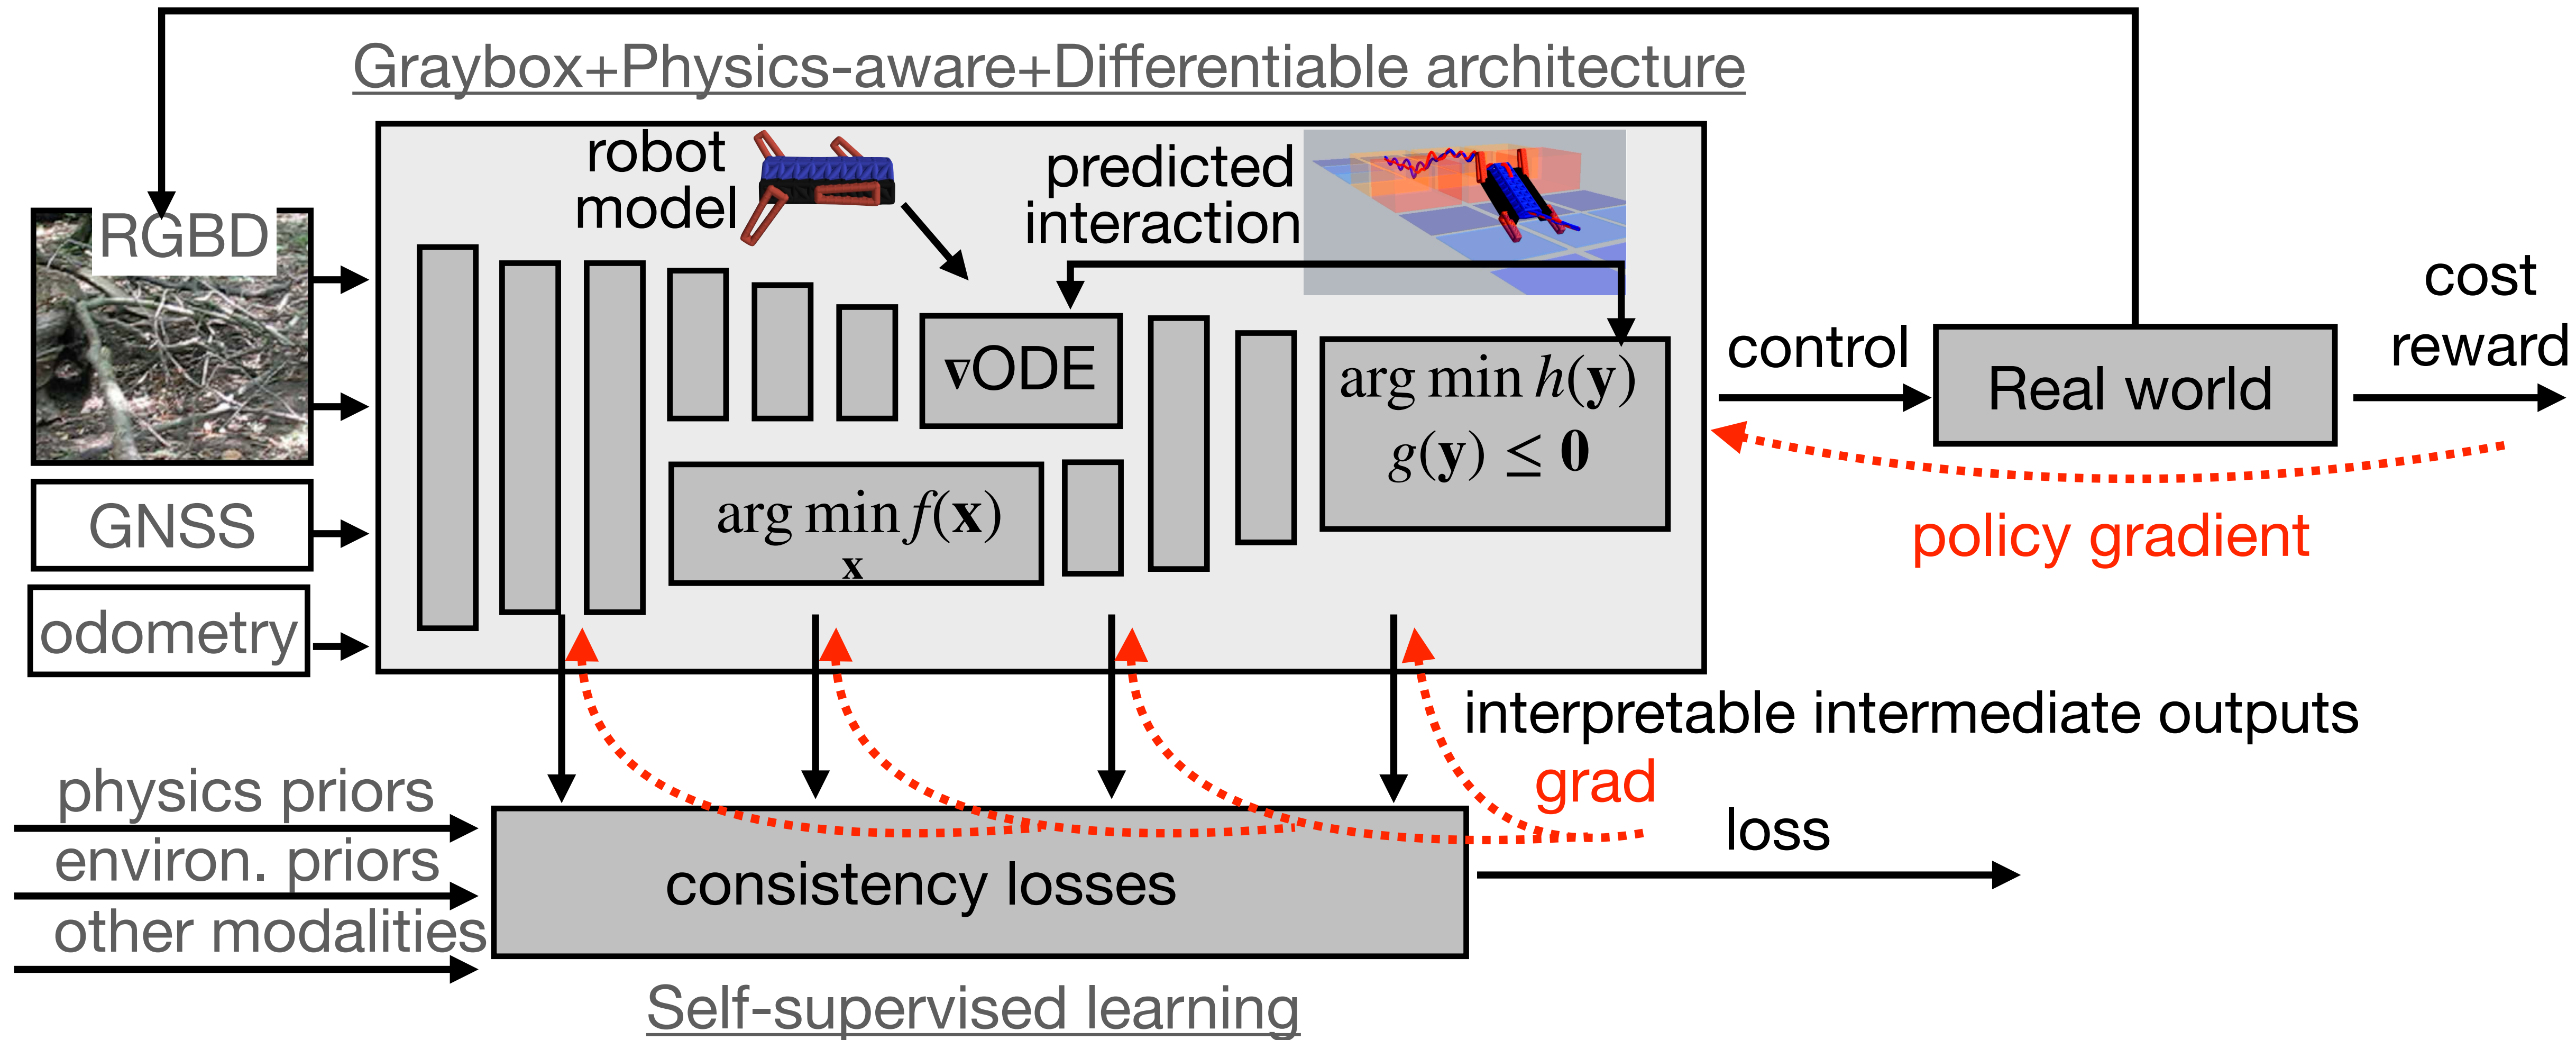


- Each component designed/learned/tuned independently:
  - huge amount of human-annotated data/human-feedbacks  
=> tuning on other domains is extremely expensive
  - domain-specific data => transfer impossible



- Allows to learn from cost/reward function via policy gradient, but it suffers from:
  - Sample inefficiency => Often trained in simulation => Simulation bias
  - Black-box architecture => Poor generalization beyond training environment
  - Physical embodiment ignored => Impossibility to transfer on other robots
  - Hard to construct dense self-evaluating cost/reward function

# Longterm goal: Explainable + self-learnable architecture





- Self-learning through “end-to-end differentiability” + “explainability” + “priors”
- Interesting issues:
  - ◆ backpropagation through non-trivial components
  - ◆ construction of consistency losses



# Backpropagation through optimization problem


$$\omega^* = \arg \min_{\omega} \|\text{SGD: } \arg \min_{\mathbf{x}} f_{\omega}(\mathbf{x})\|^2$$

$\mathbf{x}_0$  ... initial point  
 $\omega$  ... criterion params  
 $\alpha, \beta$  ... optimizer params






# Backpropagation through optimization problem



PyTorch:

```
for i in range(iter):  
    x = x - torch.autograd.grad(f, x, retain_graph=True)  
loss = x.norm()  
torch.autograd.grad(loss, omega)
```




Differentiable MPC control, [Amos et al. NIPS 2018]

<https://arxiv.org/abs/1810.13400>

Cost

System  
Dynamics

Initial State




QP Iterate  $i$

$$\tau_{1:T}^i = \operatorname{argmin}_{\tau_{1:T}} \sum_t \tilde{C}_\theta^i(\tau_t)$$


subject to  $x_1 = x_{init}$

$$x_{t+1} = \tilde{f}_\theta^i(\tau_t)$$

$$\underline{u} \leq u \leq \bar{u}$$



Optimal actions  
to take next





# PyTorch embedded modeling language for convex optimization problems

<https://www.cvxpy.org/>

[Boyd, et al, NIPS, 2019]

```
cvxpylayer = CvxpyLayer(problem, parameters=[A, b], variables=[x])  
solution, = cvxpylayer(A, b) # feed-forward pass (solve the problem)  
solution.sum().backward() # backward pass (how A,b influence solution)
```

The logo for CVXpy, featuring the word "CVX" in a stylized font with a curved line underneath, followed by "py".

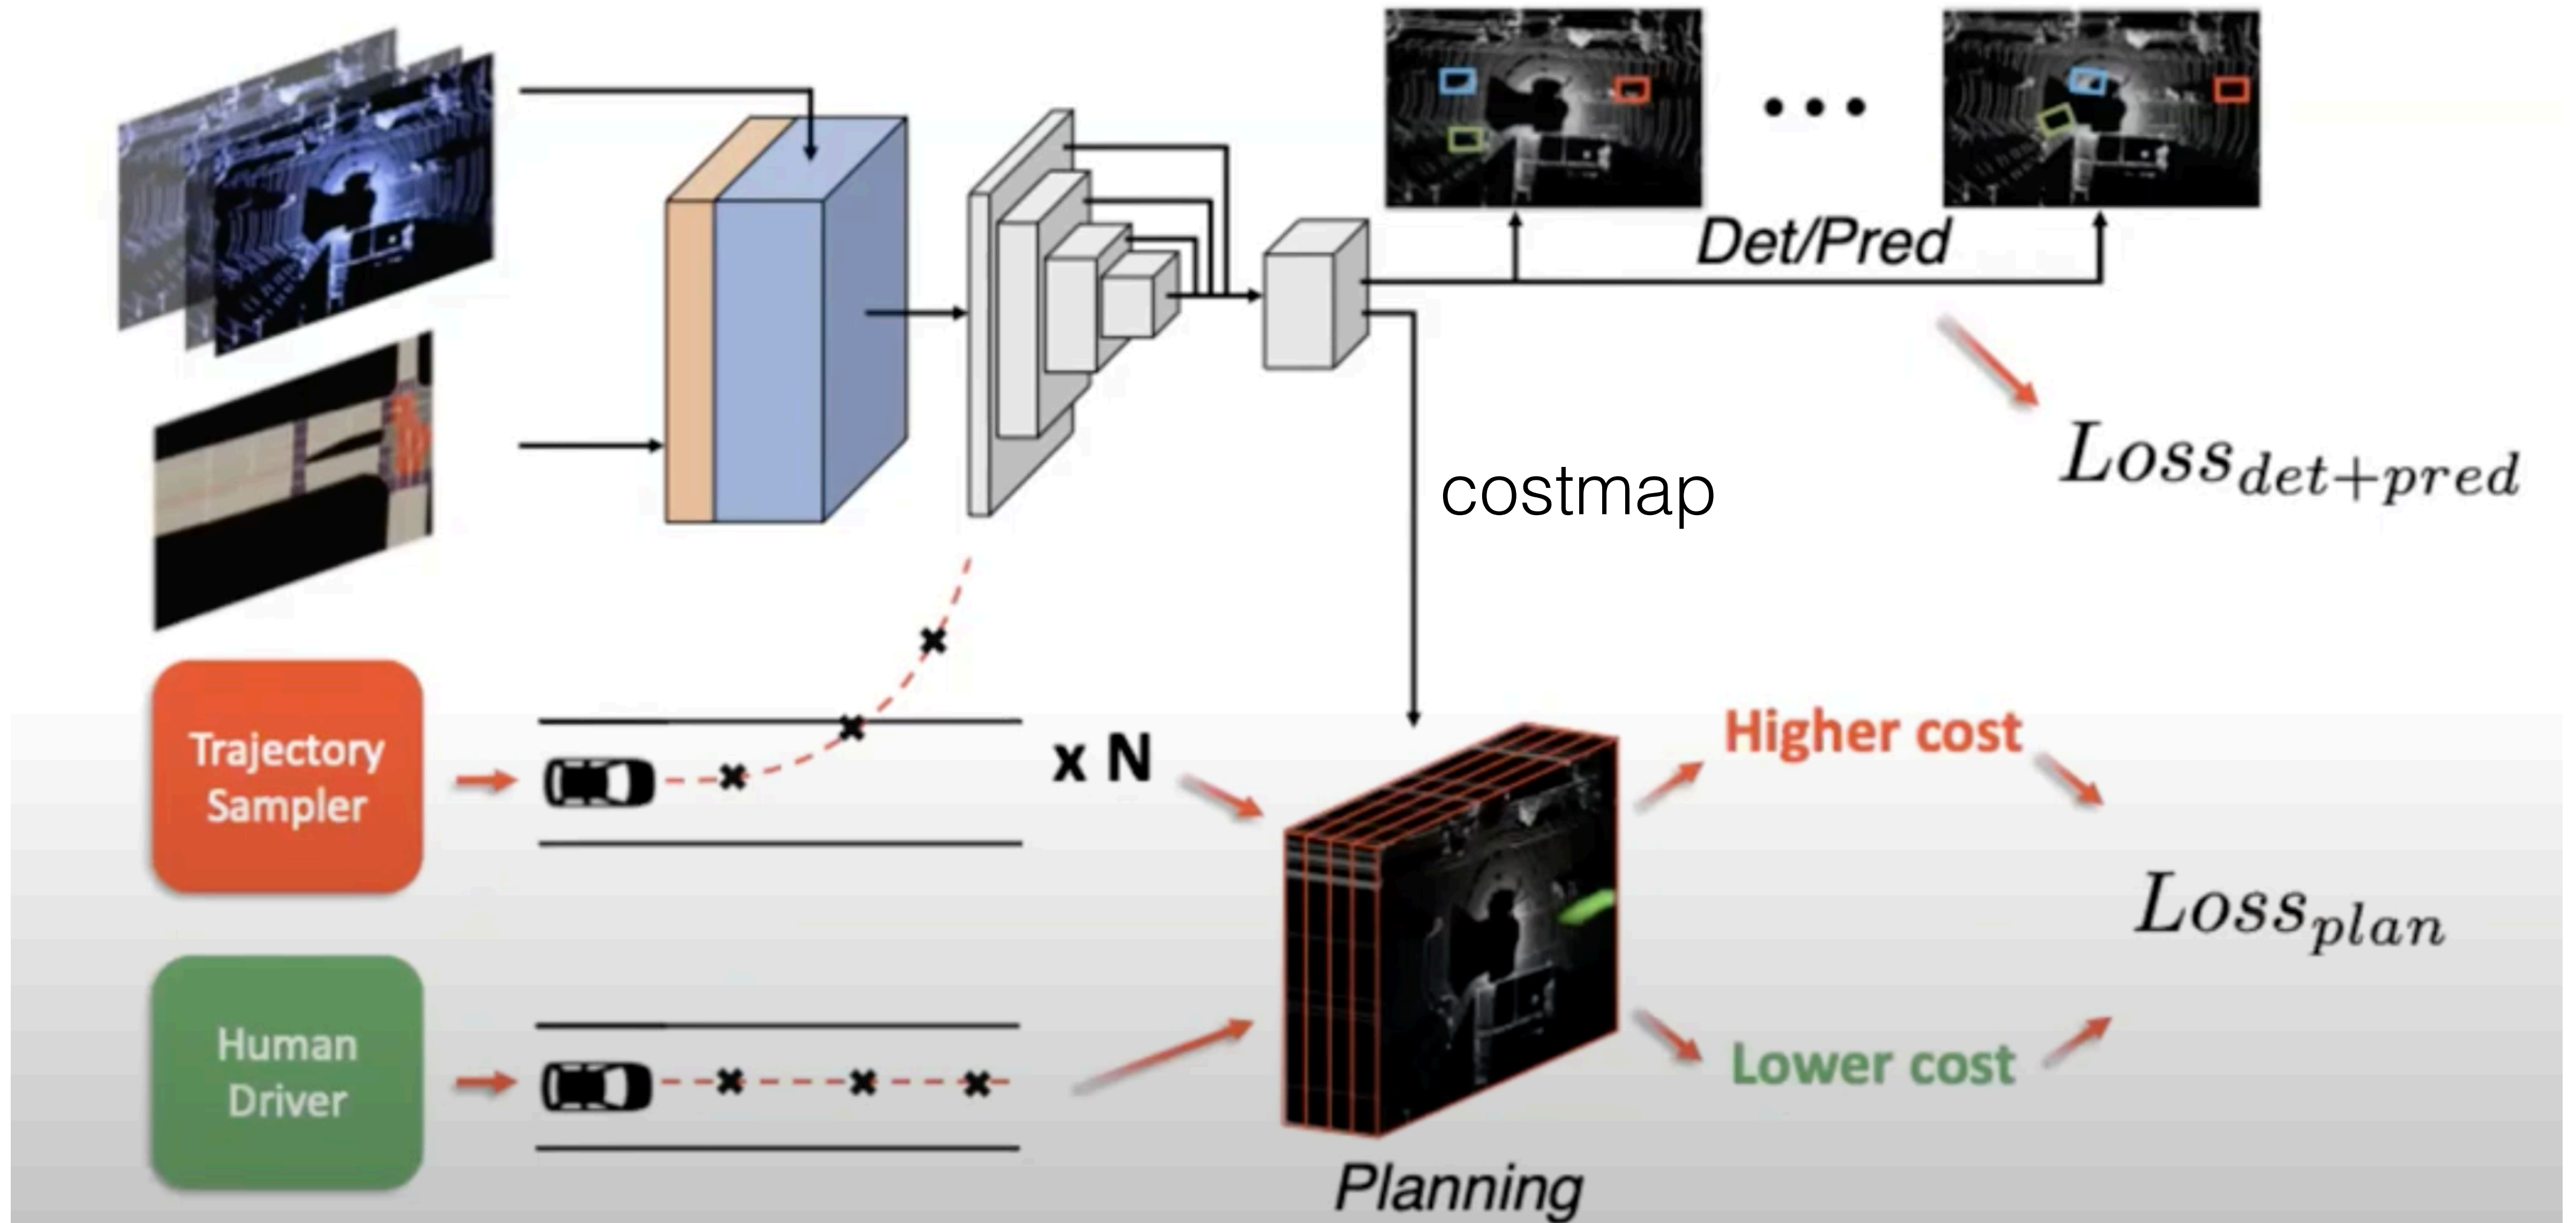
$$x^*(\theta) = \underset{x}{\operatorname{argmin}} f(x; \theta)$$

subject to  $g(x; \theta) \leq 0$   
 $h(x; \theta) = 0$

The PyTorch logo, consisting of a red flame-like icon and the text "PyTorch".

TensorFlow


Interpretable motion planning  
[Zeng,.. Urtasun, Uber, CVPR, 2019]  
<http://www.cs.toronto.edu/~wenjie/>





# Grad SLAM [Murthy, ICRA, 2021]

<https://gradslam.github.io/>



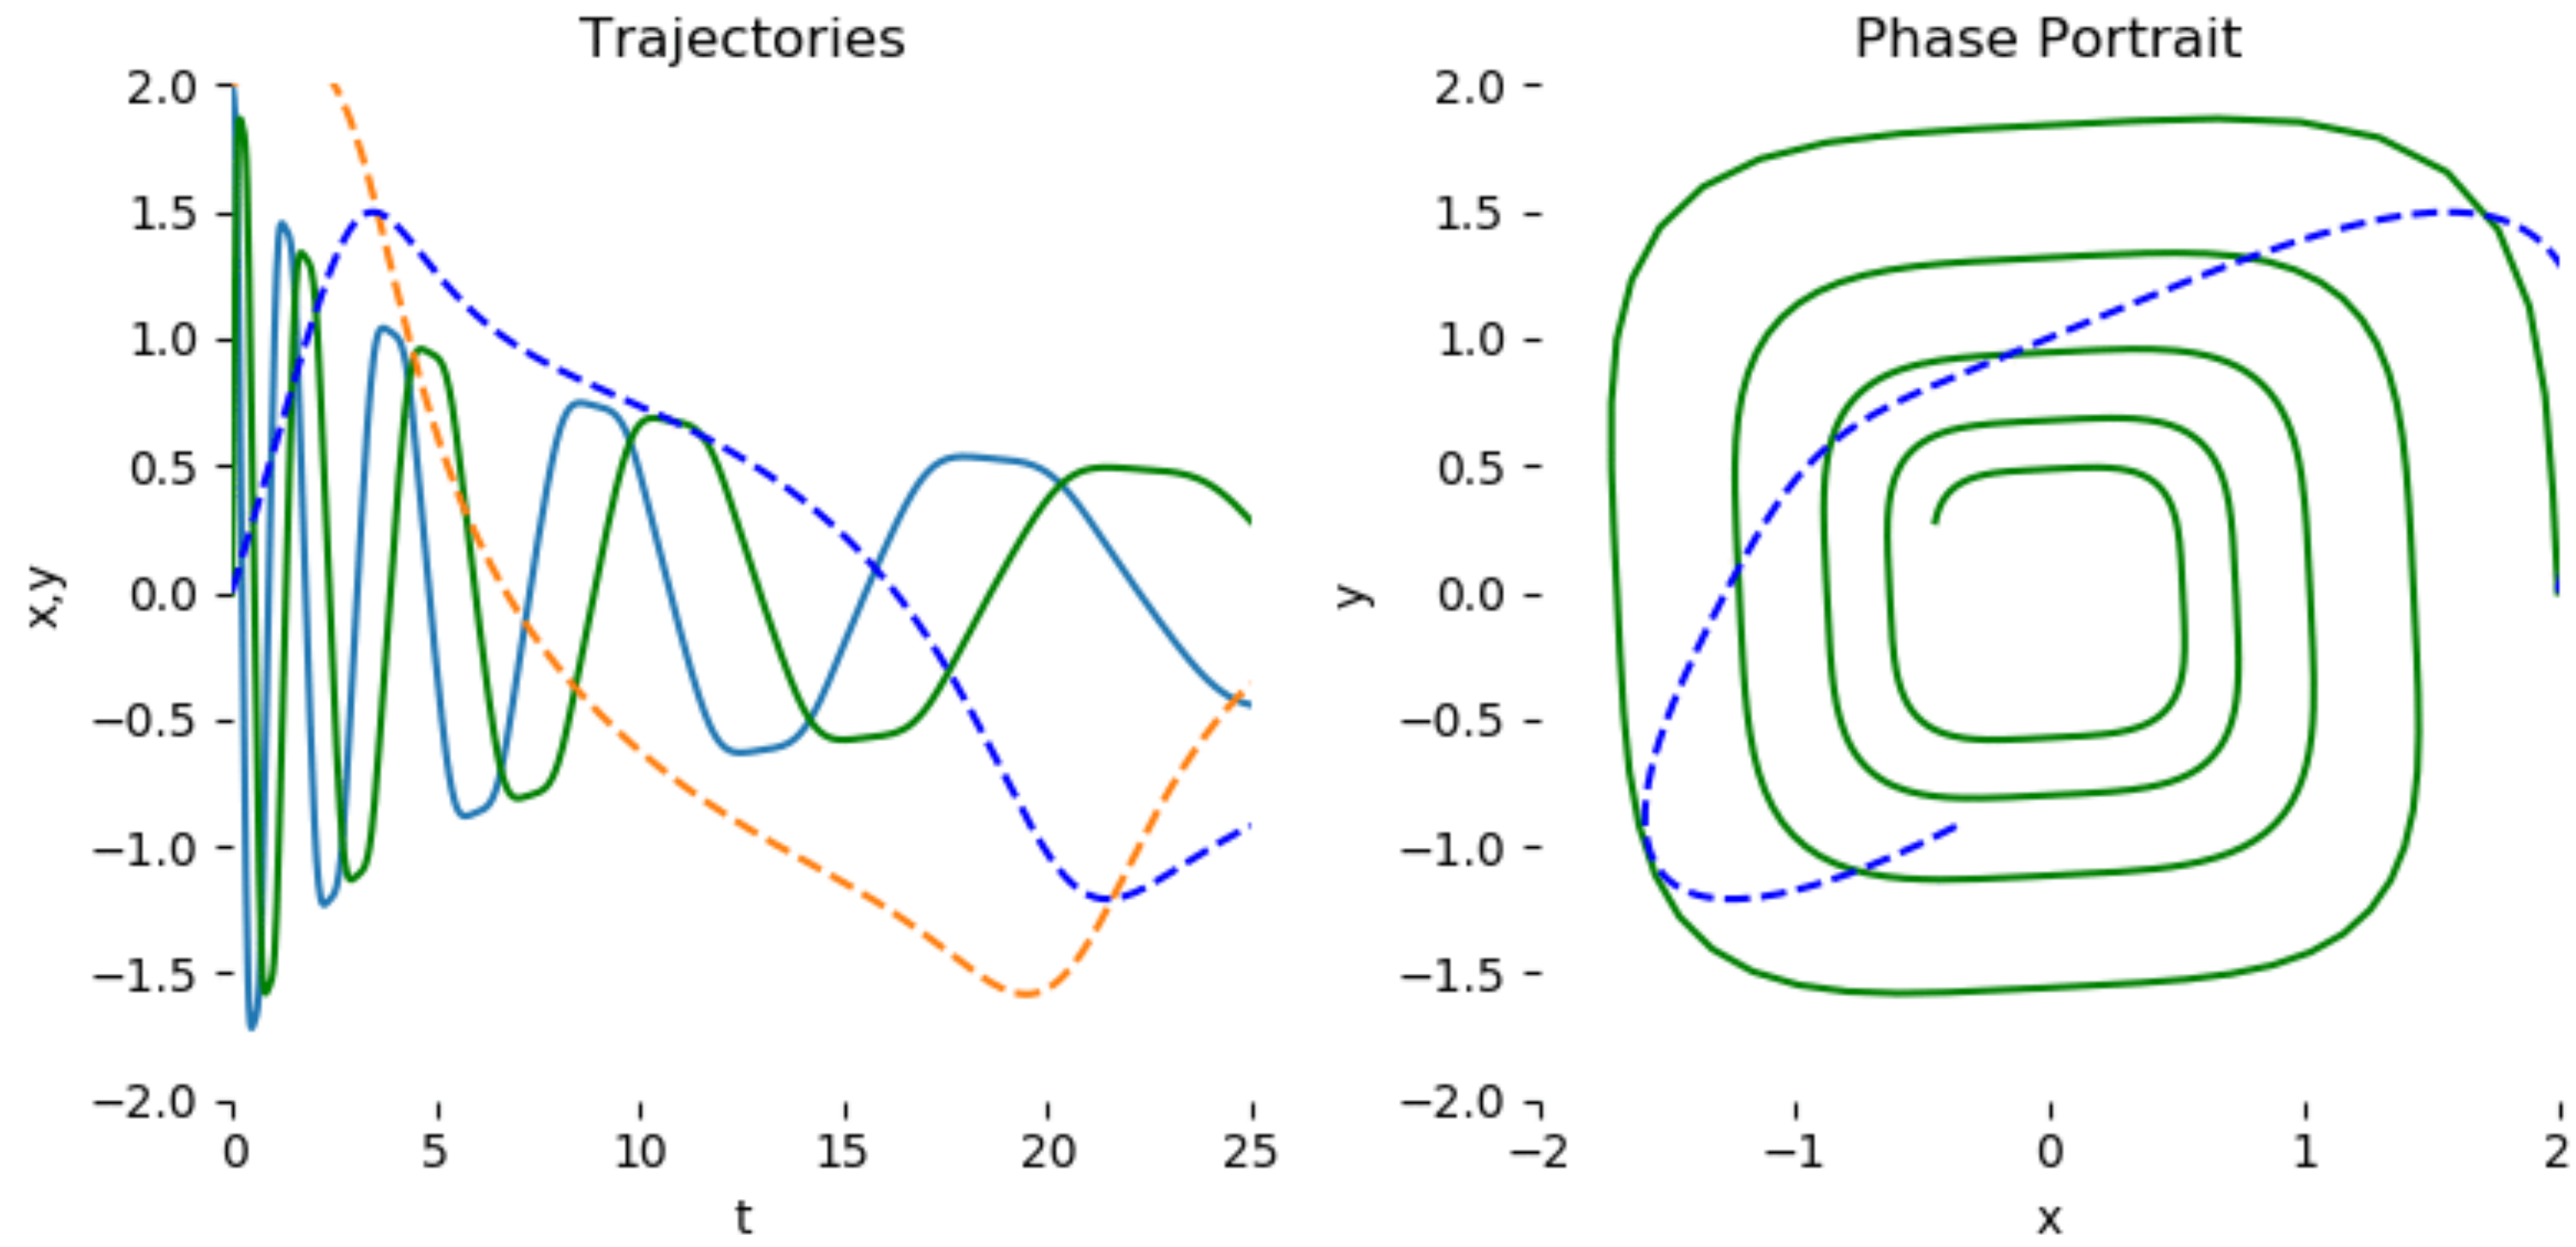
map:

pose:  $\Phi$

$$\frac{\partial \text{SLAM}(\mathbf{x}_1, \dots, \mathbf{x}_t)}{\partial \mathbf{x}_1, \dots, \mathbf{x}_t} = ?$$


# Pytorch differentiable ODE solver

<https://github.com/rtqichen/torchdiffeq>



# Google's BRAX/ Nvidia's WARP - differentiable physics engine

<https://github.com/google/brax>

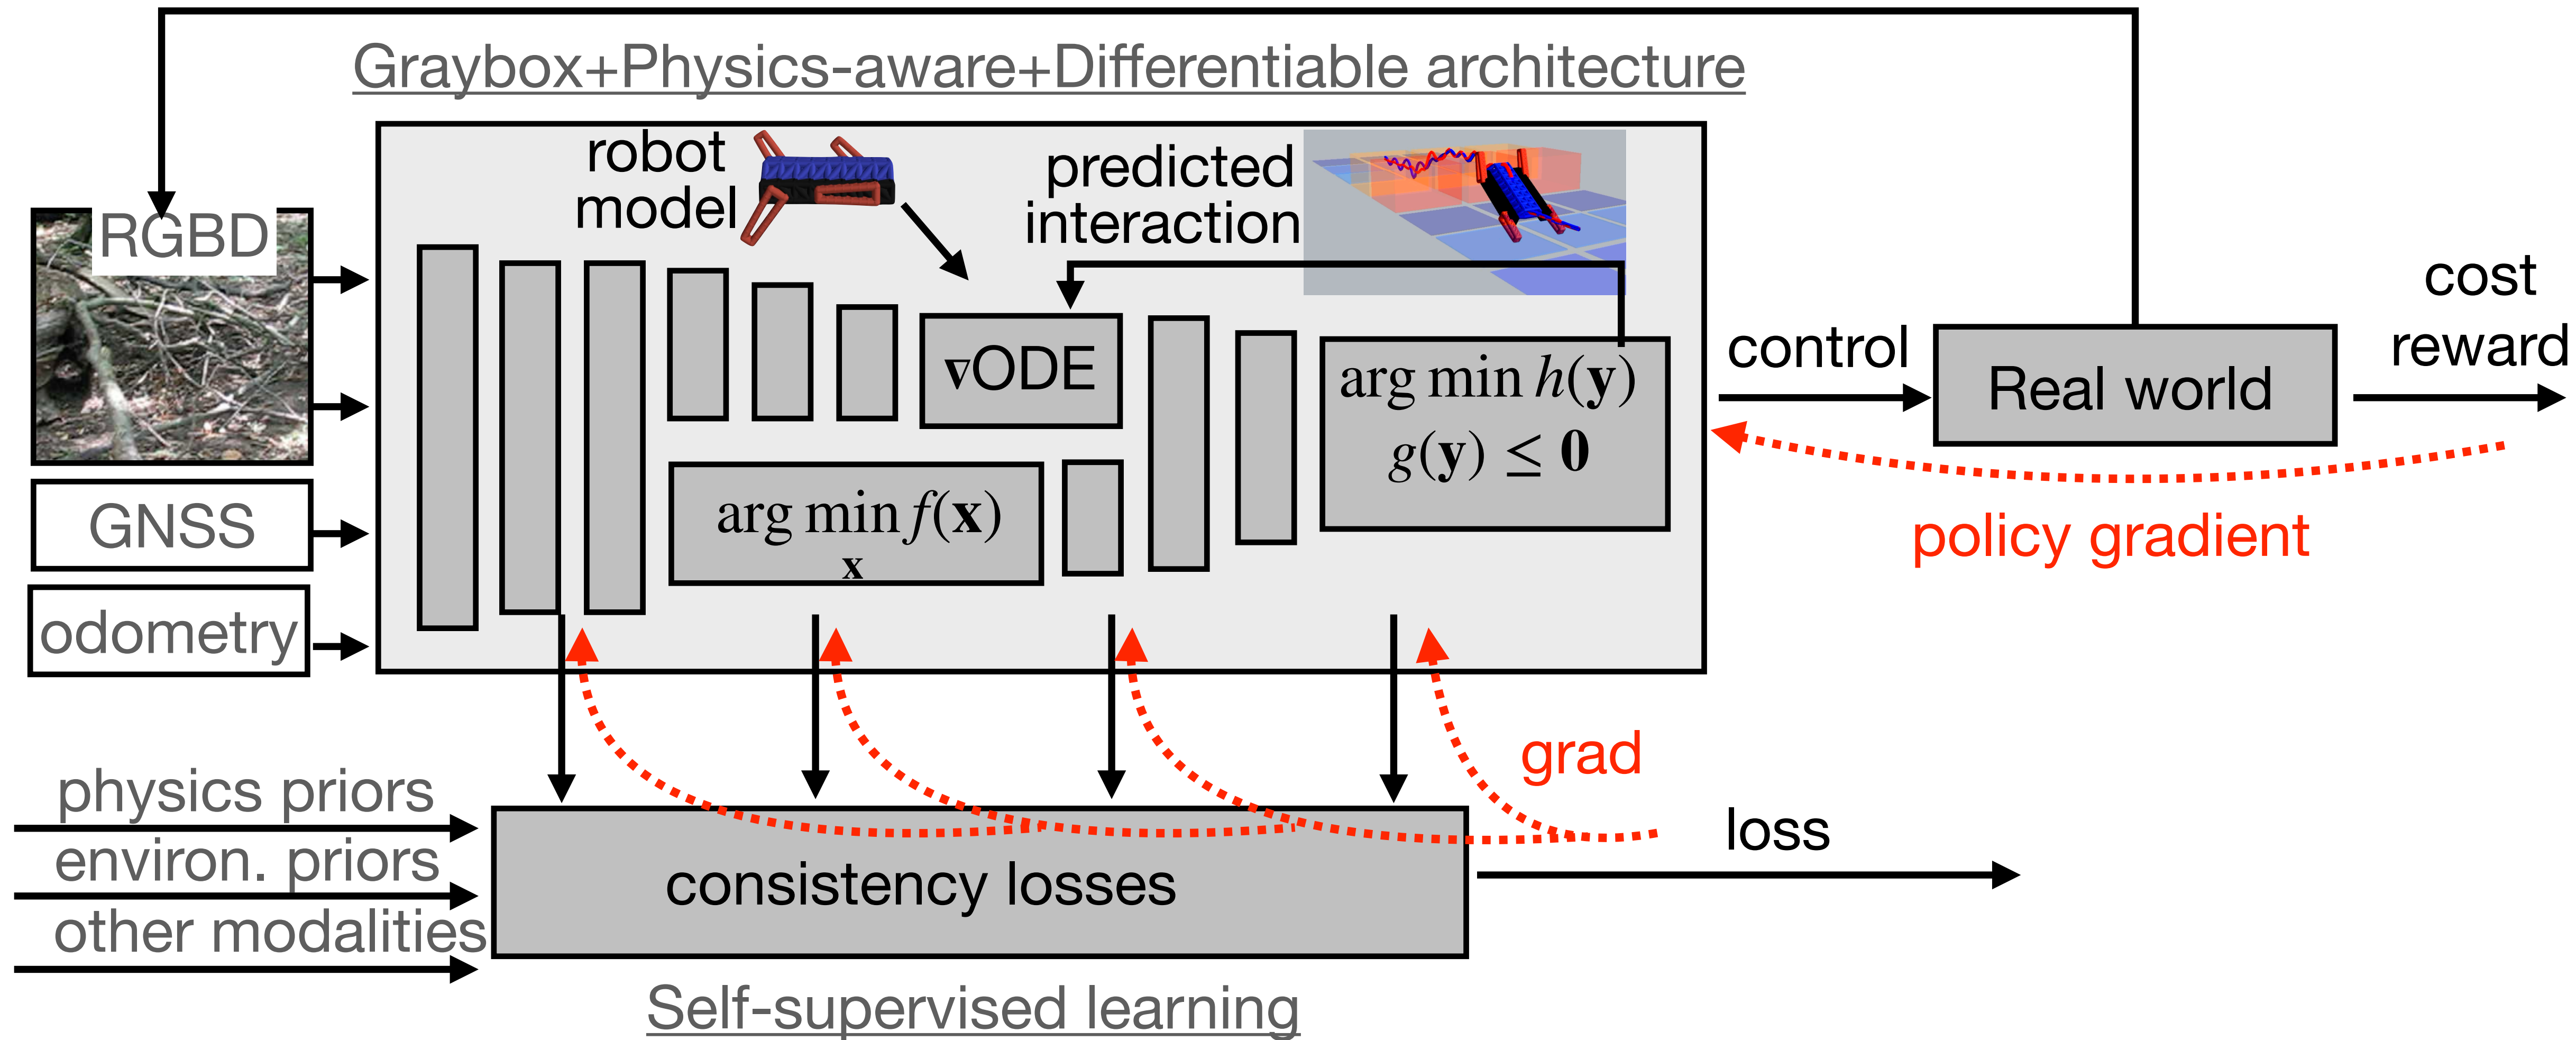


Brax simulates these environments at millions of physics steps per second on TPU

 **BRAX**




# Longterm goal: Explainable + self-learnable architecture



- Self-learning through “end-to-end differentiability” + “explainability” + “priors”
- Interesting issues:
  - ◆ backpropagation through non-trivial components
  - ◆ construction of consistency losses

# Longterm goal: Explainable + self-learnable architecture



- Self-learning through “end-to-end differentiability” + “explainability” + “priors”
- Interesting issues:
  - ◆ backpropagation through non-trivial components
  - ◆ construction of consistency losses




## Real robots

Traversability assessment: ability to predict robots behaviour on complex terrains perceived by exteroceptive onboard sensors such as camera or lidar.





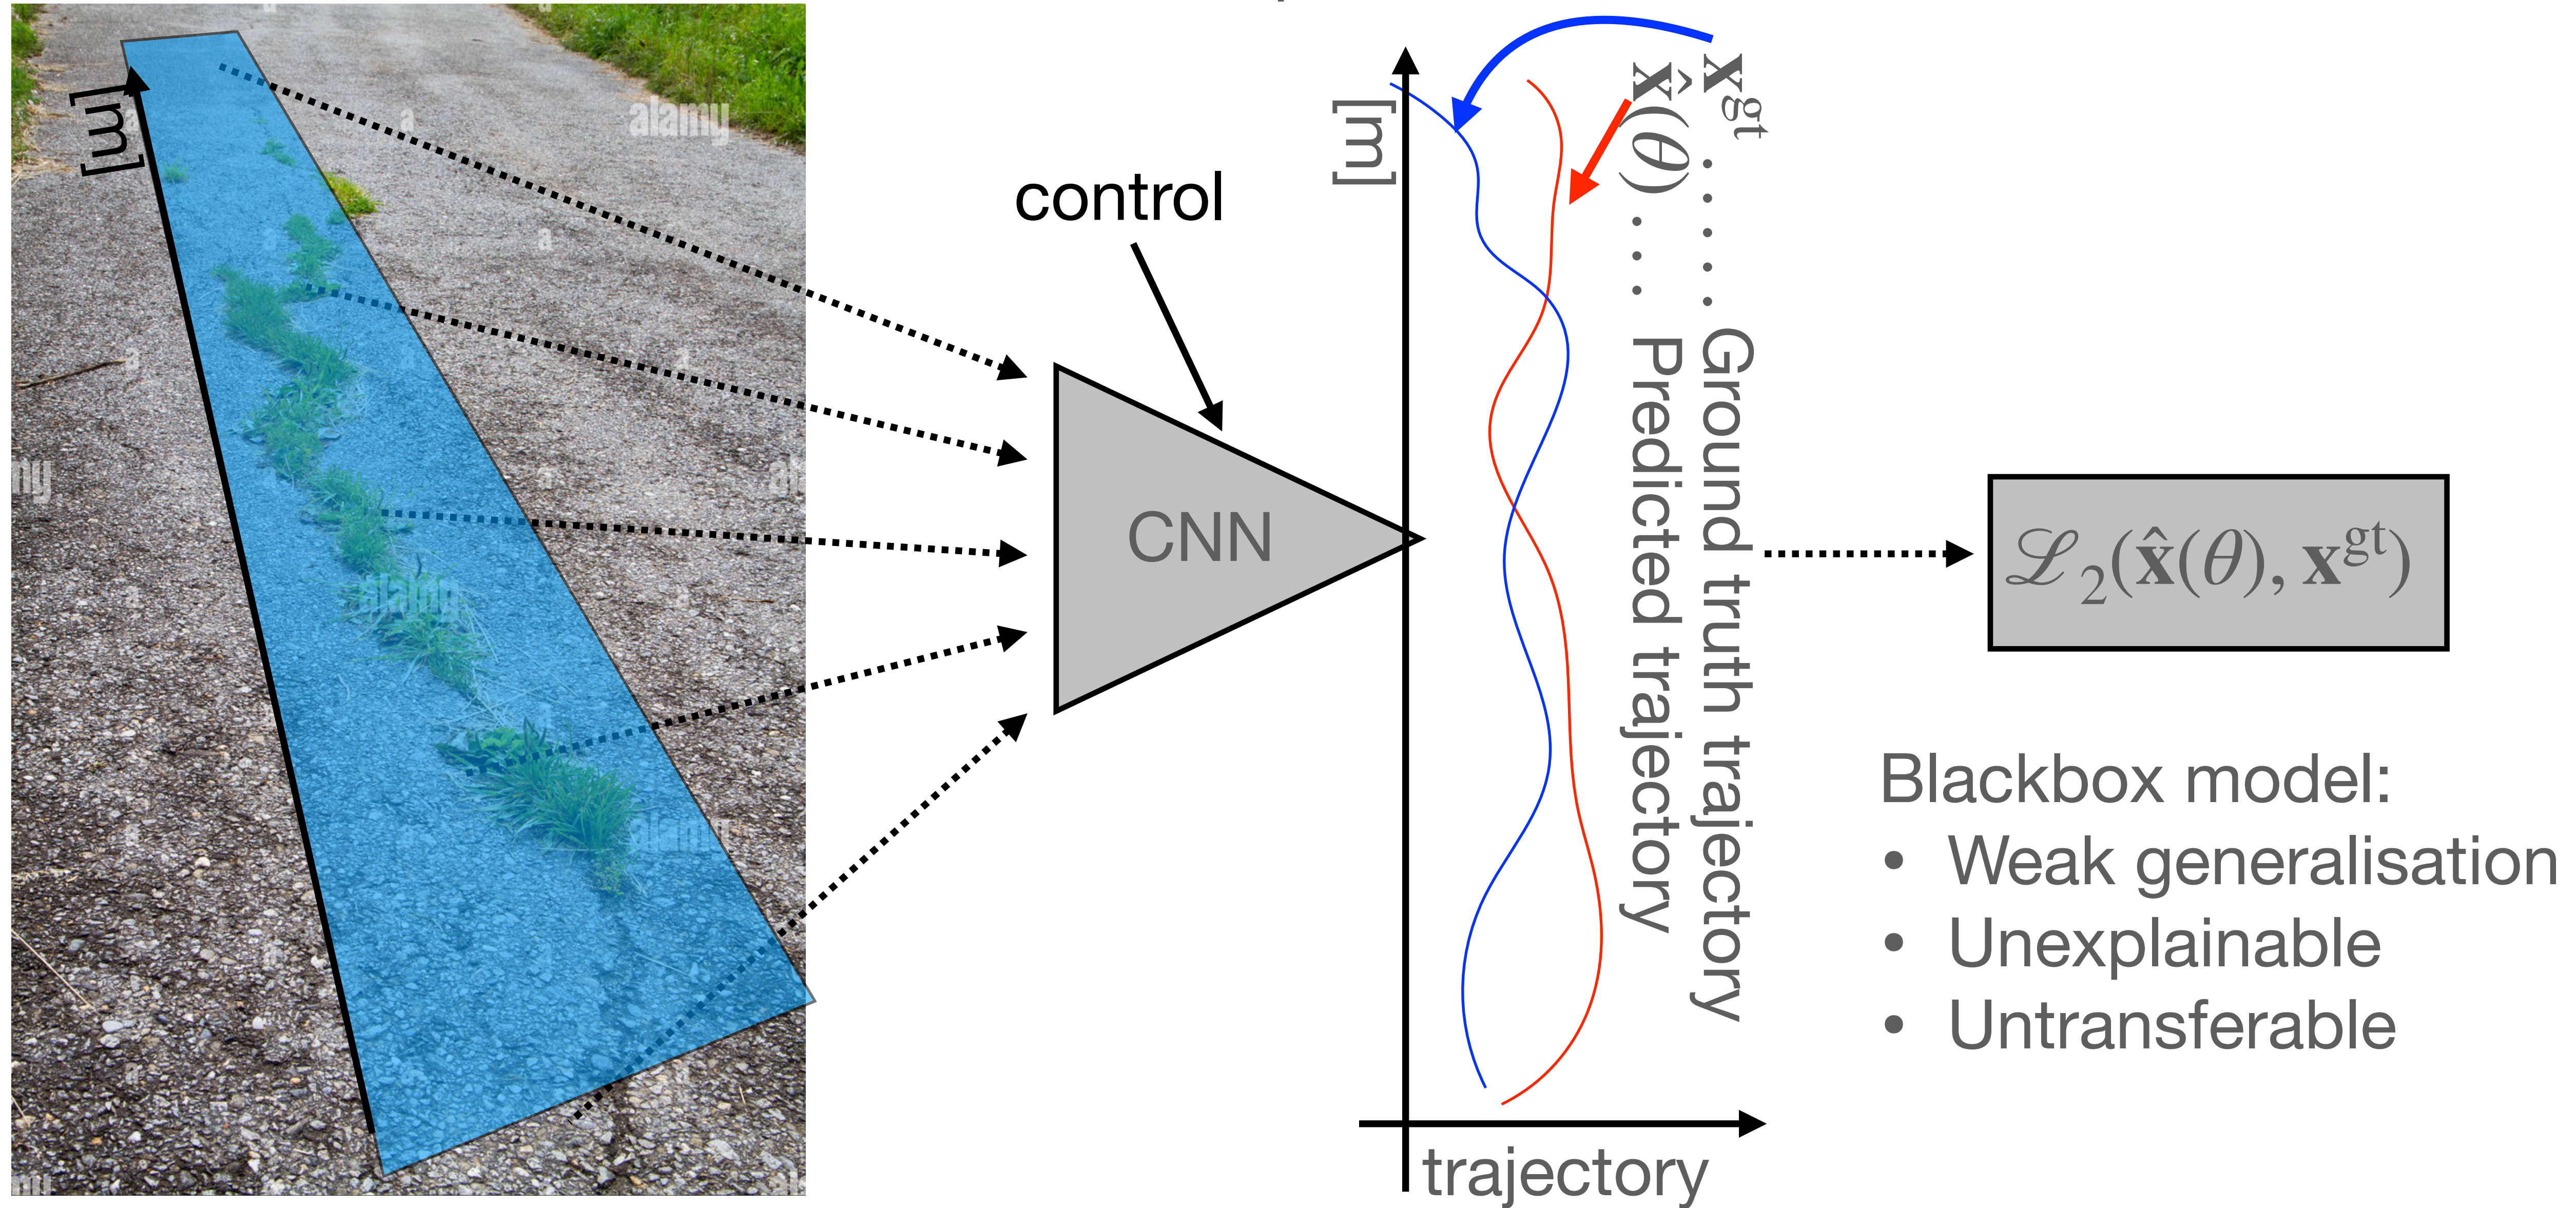
# Learning to predict robot-terrain interaction from RGB(D) Trivial self-supervision





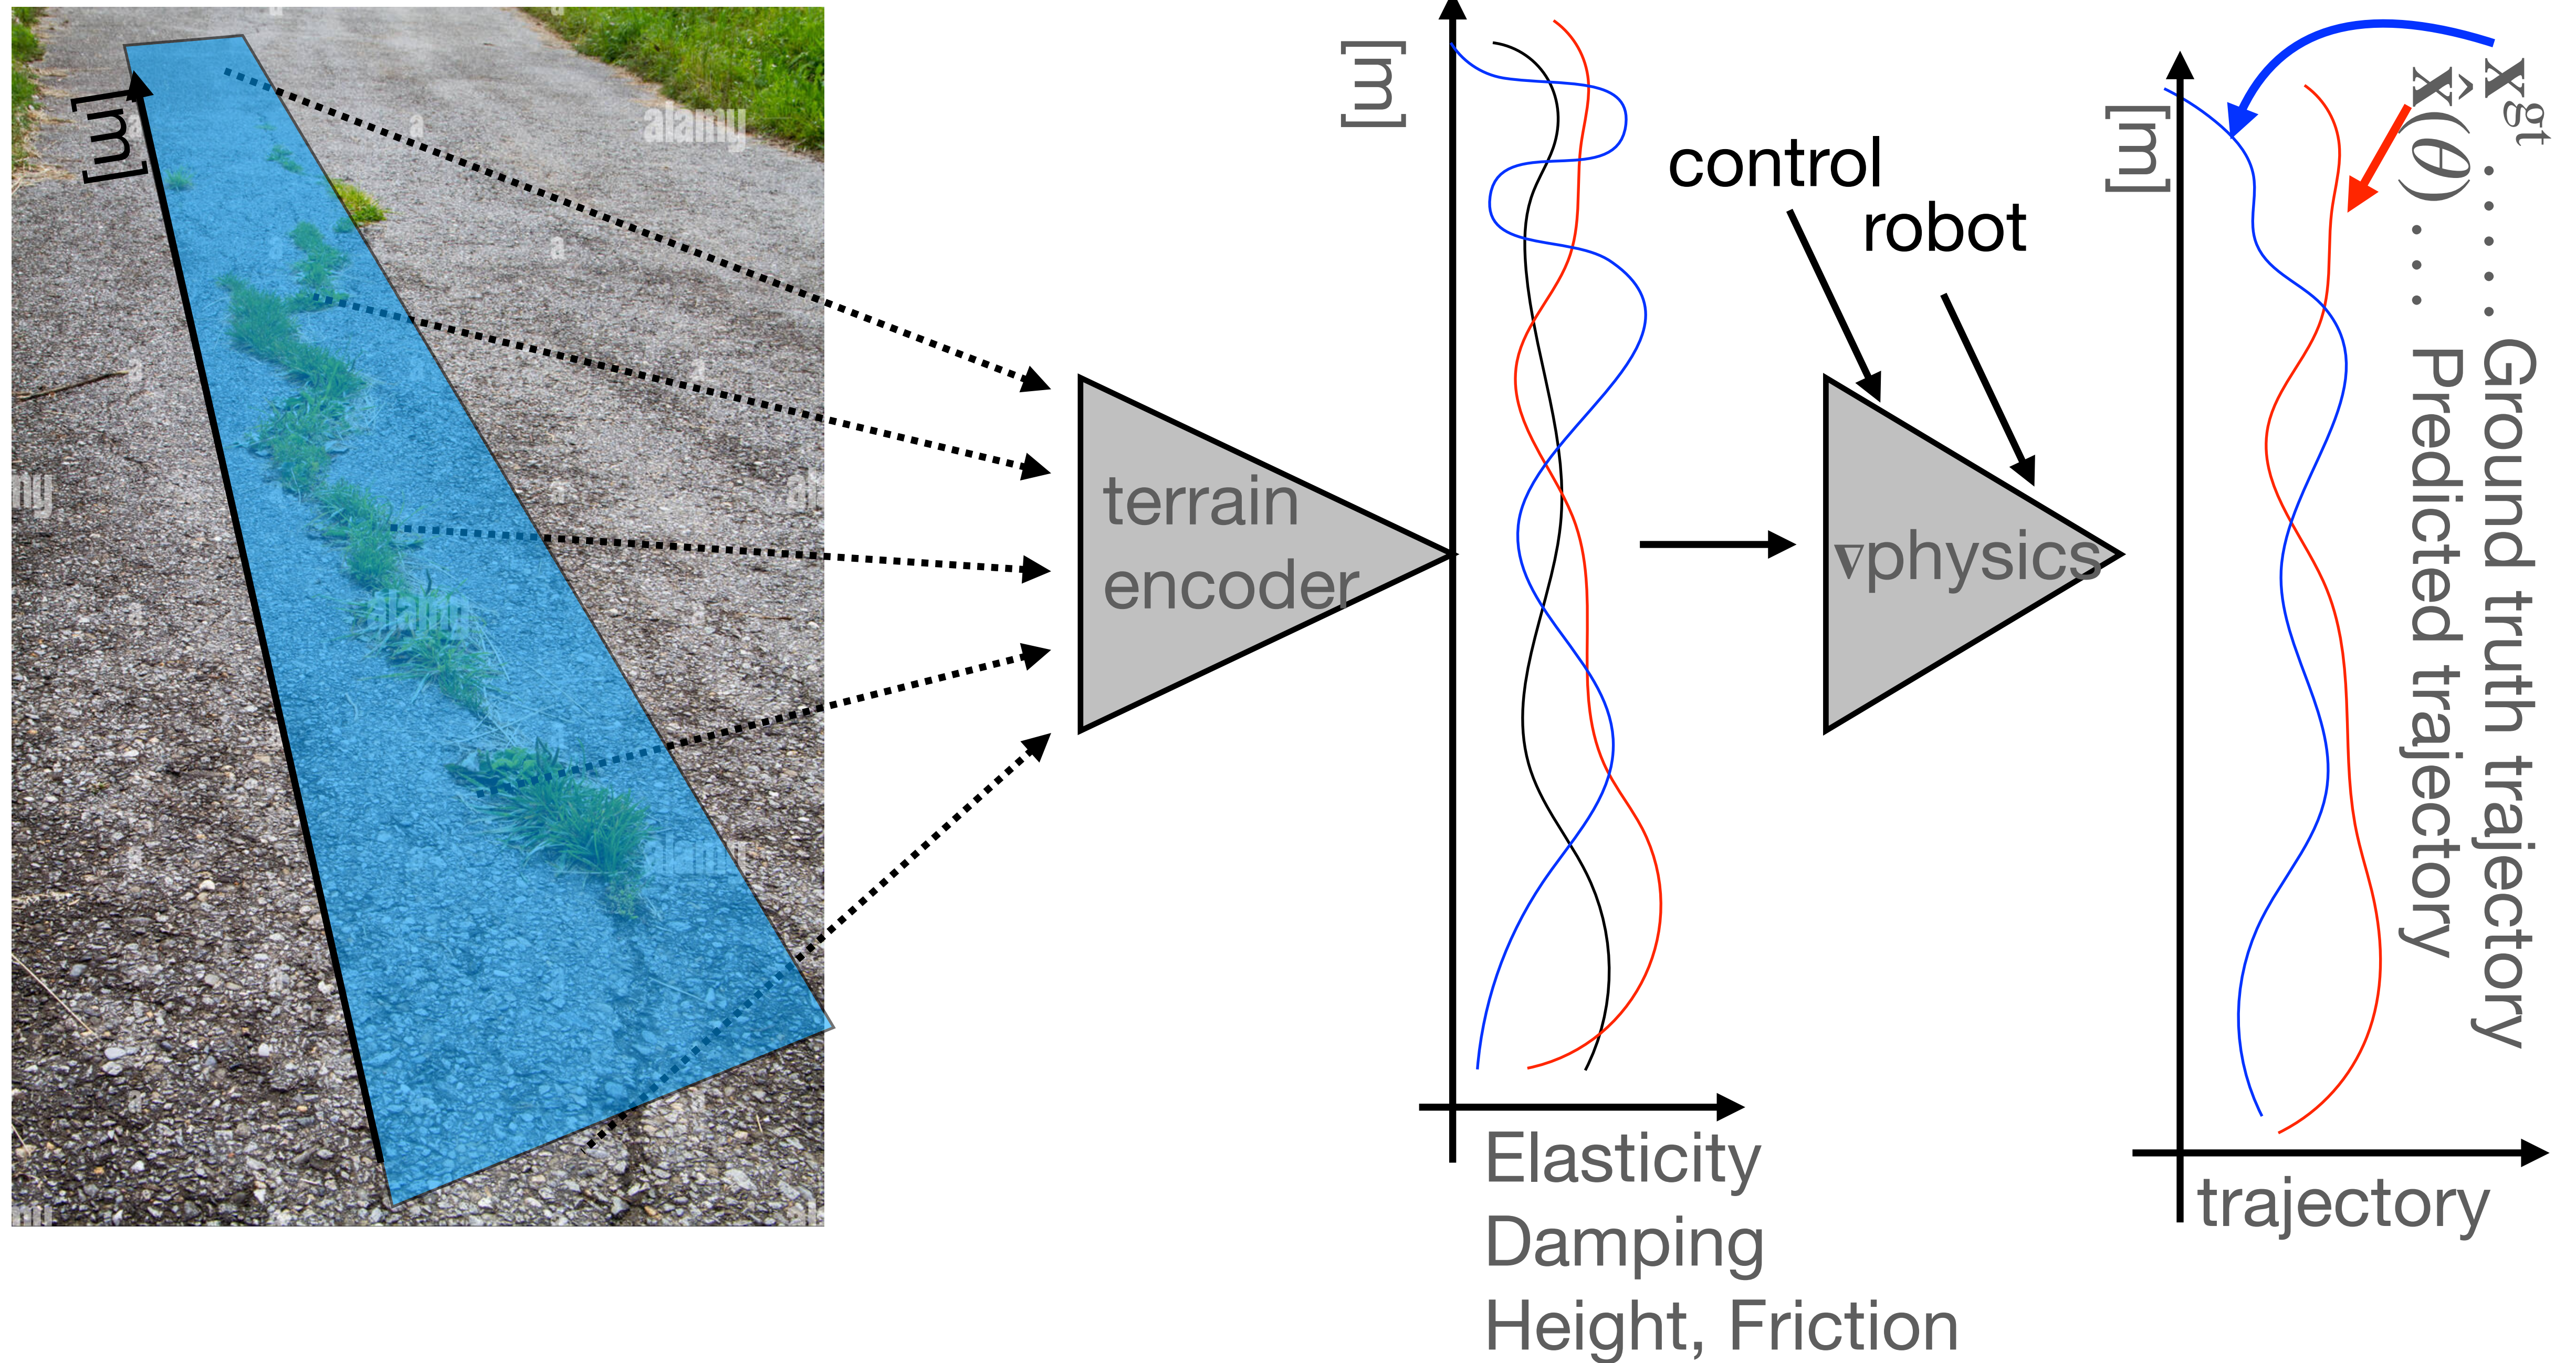
# Learning to predict robot-terrain interaction from RGB(D)

## Trivial self-supervision





# Learning to predict robot-terrain interaction from RGB(D)






# Learning to predict terrain properties


loss = (10.7341478)

lidar camera

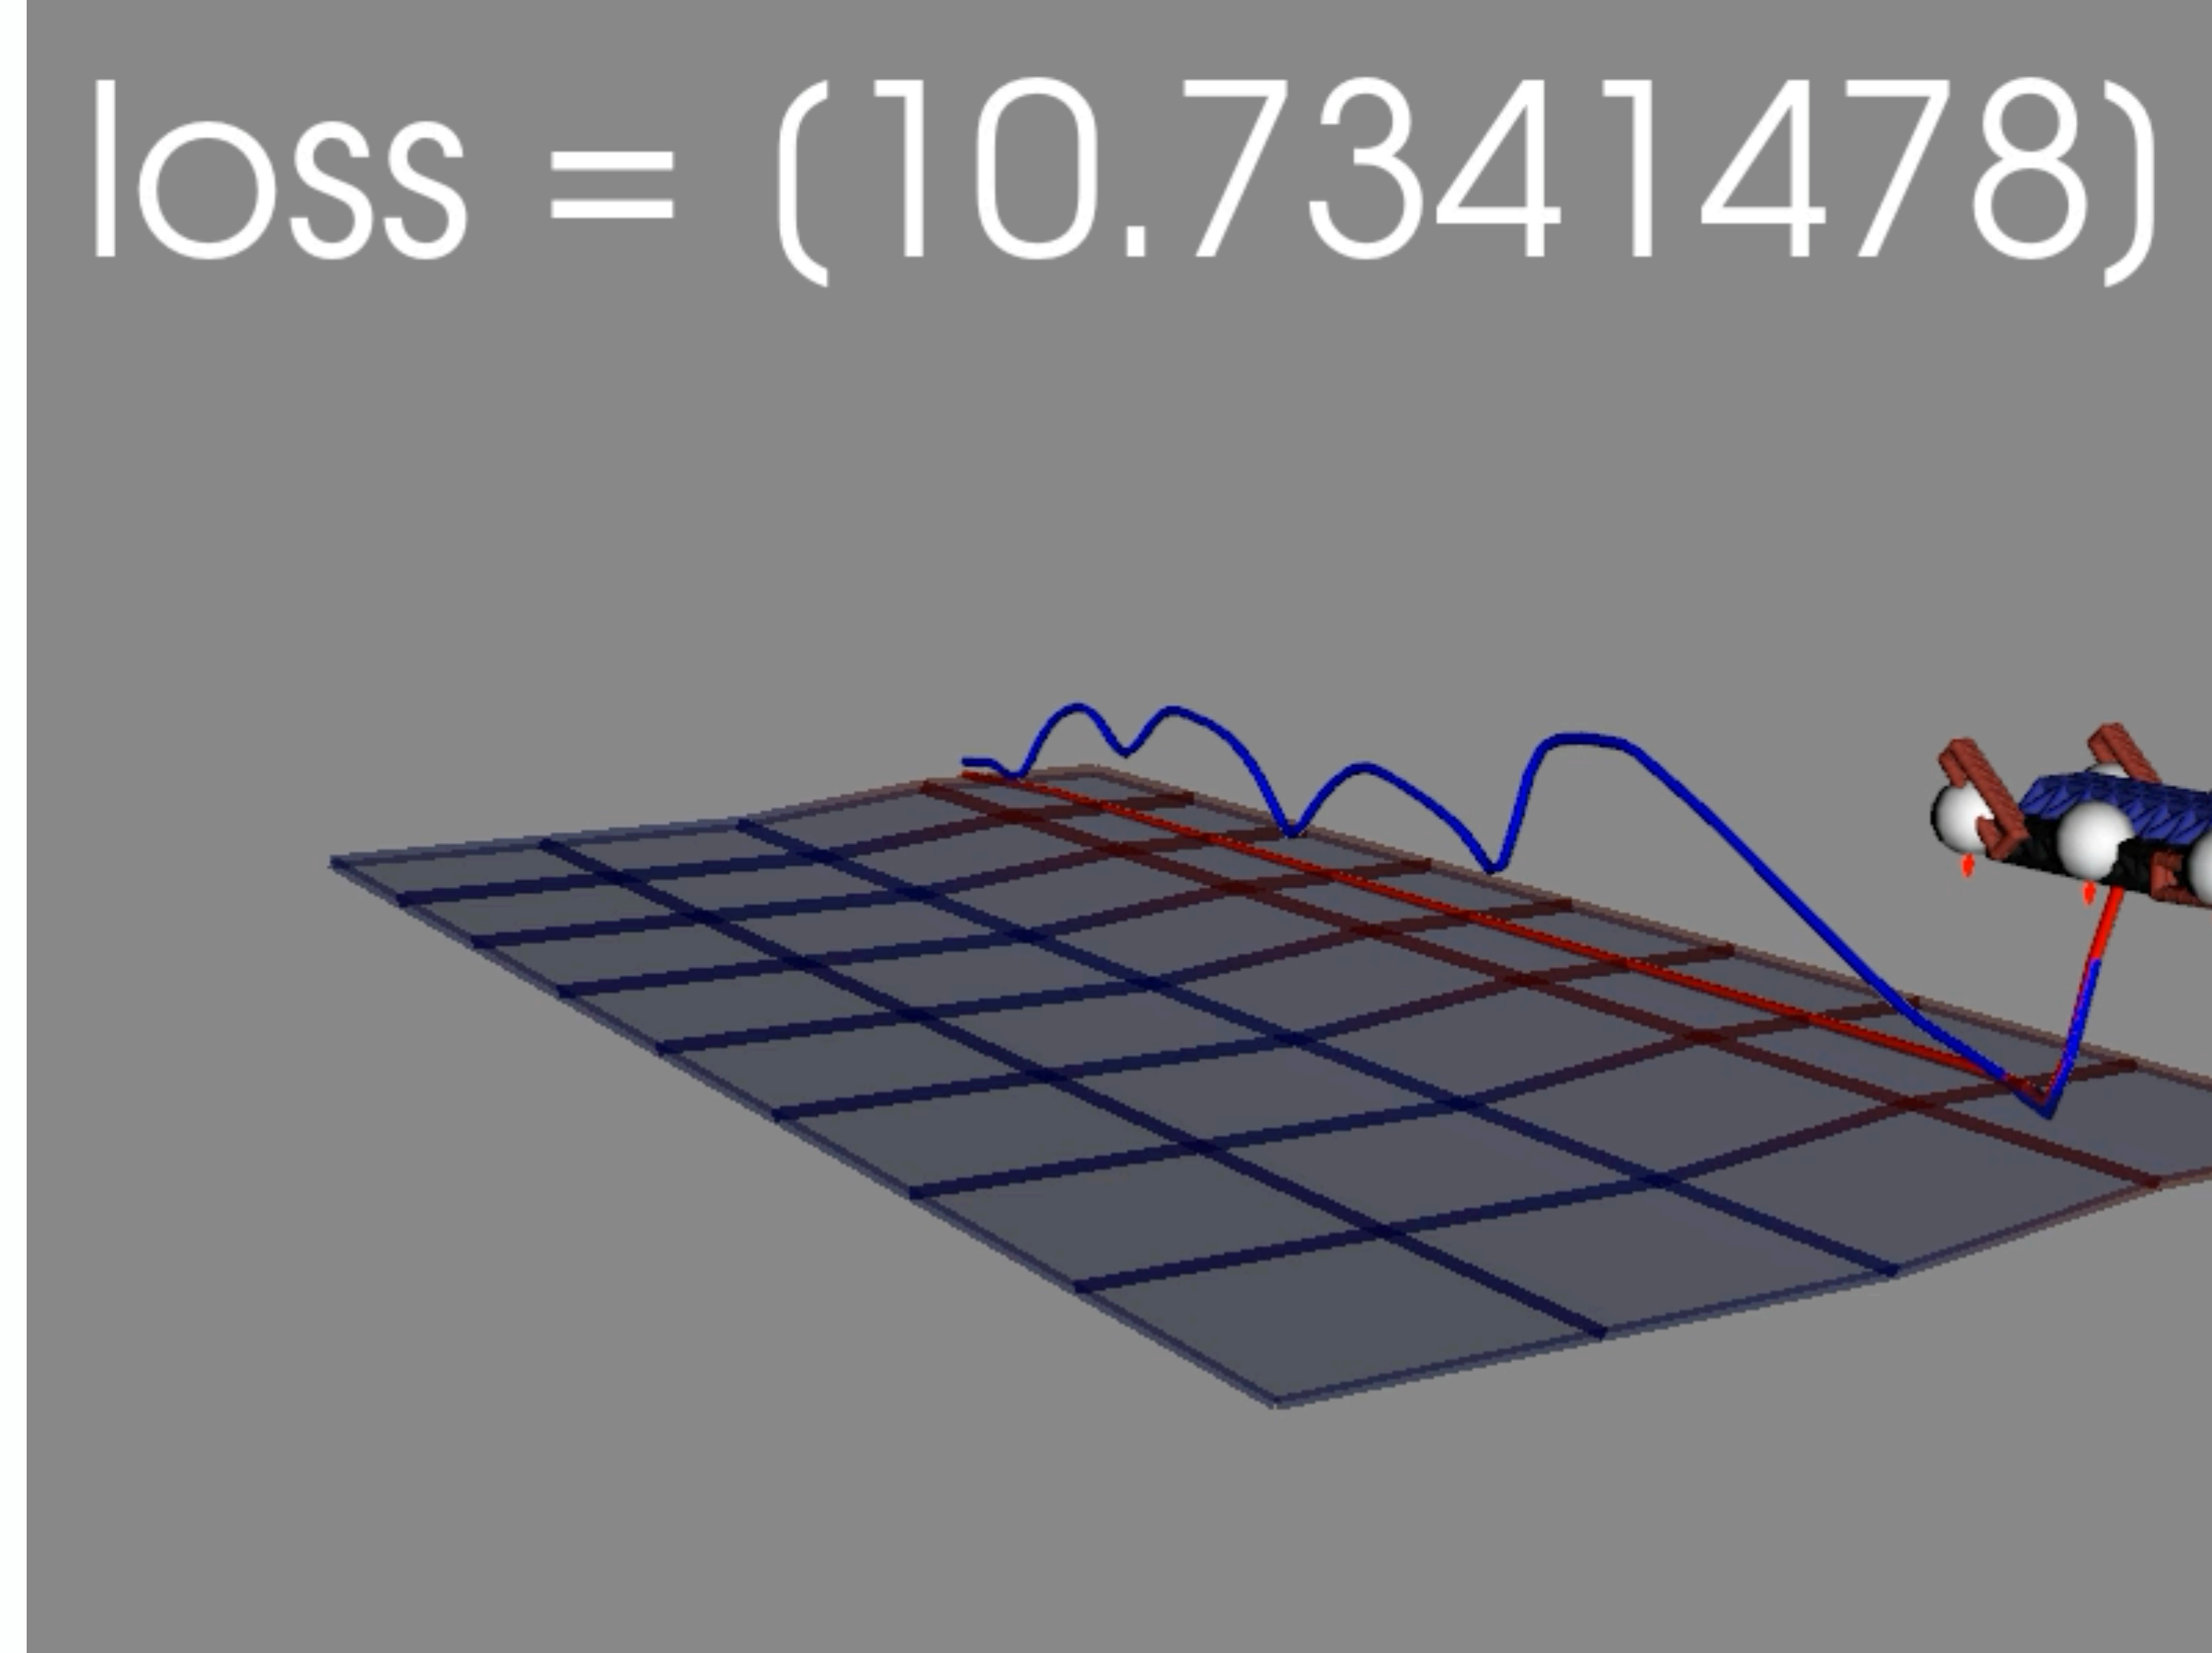

$\mathbf{x}$



estimated terrain  $\hat{\mathbf{h}}$   
heightmap, elasticity, damping,  
friction, interaction forces



estim. trajectory



real robot trajectory

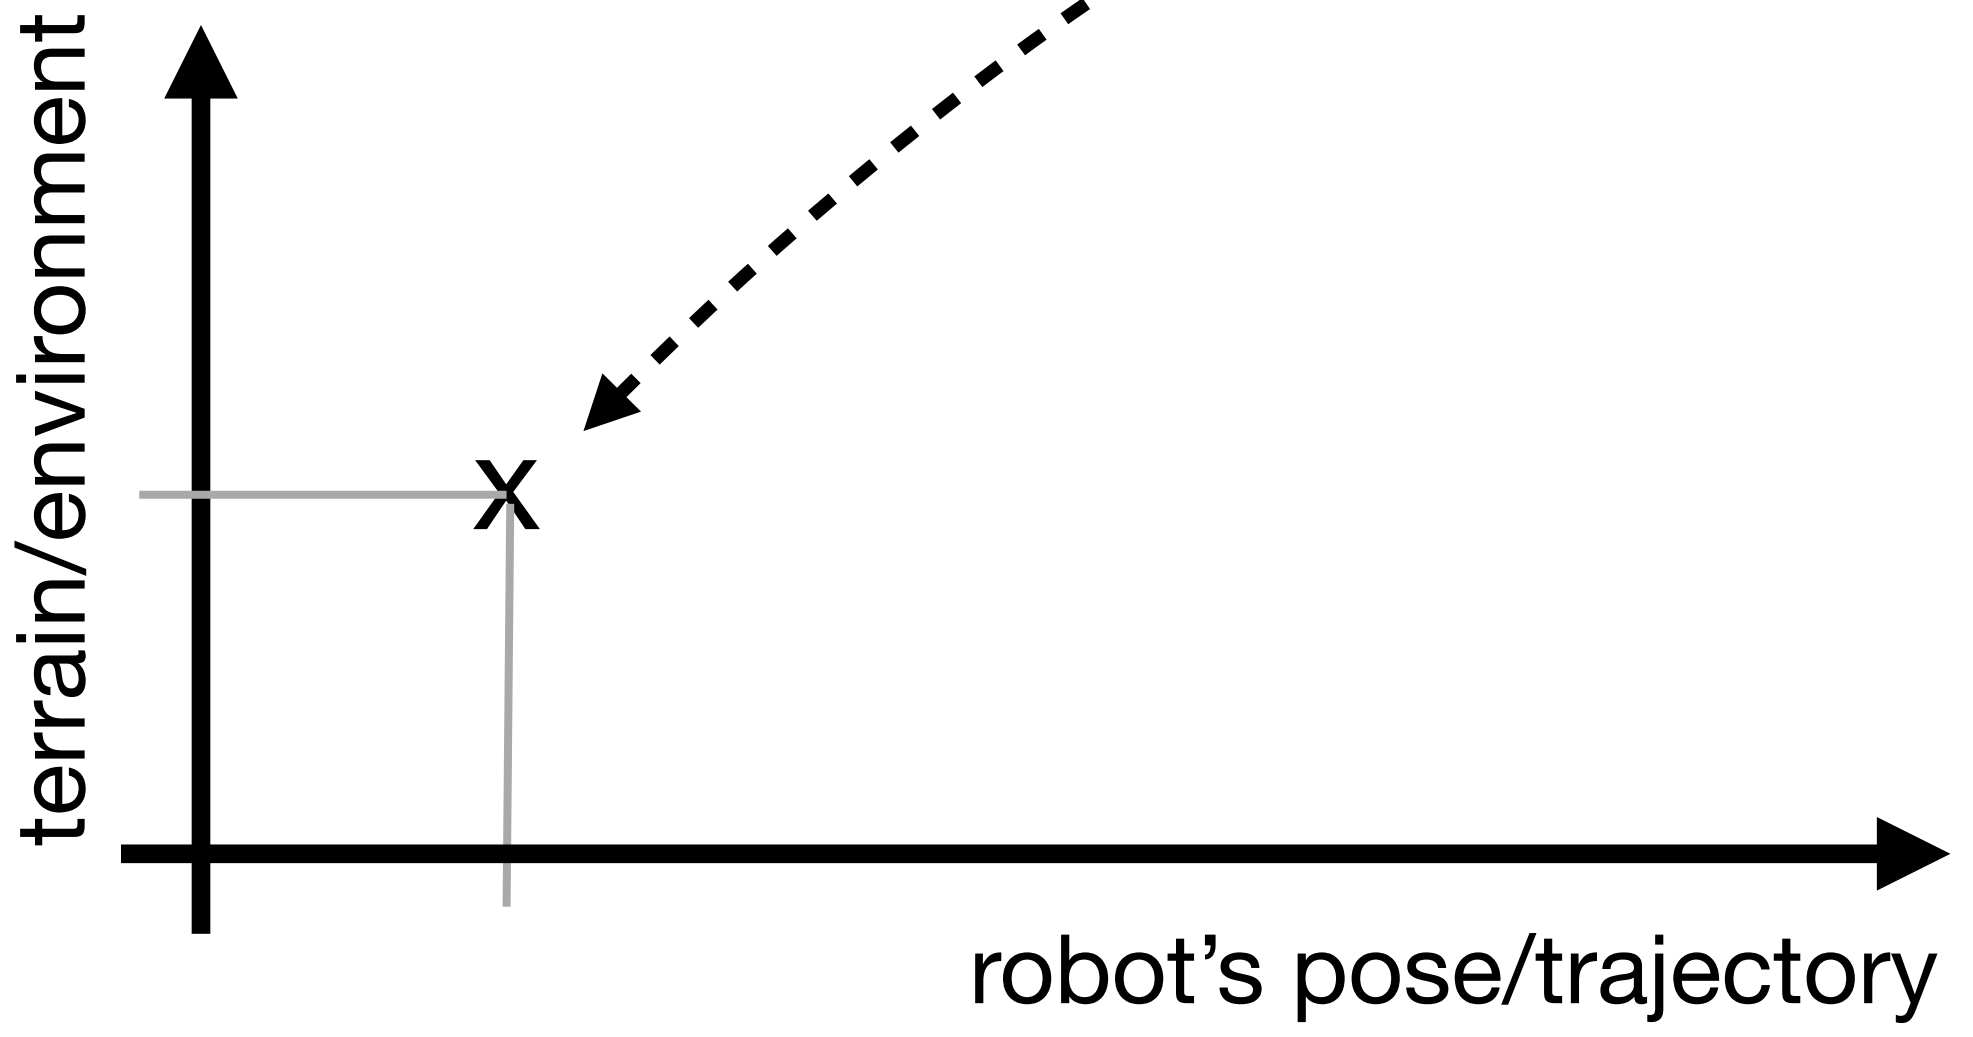

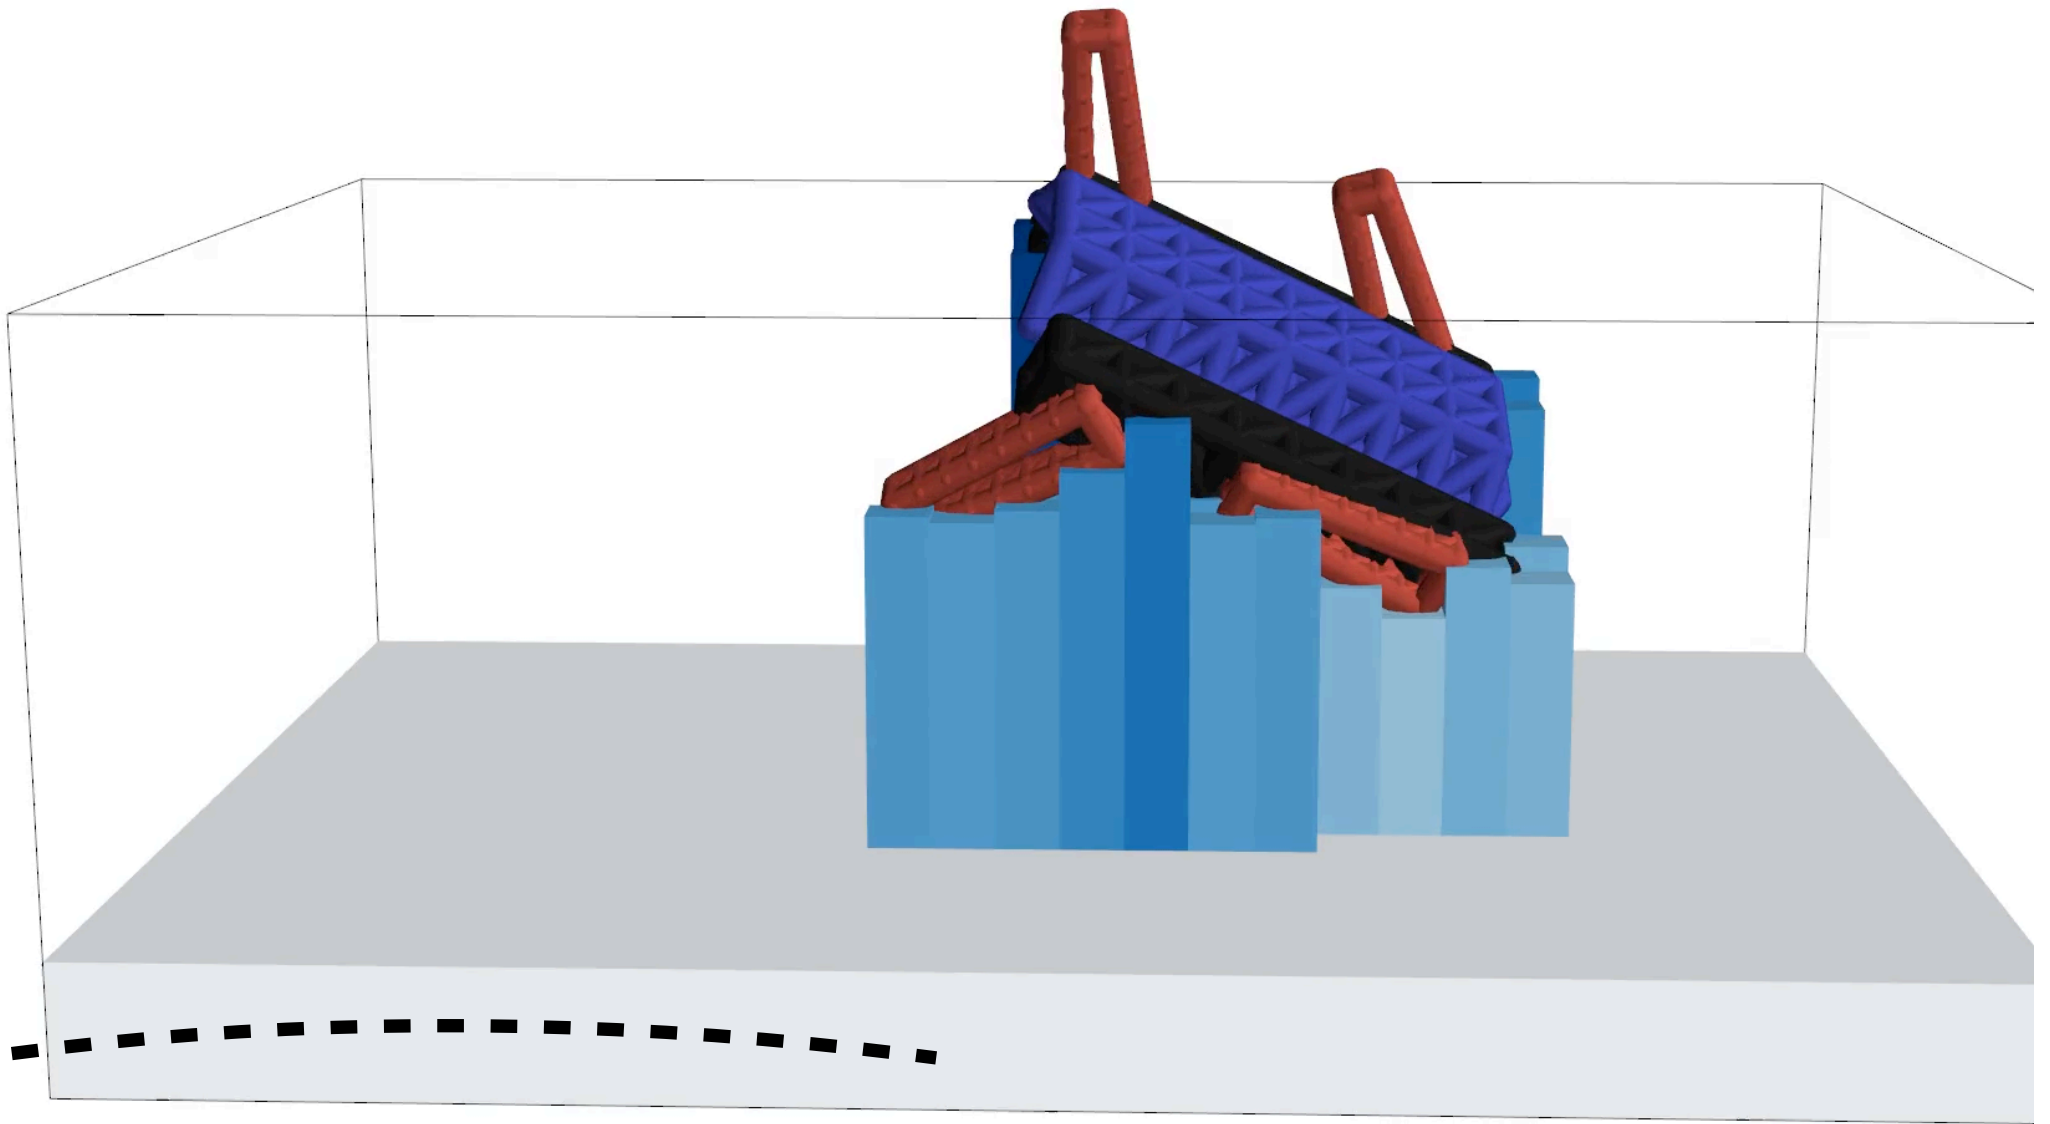
$\mathbf{x}^*$

$$\frac{\partial h_{\theta}(\mathbf{x})}{\partial \theta}$$




$$\frac{\partial \nabla \text{physics}(\hat{\mathbf{h}})}{\partial \hat{\mathbf{h}}}$$

$$\frac{\partial \|\mathbf{x} - \mathbf{x}^*\|}{\partial \mathbf{x}}$$


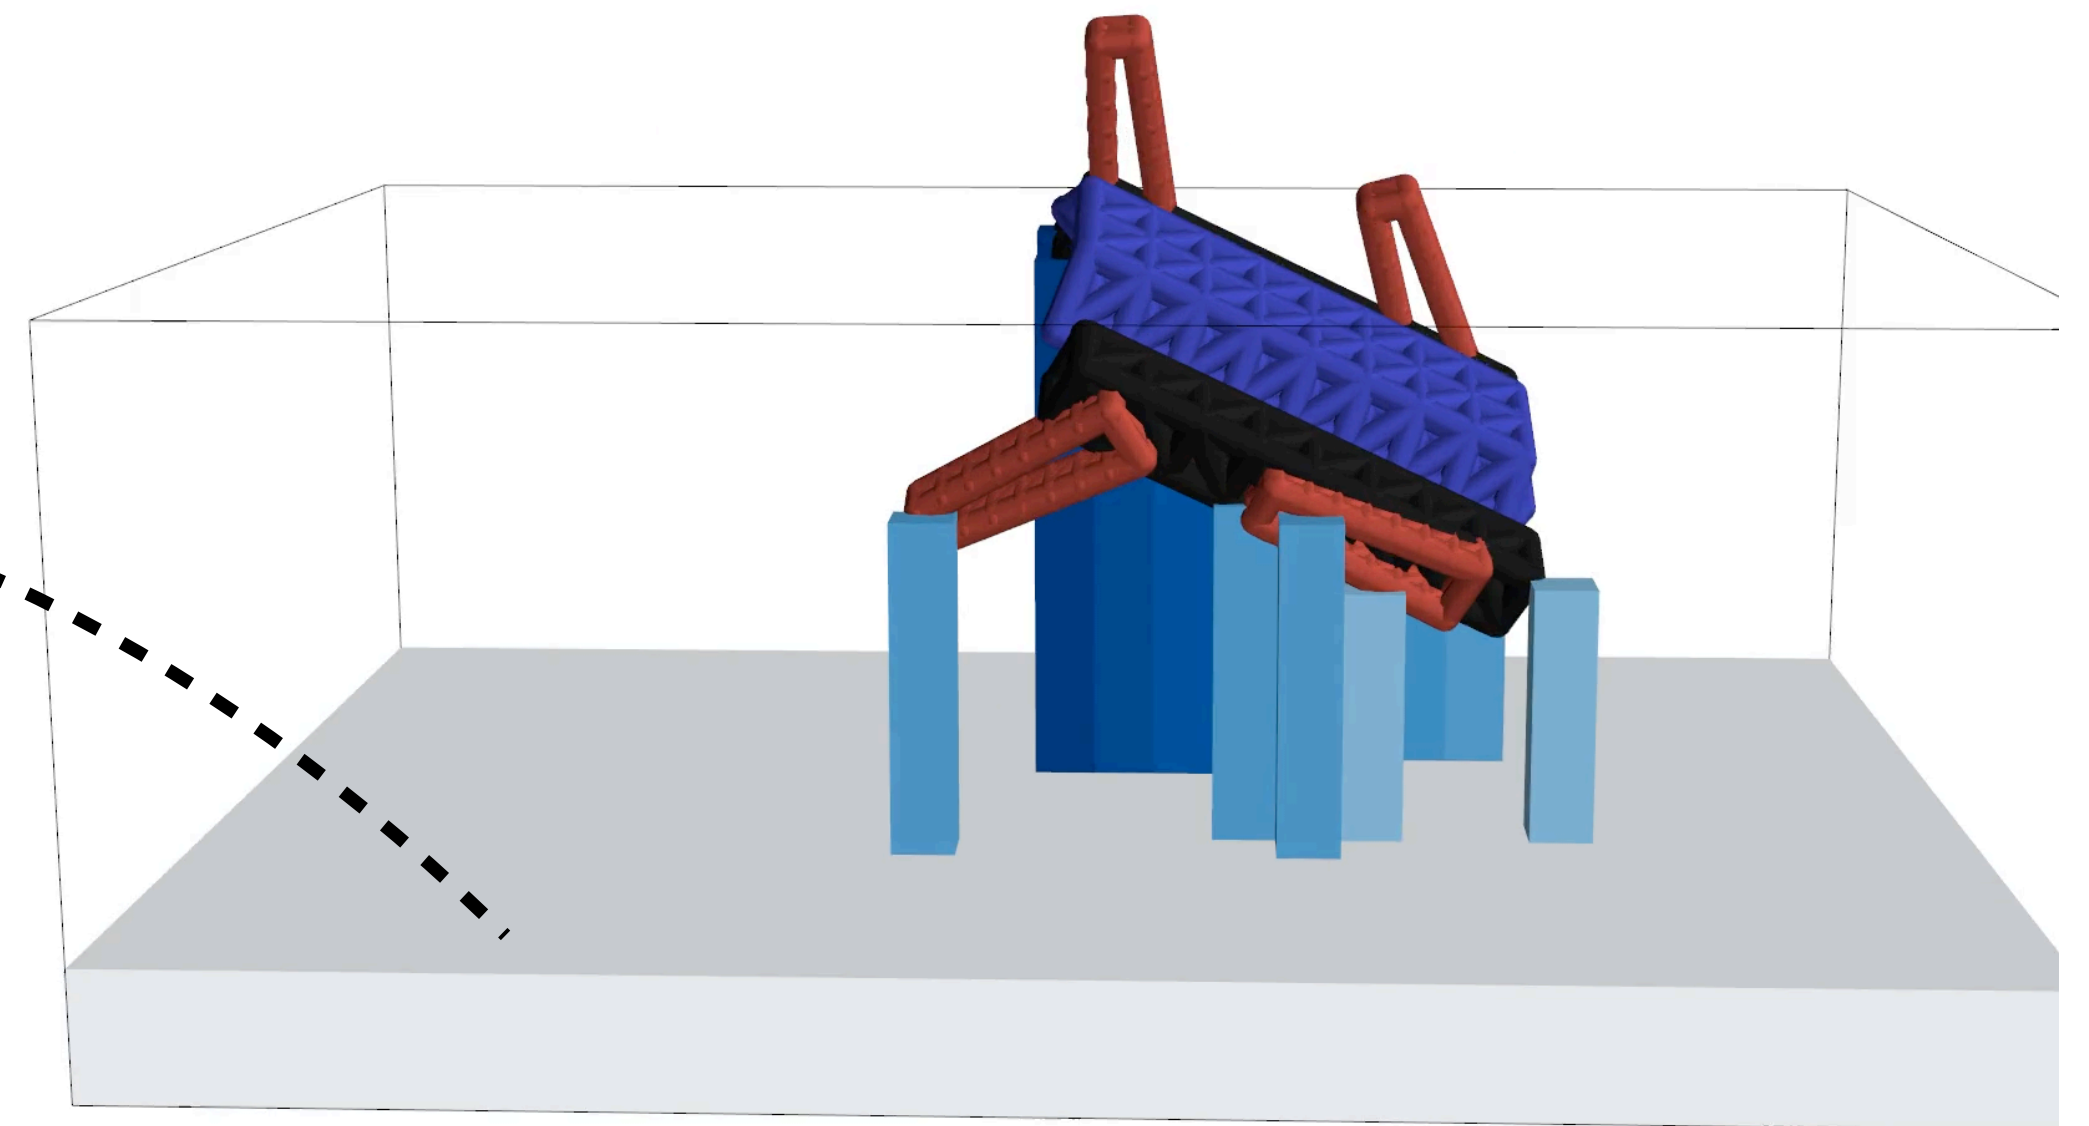









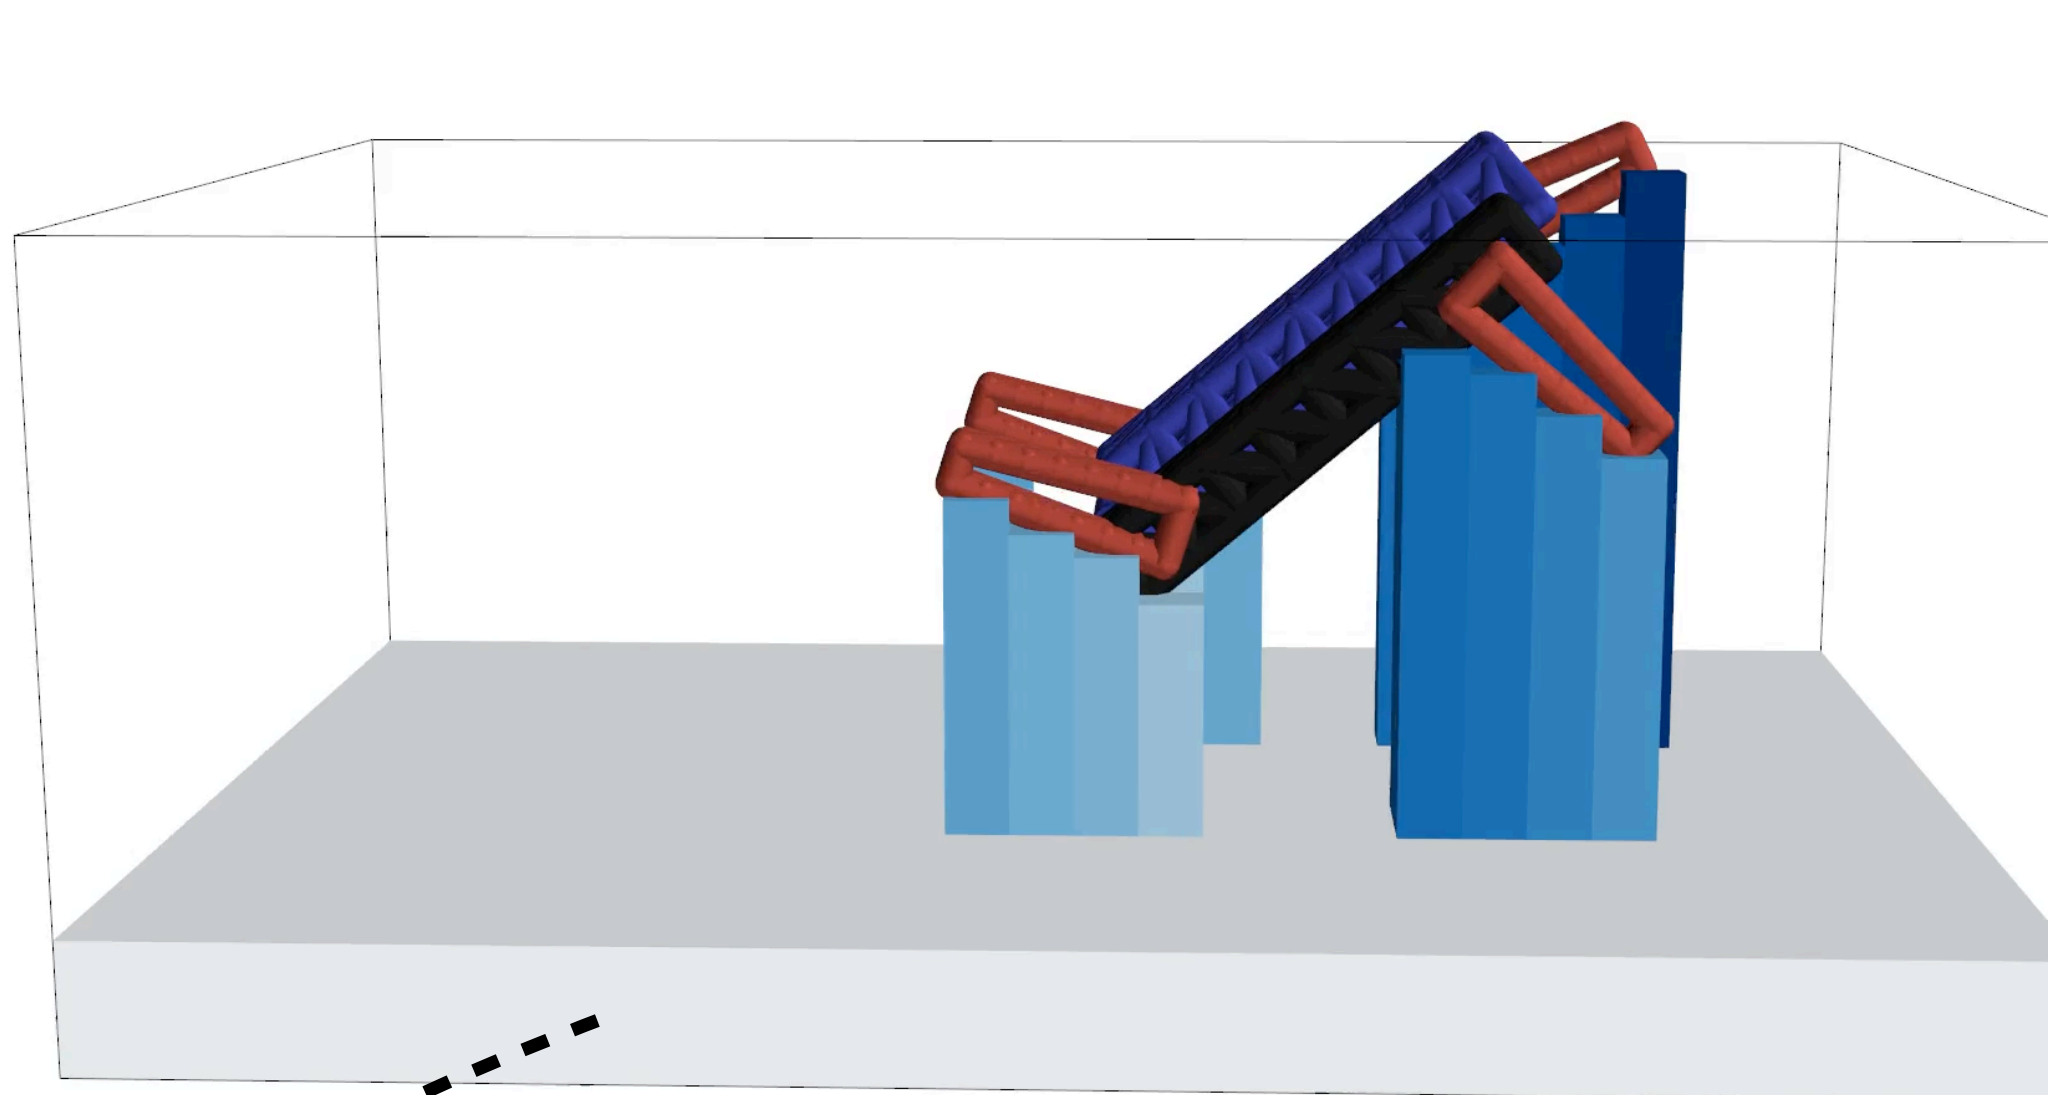
Many heightmaps yields the same pose




Many heightmaps yields the same pose



Many heightmaps yields the same pose

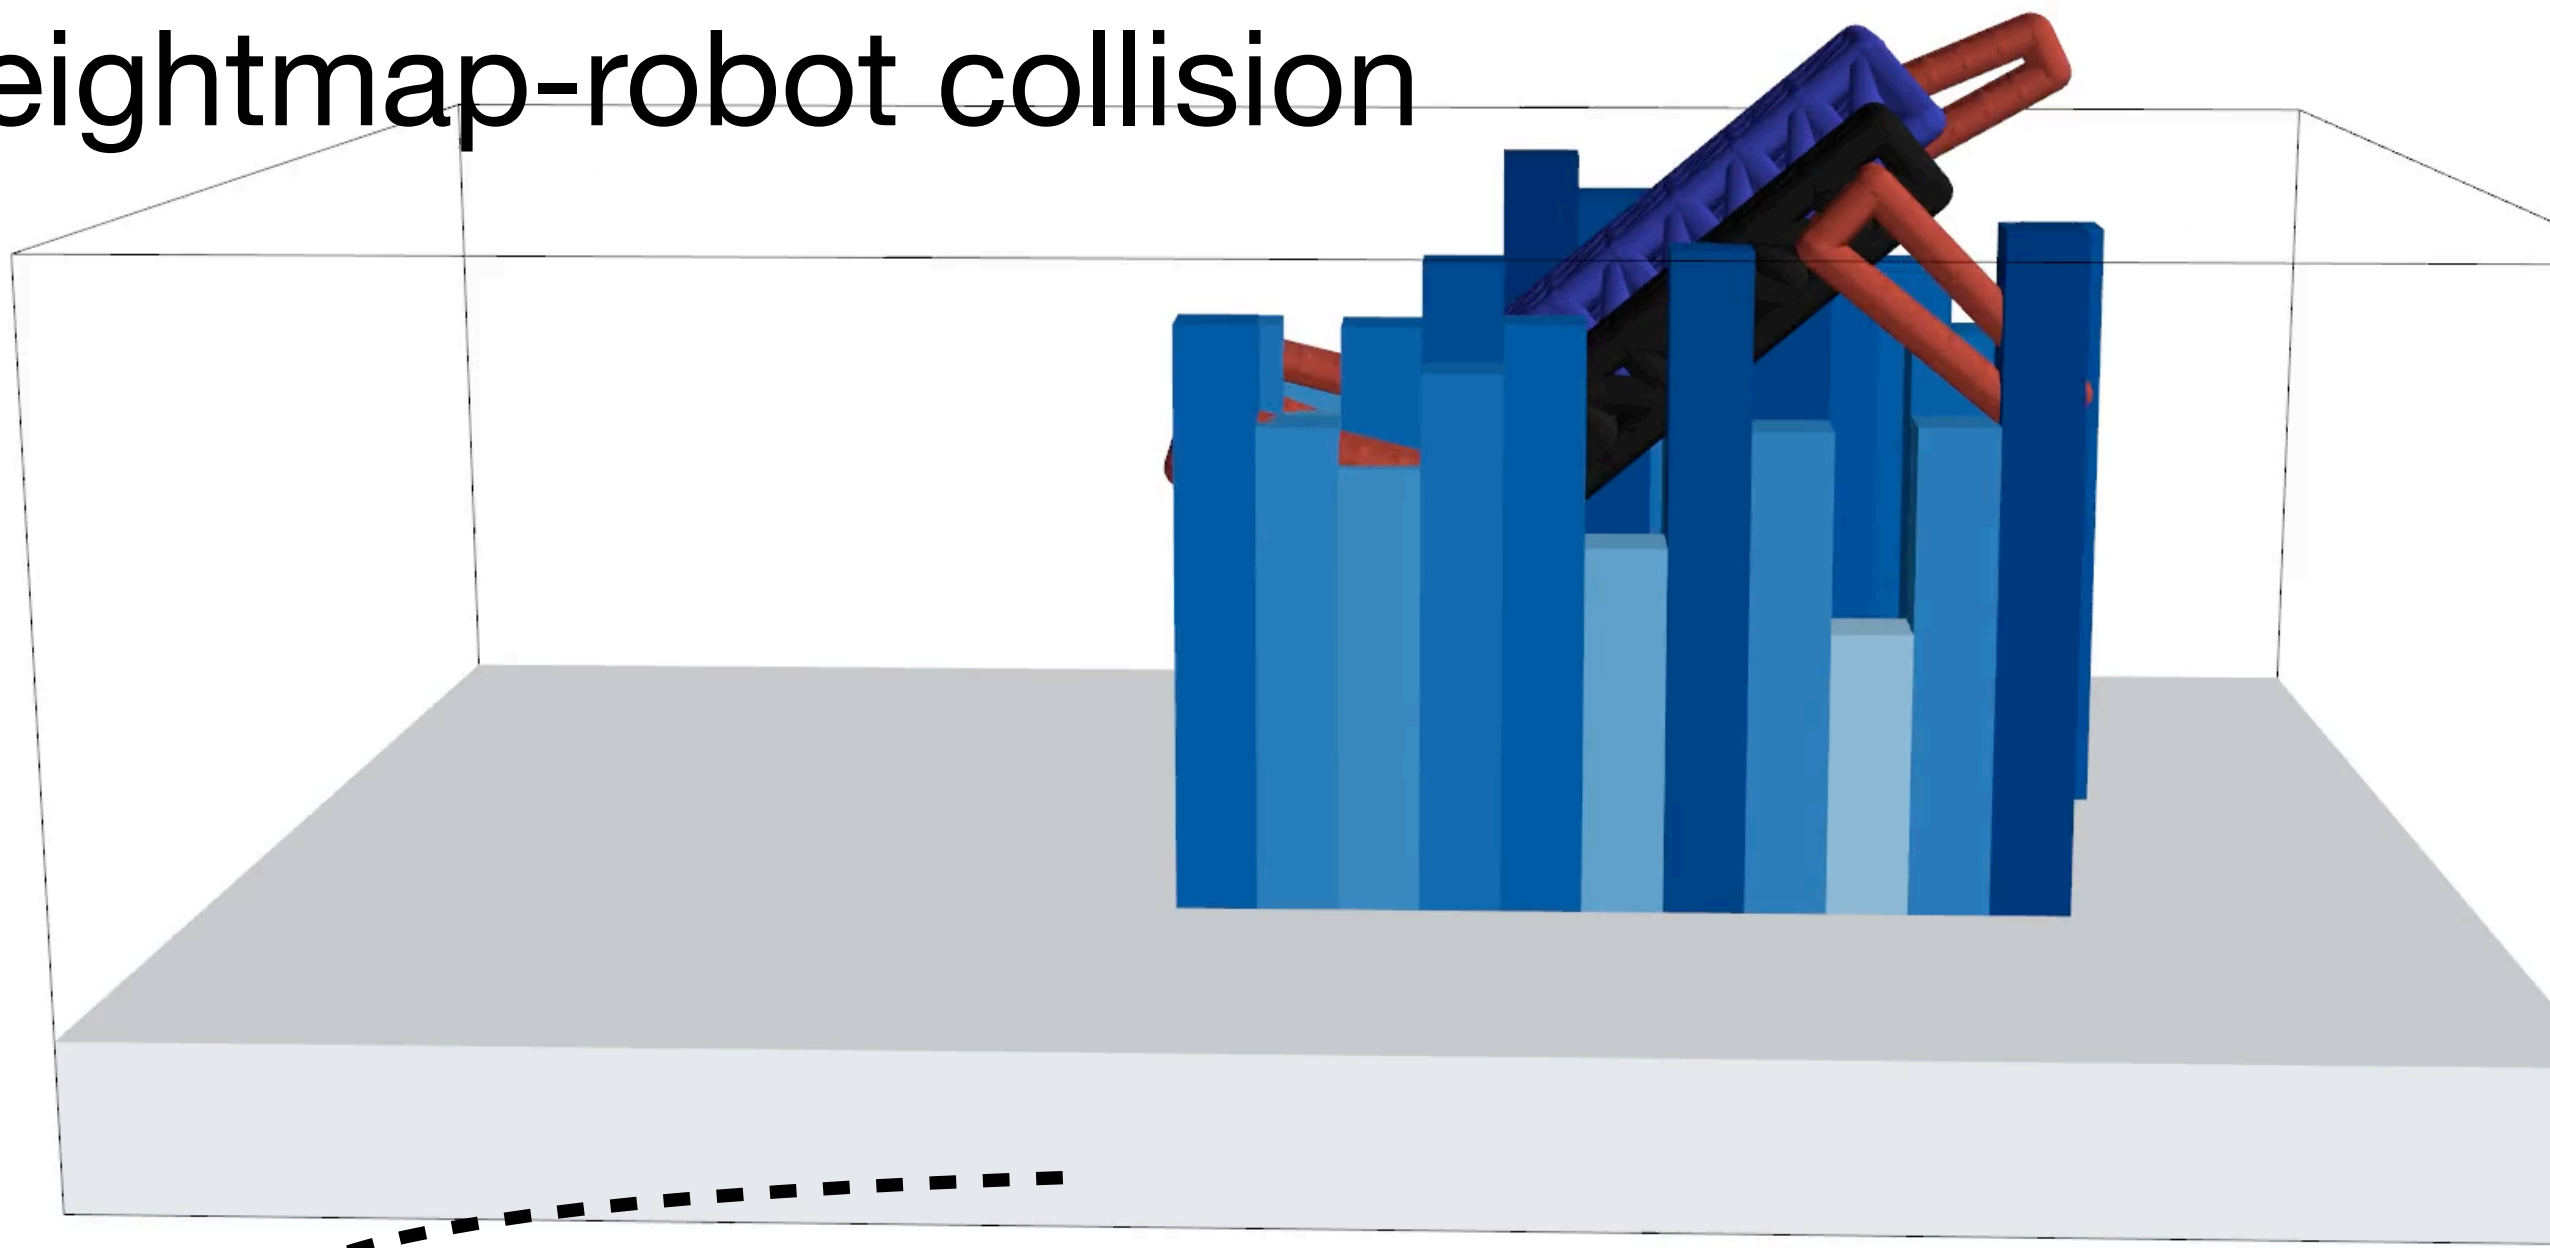









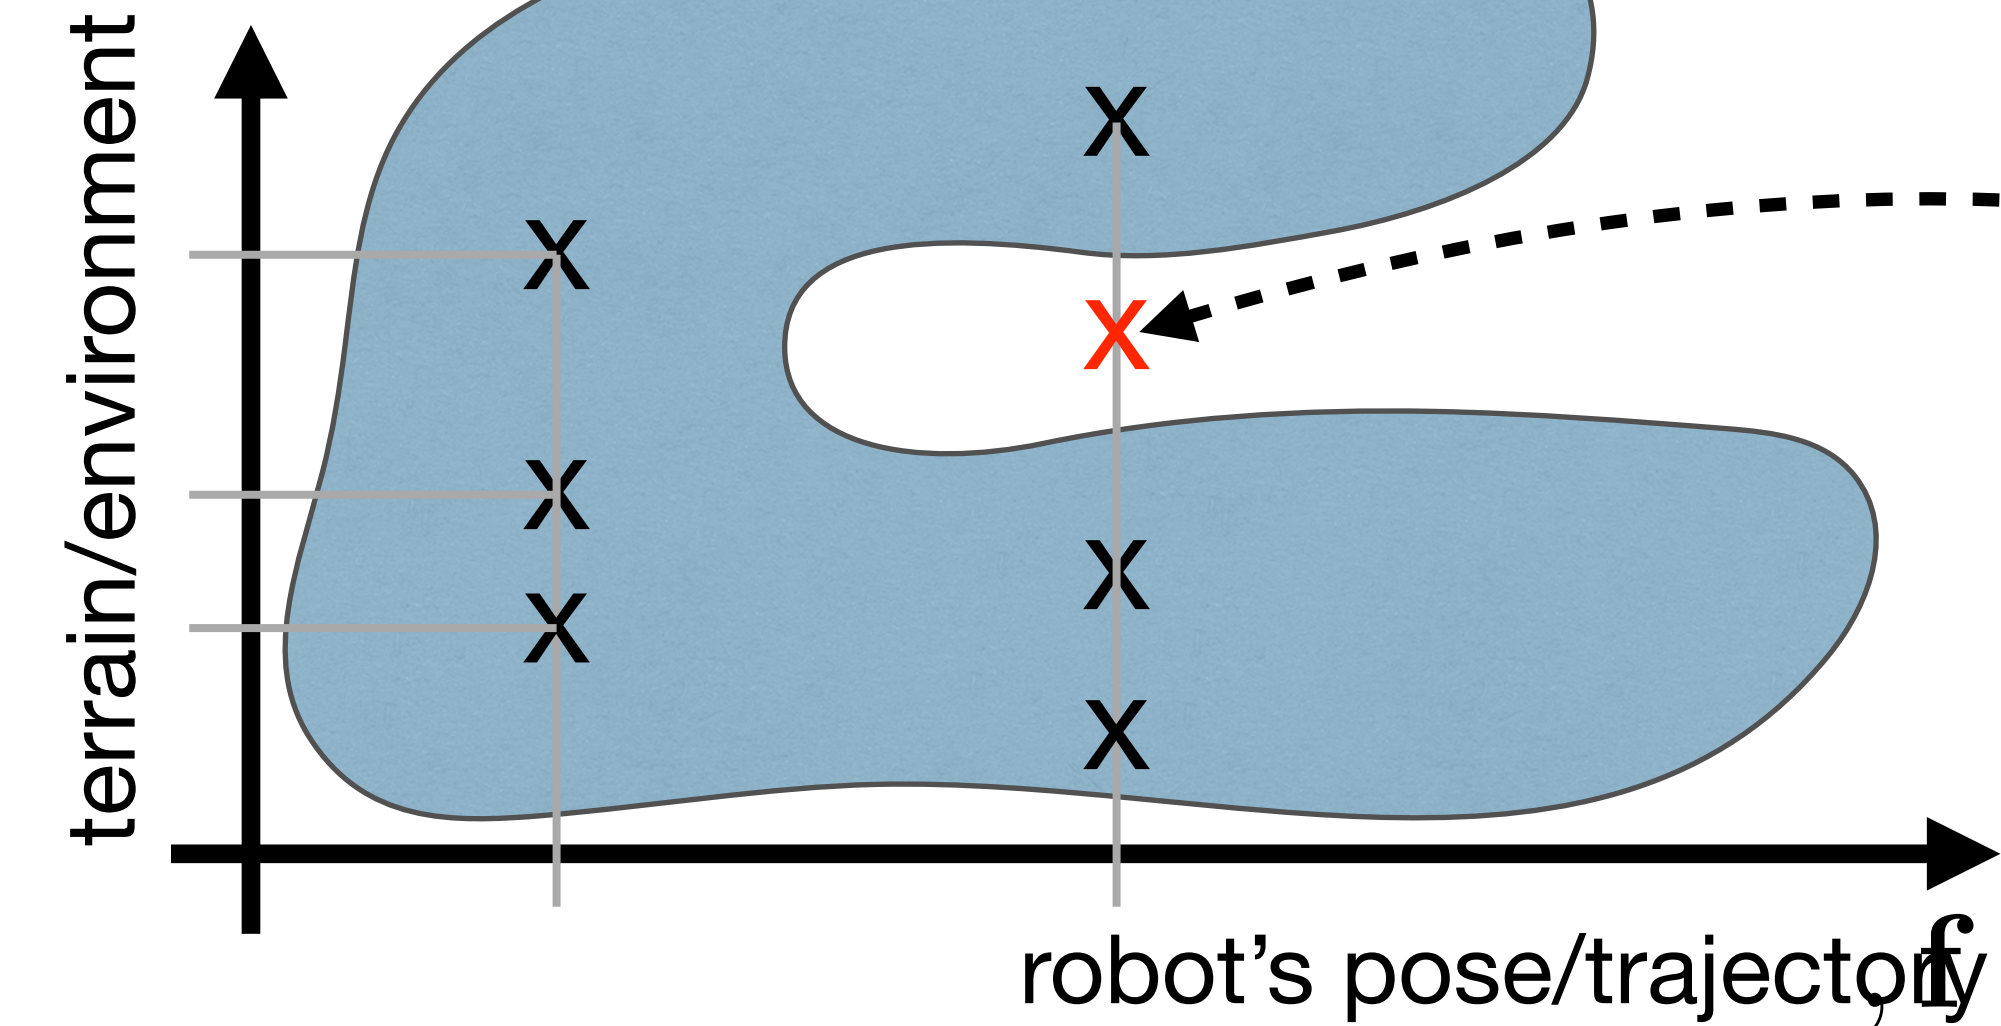
heightmap-robot collision



insufficiently supported pose




The remaining tuples are physically inconsistent





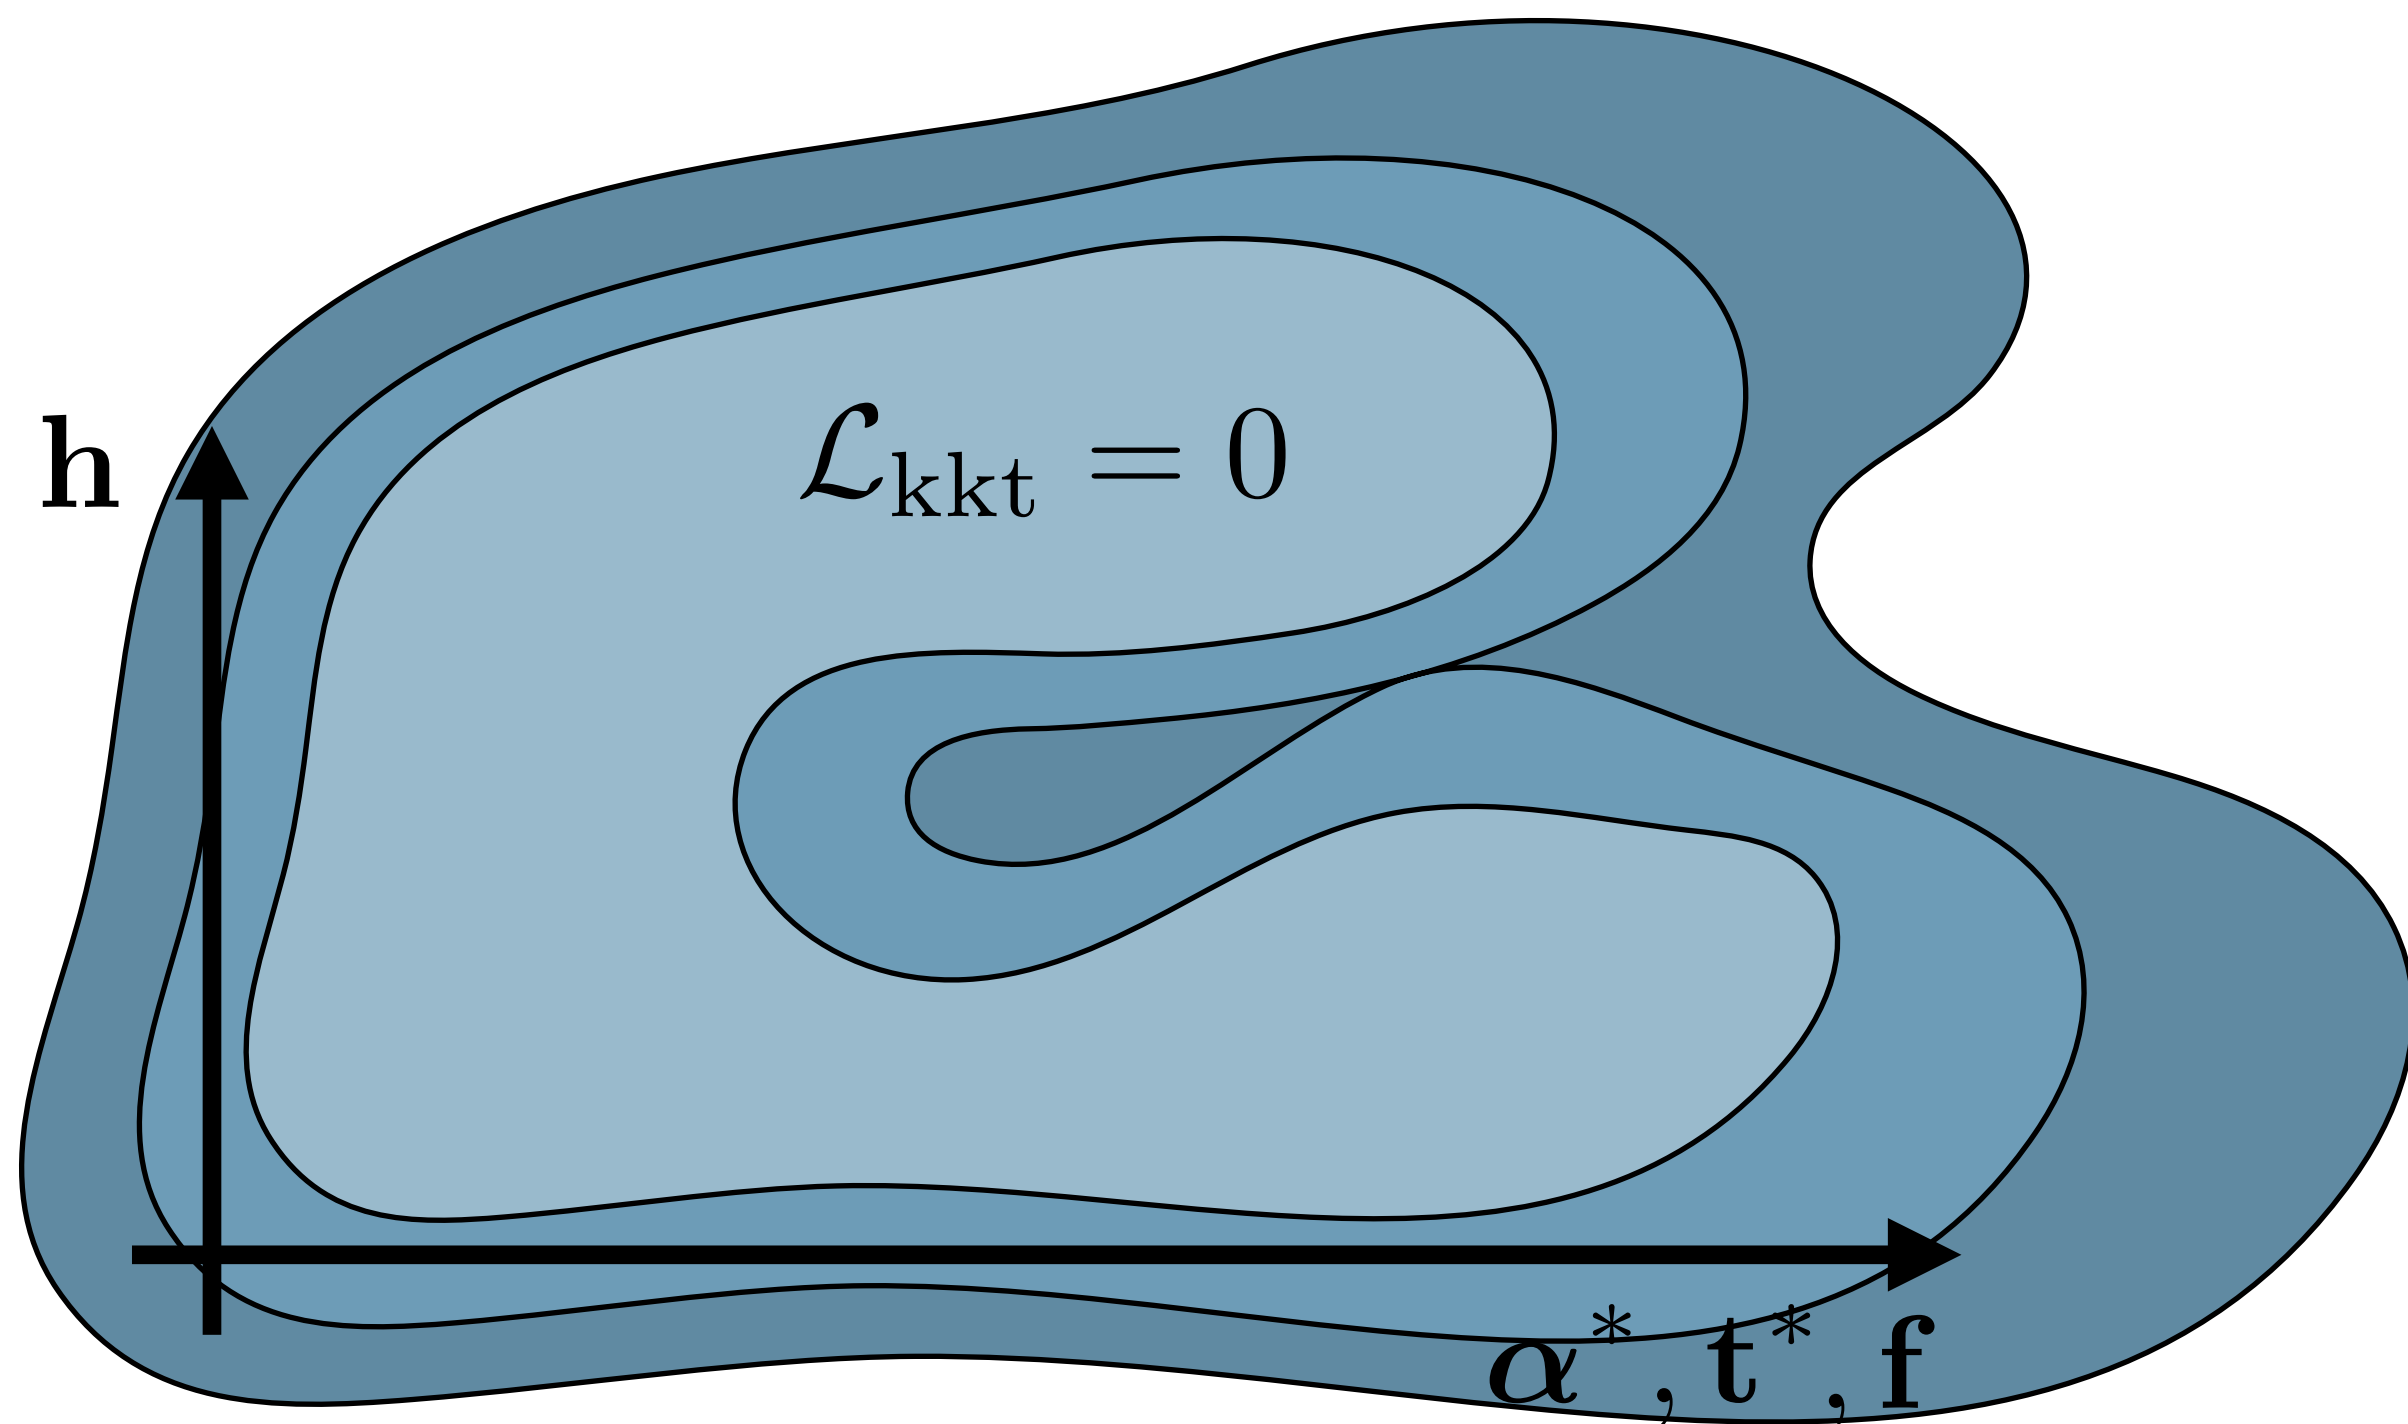


real world physical inconsistencies



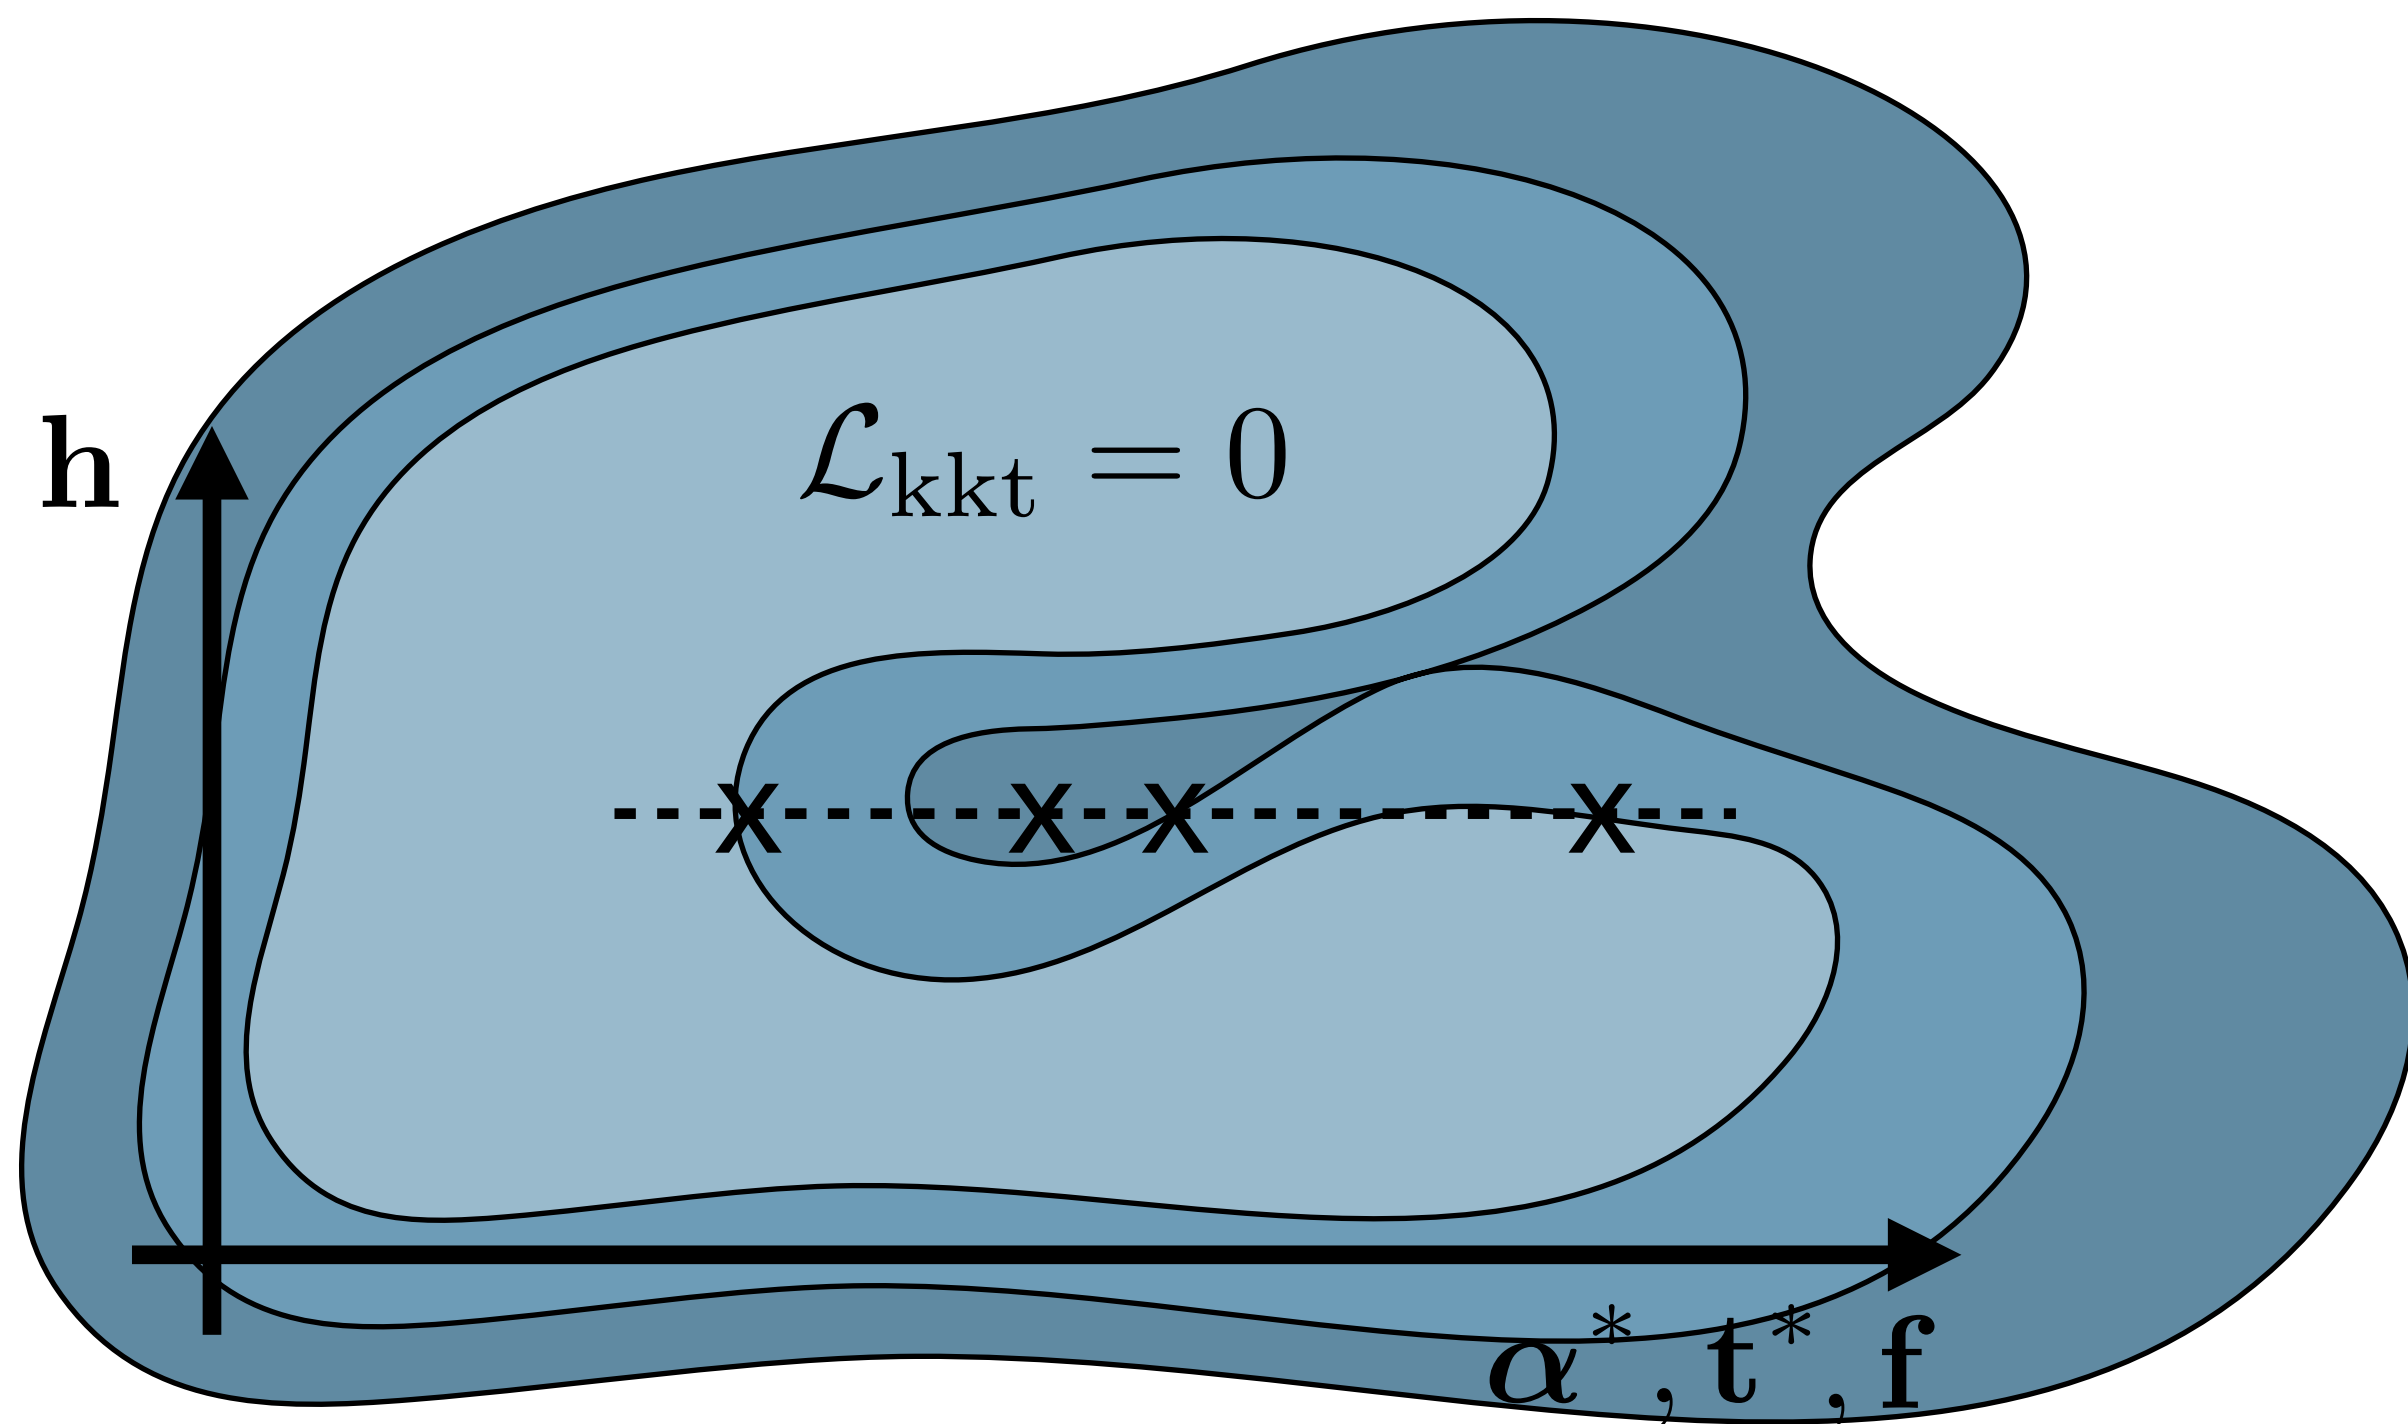


Construct differentiable loss that measures physical trajectory/terrain-consistency



Construct differentiable loss that measures physical trajectory/terrain-consistency

Given the ground truth trajectory one can optimize terrain



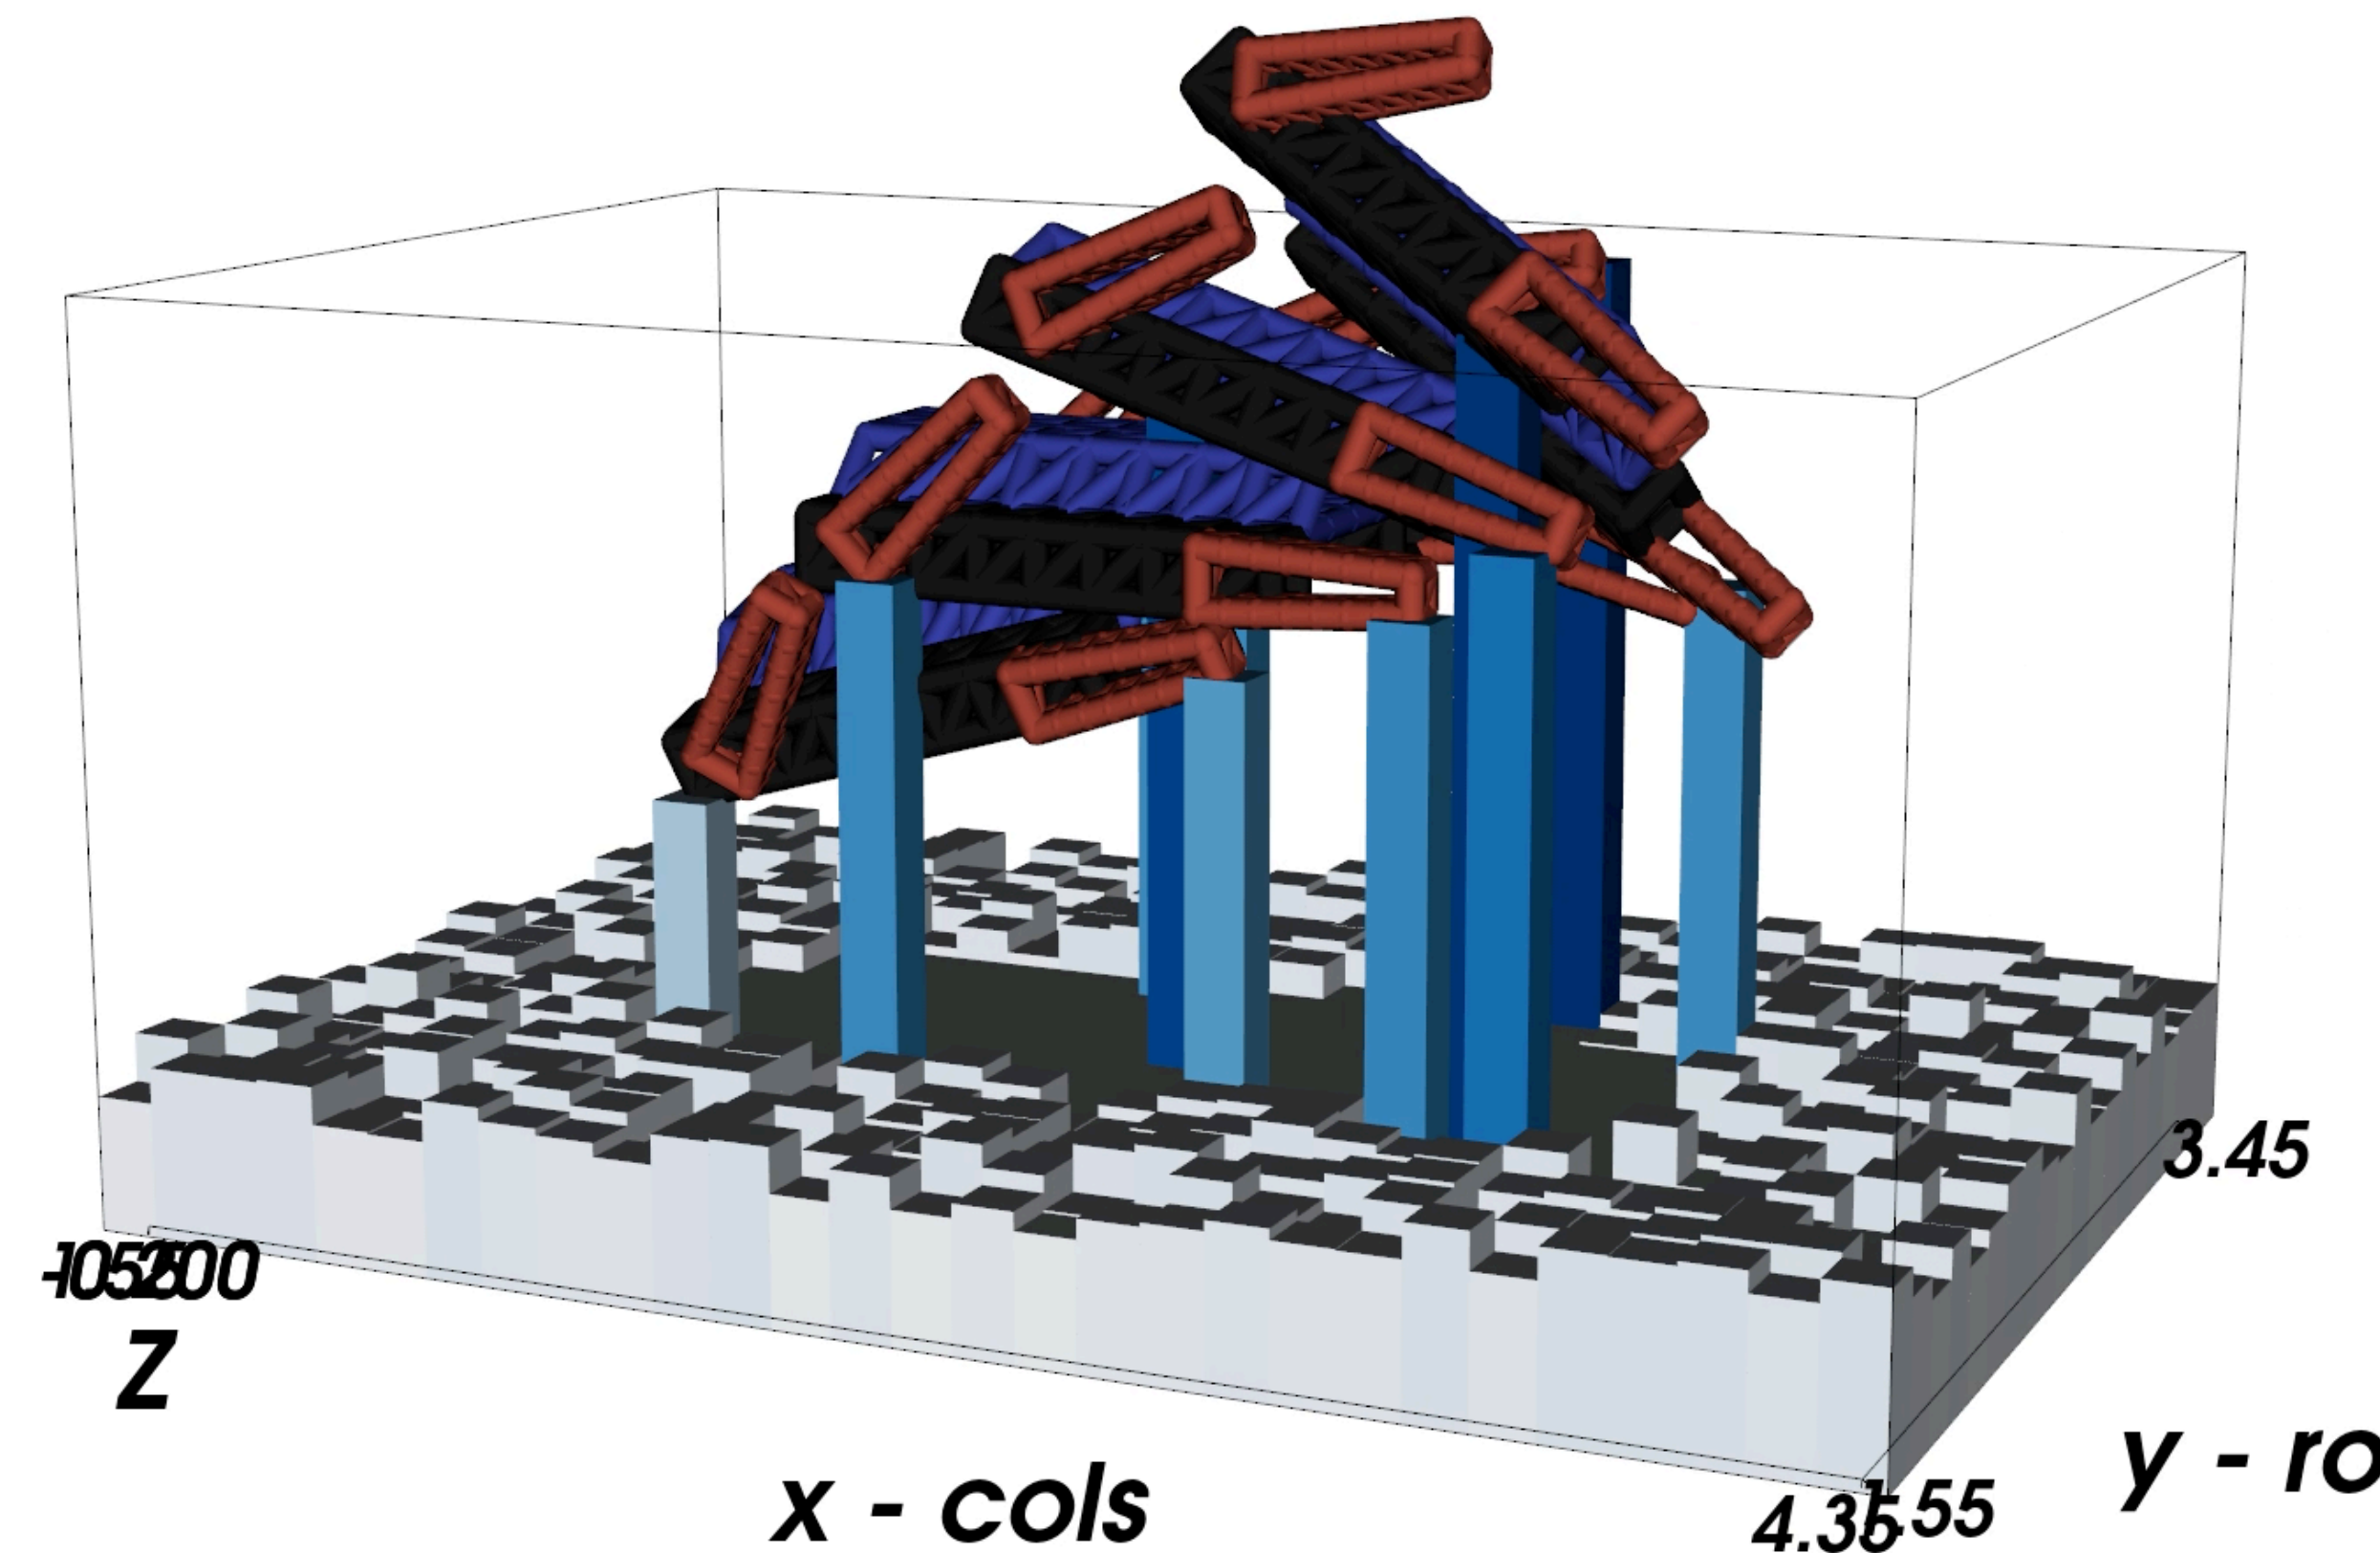



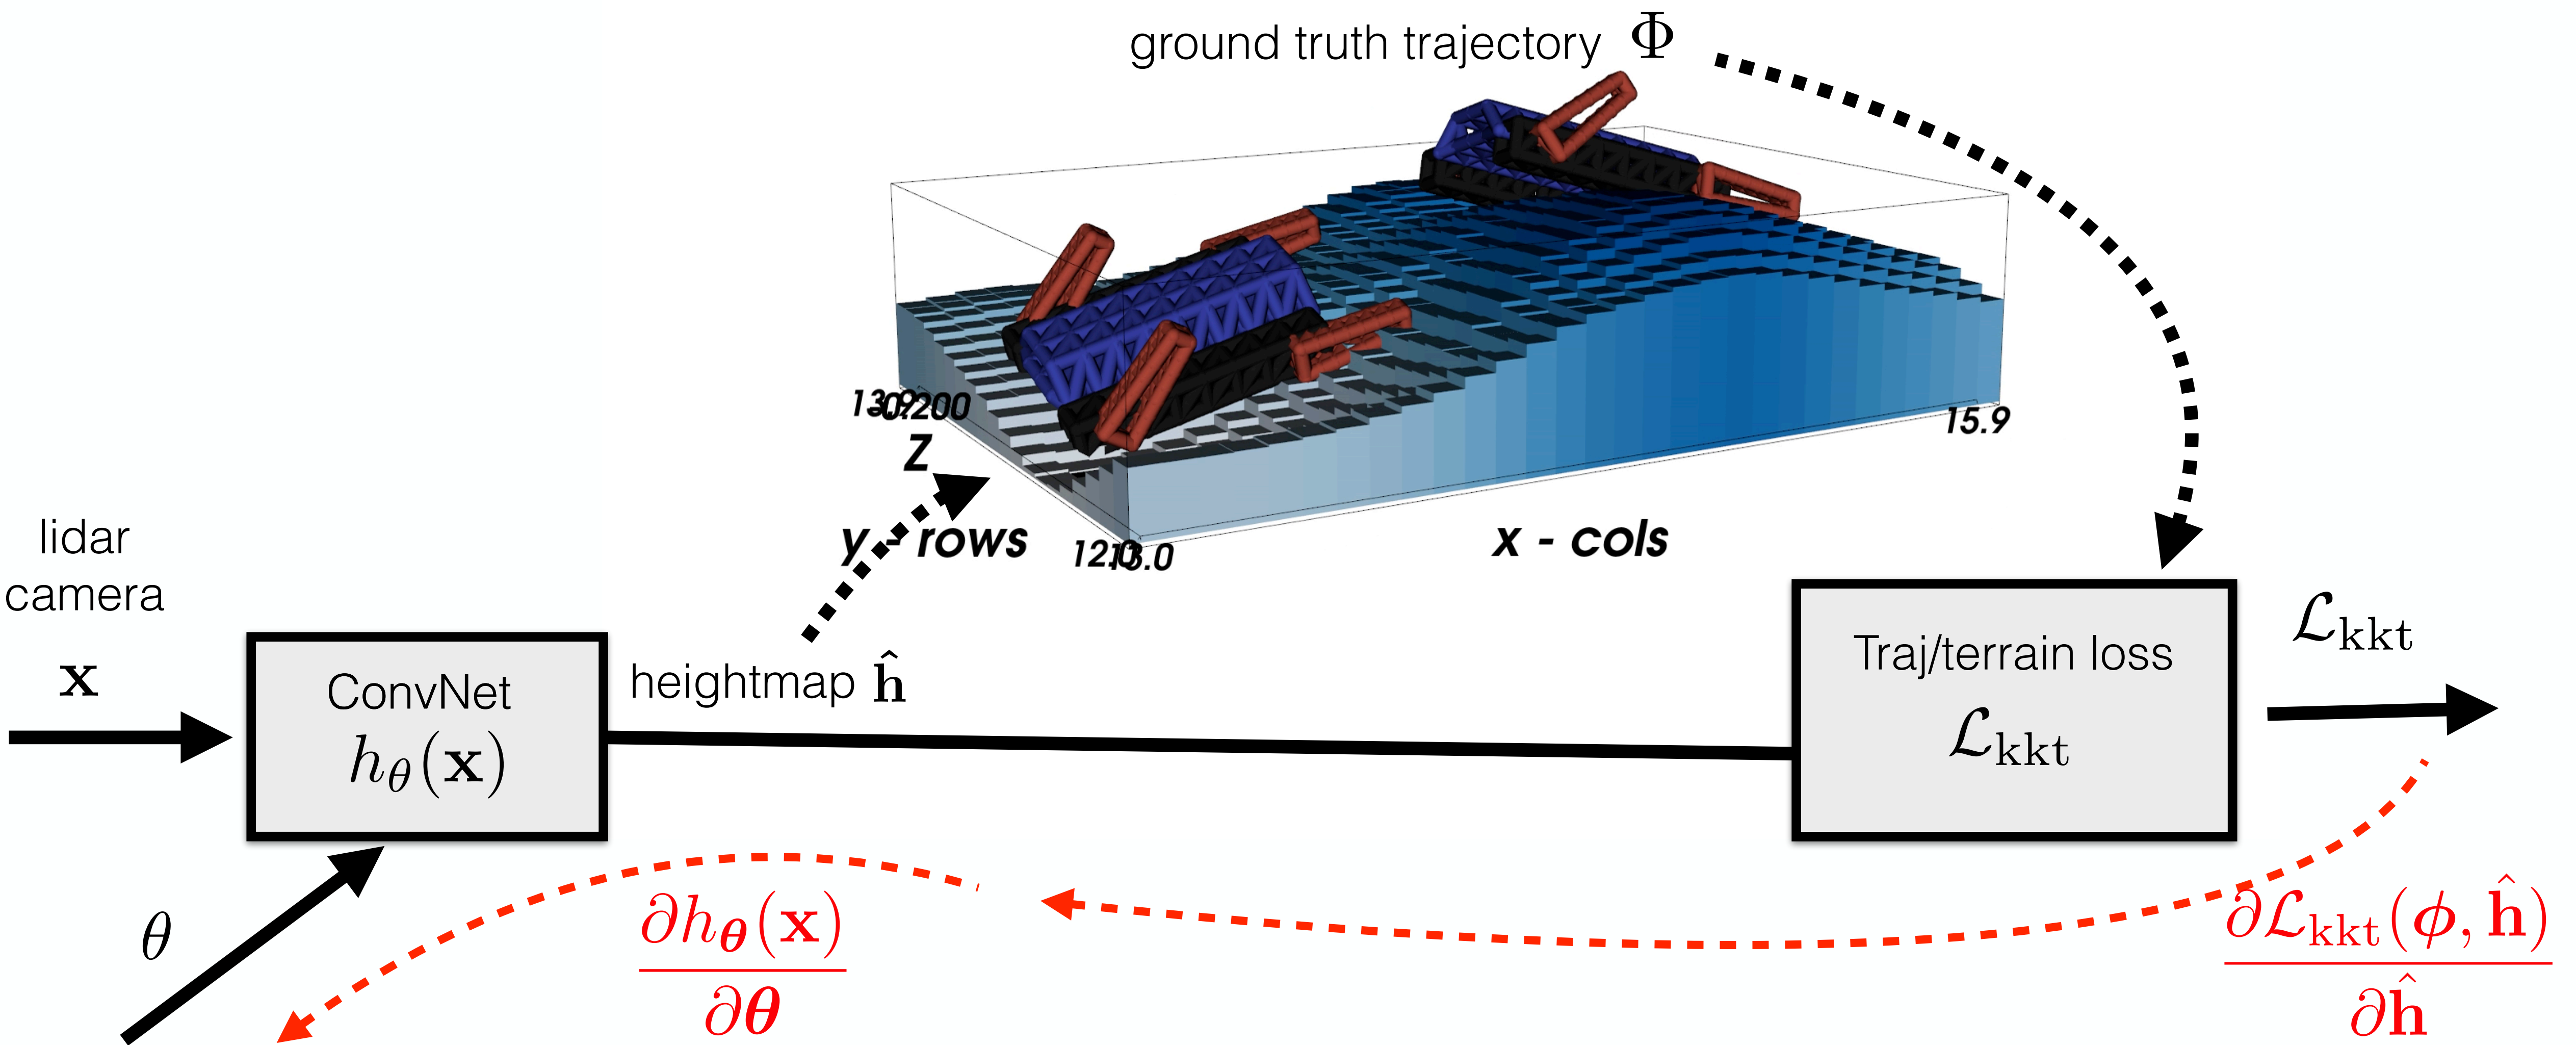
Construct differentiable loss that measures physical trajectory/terrain-consistency

Given the ground truth trajectory one can optimize terrain

Obviously huge sub-space of physically consistent terrains

Regularization needed

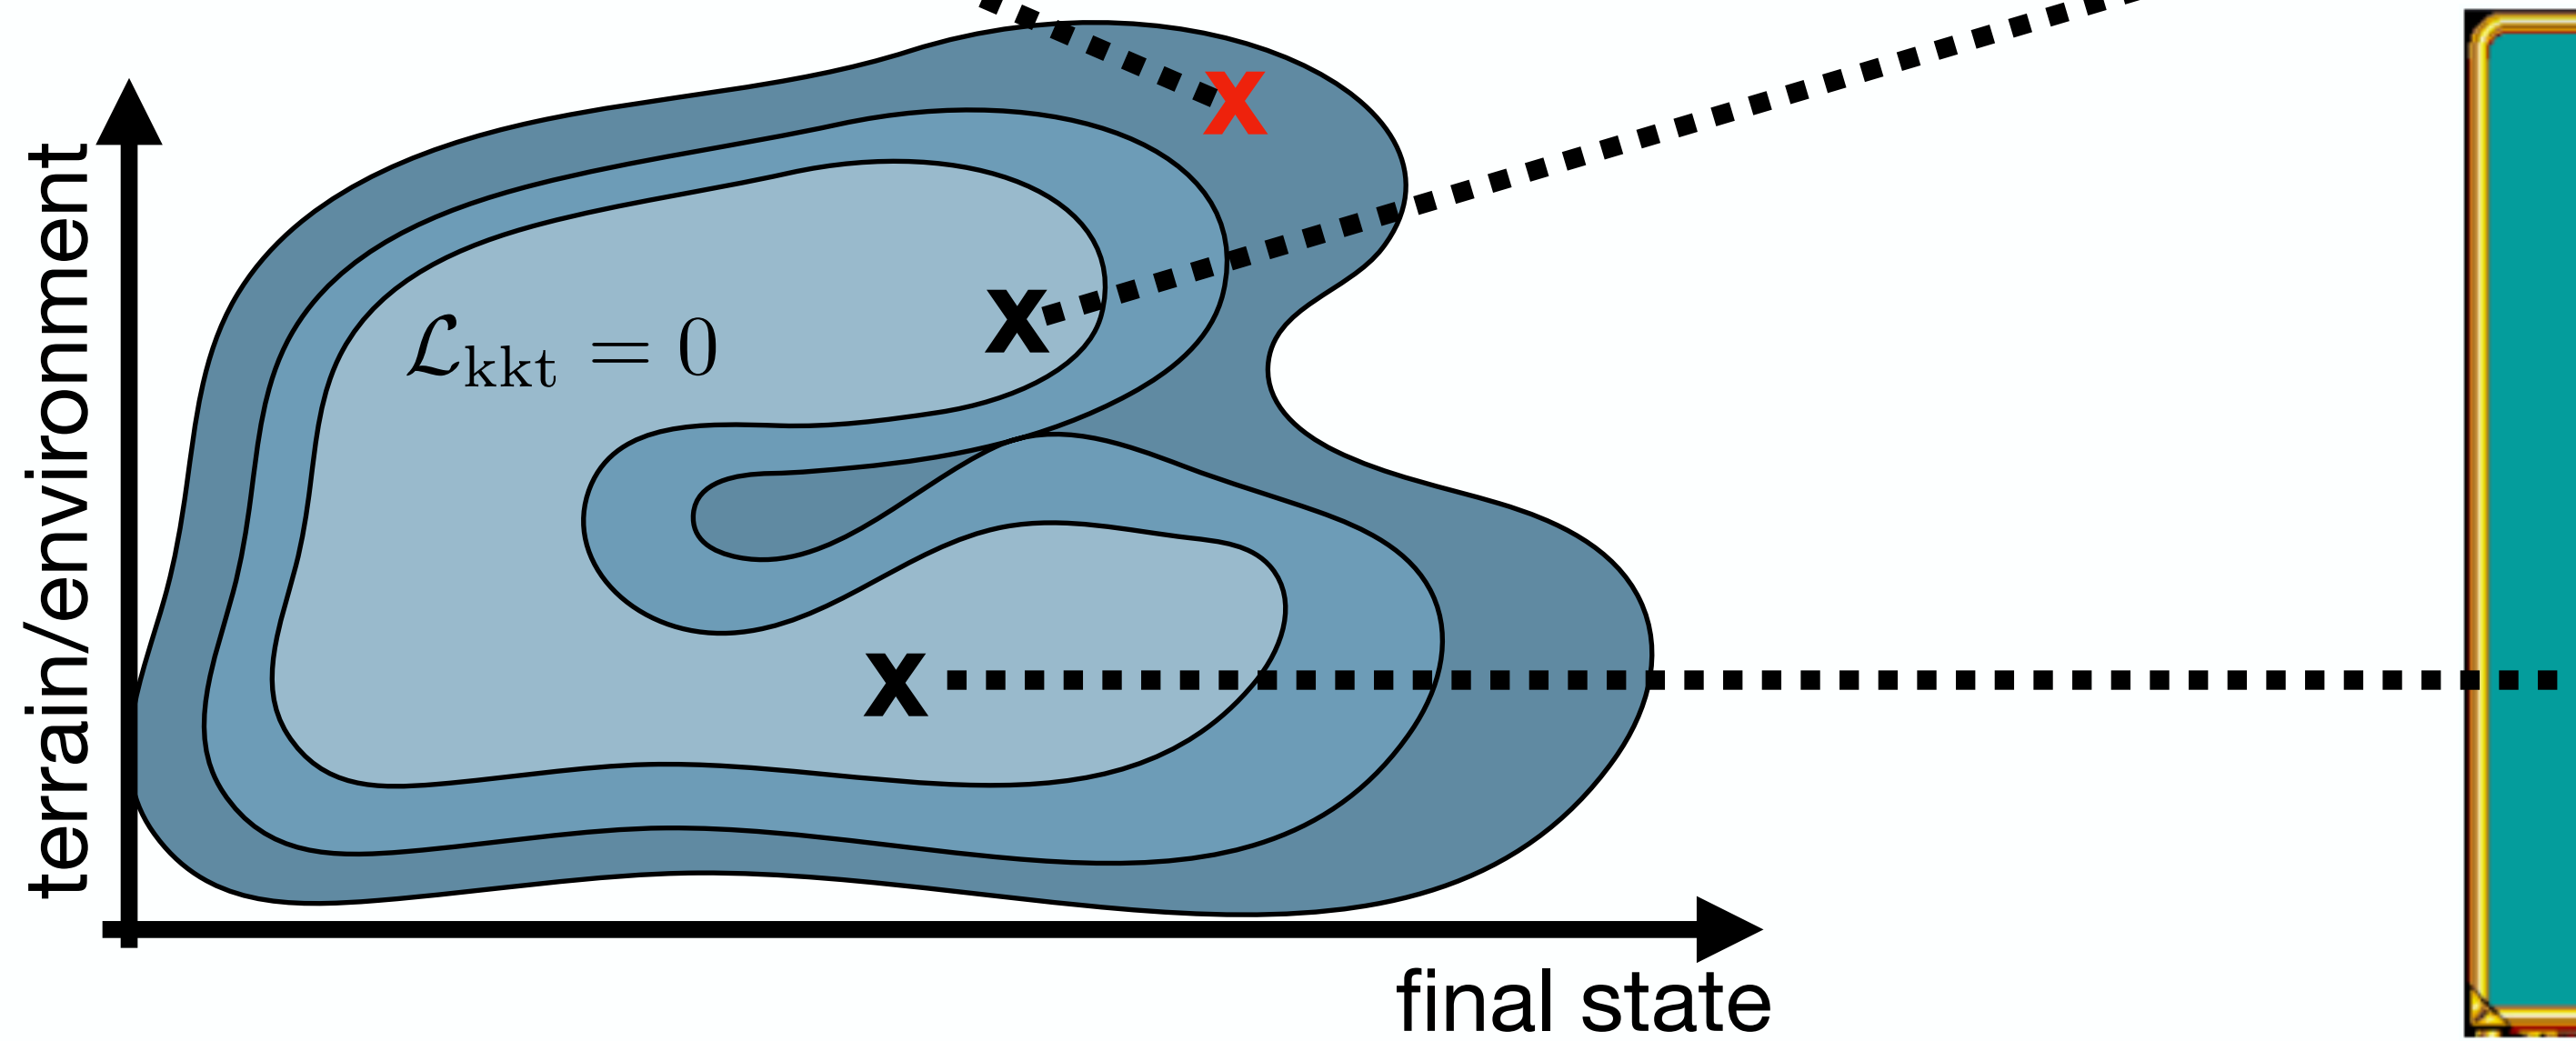







```
loss = loss_kkt(robot_pose, net(input), robot_model)
loss.backward()
```



# Incredible Machine (DOS 1993)




Huge subspace of terrains physically consistent with the final state (ball-in-bucket)




**Solution => Environment:** Hard, ill-conditioned, combinatorial optimisation problem  
a **solution** has huge number of **ways** that leads into it.  
e.g. all “Rube Goldberg’s Incredible Machine”

Incredible Machine (DOS 1993)



Rube Goldberg’s self-operating napkin (1931)

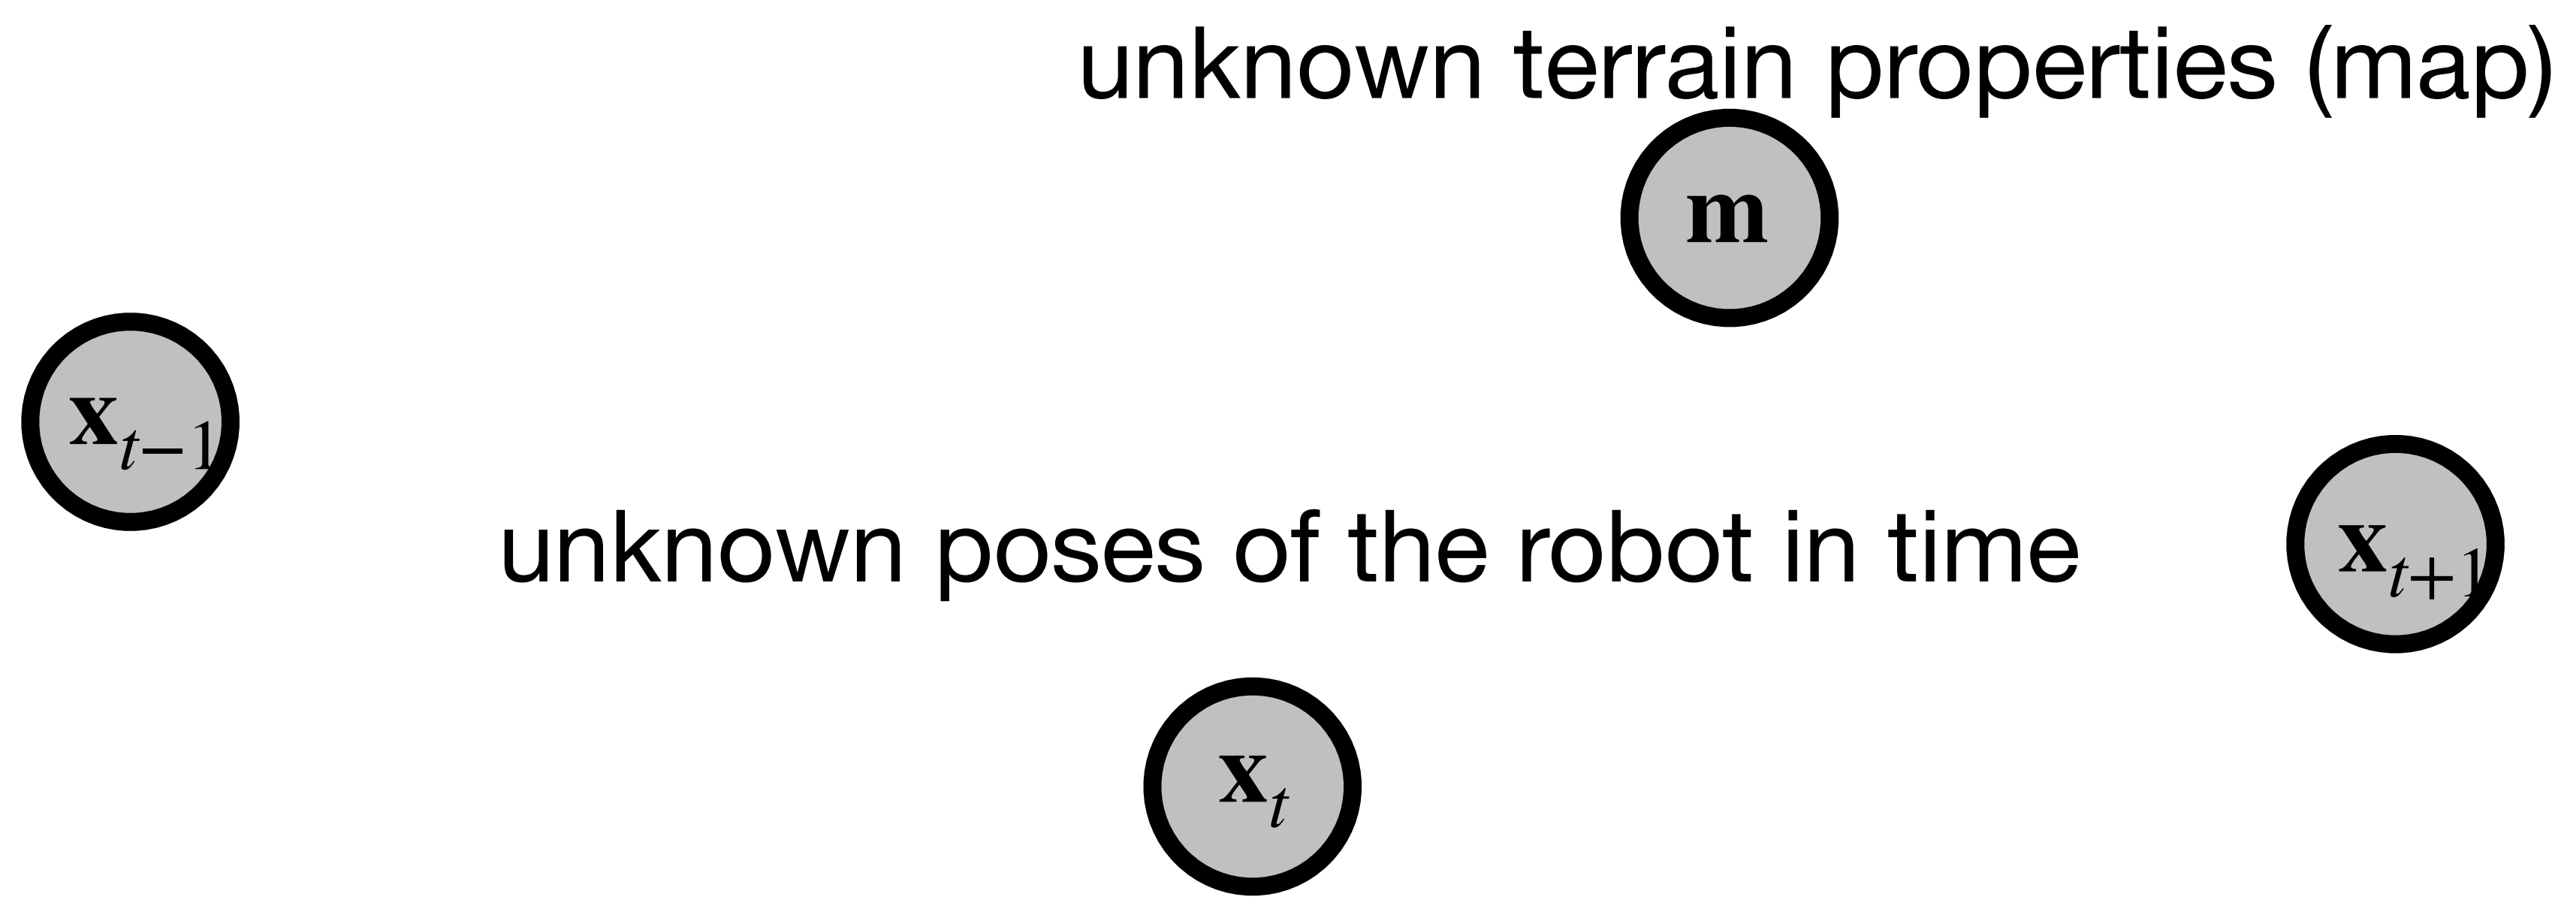


# Application II: robot-terrain-interaction-aware SLAM

$$\text{MAP: } \mathbf{x}^* = \arg \max_{\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m}} p(\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m} \mid \mathbf{z}_1 \dots \mathbf{z}_t, \mathbf{u}_1 \dots \mathbf{u}_t) = \arg \min_{\mathbf{x}_0, \dots, \mathbf{x}_t, \mathbf{m}} \sum_j f_j(\mathbf{x}, \mathbf{m}, \mathbf{z})^2 \Rightarrow \text{LM}$$

The equation is annotated with two hand-drawn shapes: a blue cloud-like shape containing the text "estimate this" pointing to the variables  $\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m}$  in the numerator, and a yellow cloud-like shape containing the text "given this" pointing to the variables  $\mathbf{z}_1 \dots \mathbf{z}_t, \mathbf{u}_1 \dots \mathbf{u}_t$  in the denominator.

Factorgraph: graphical model of conditional independence among unknown variables




# Application II: robot-terrain-interaction-aware SLAM

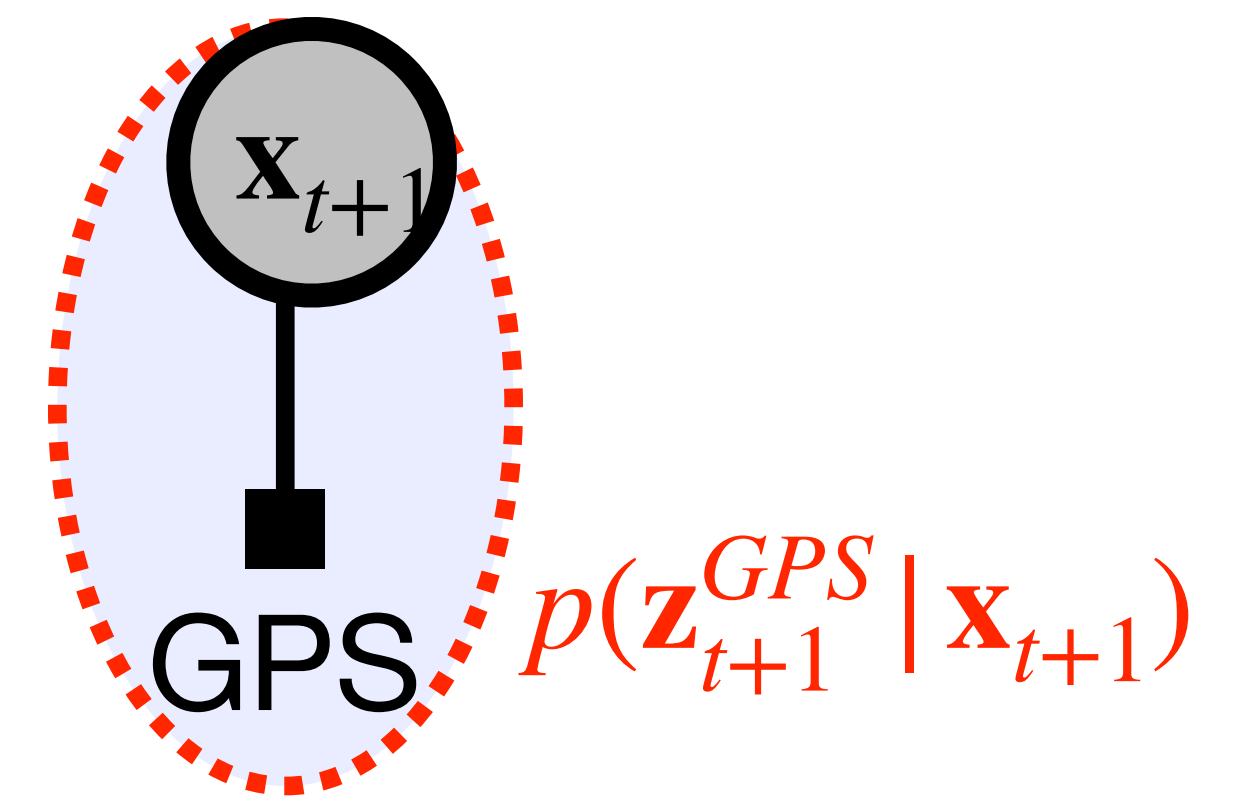
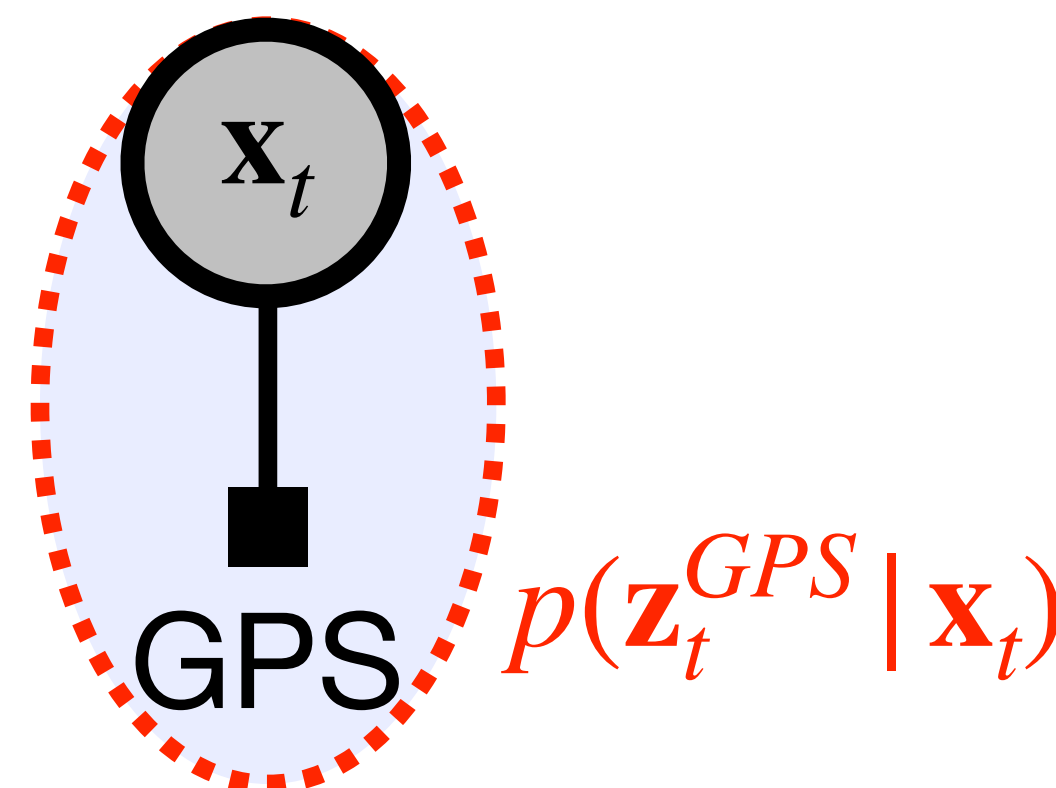
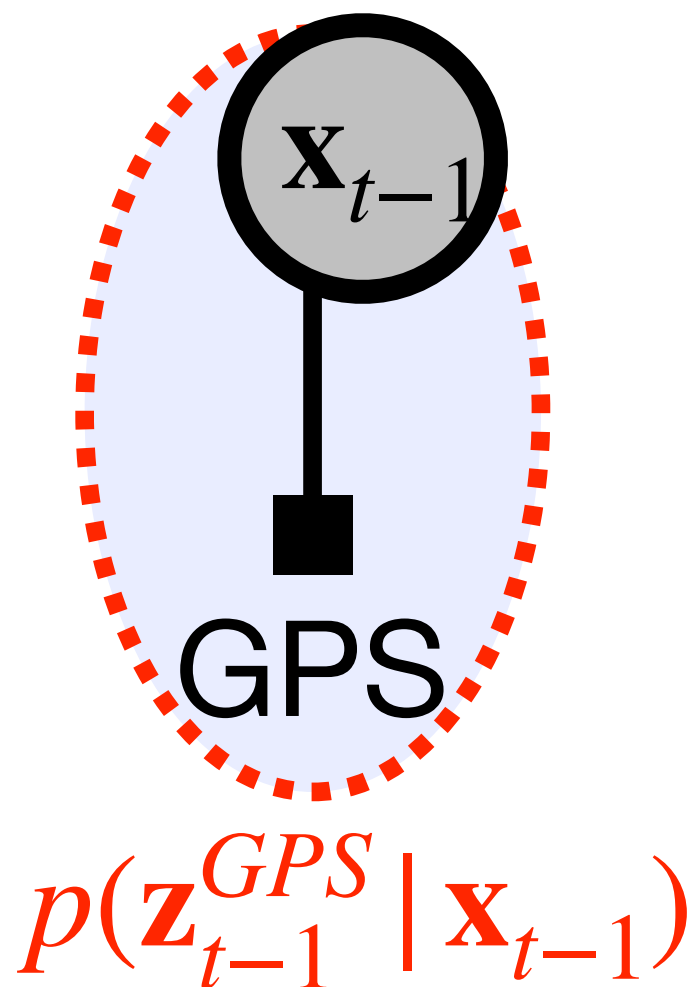
$$\text{MAP: } \mathbf{x}^* = \arg \max_{\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m}} p(\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m} \mid \mathbf{z}_1 \dots \mathbf{z}_t, \mathbf{u}_1 \dots \mathbf{u}_t) = \arg \min_{\mathbf{x}_0, \dots, \mathbf{x}_t, \mathbf{m}} \sum_j f_j(\mathbf{x}, \mathbf{m}, \mathbf{z})^2 \Rightarrow \text{LM}$$

Factorgraph: graphical model of conditional independence among unknown variables

unknown terrain properties (map)



unknown poses of the robot in time

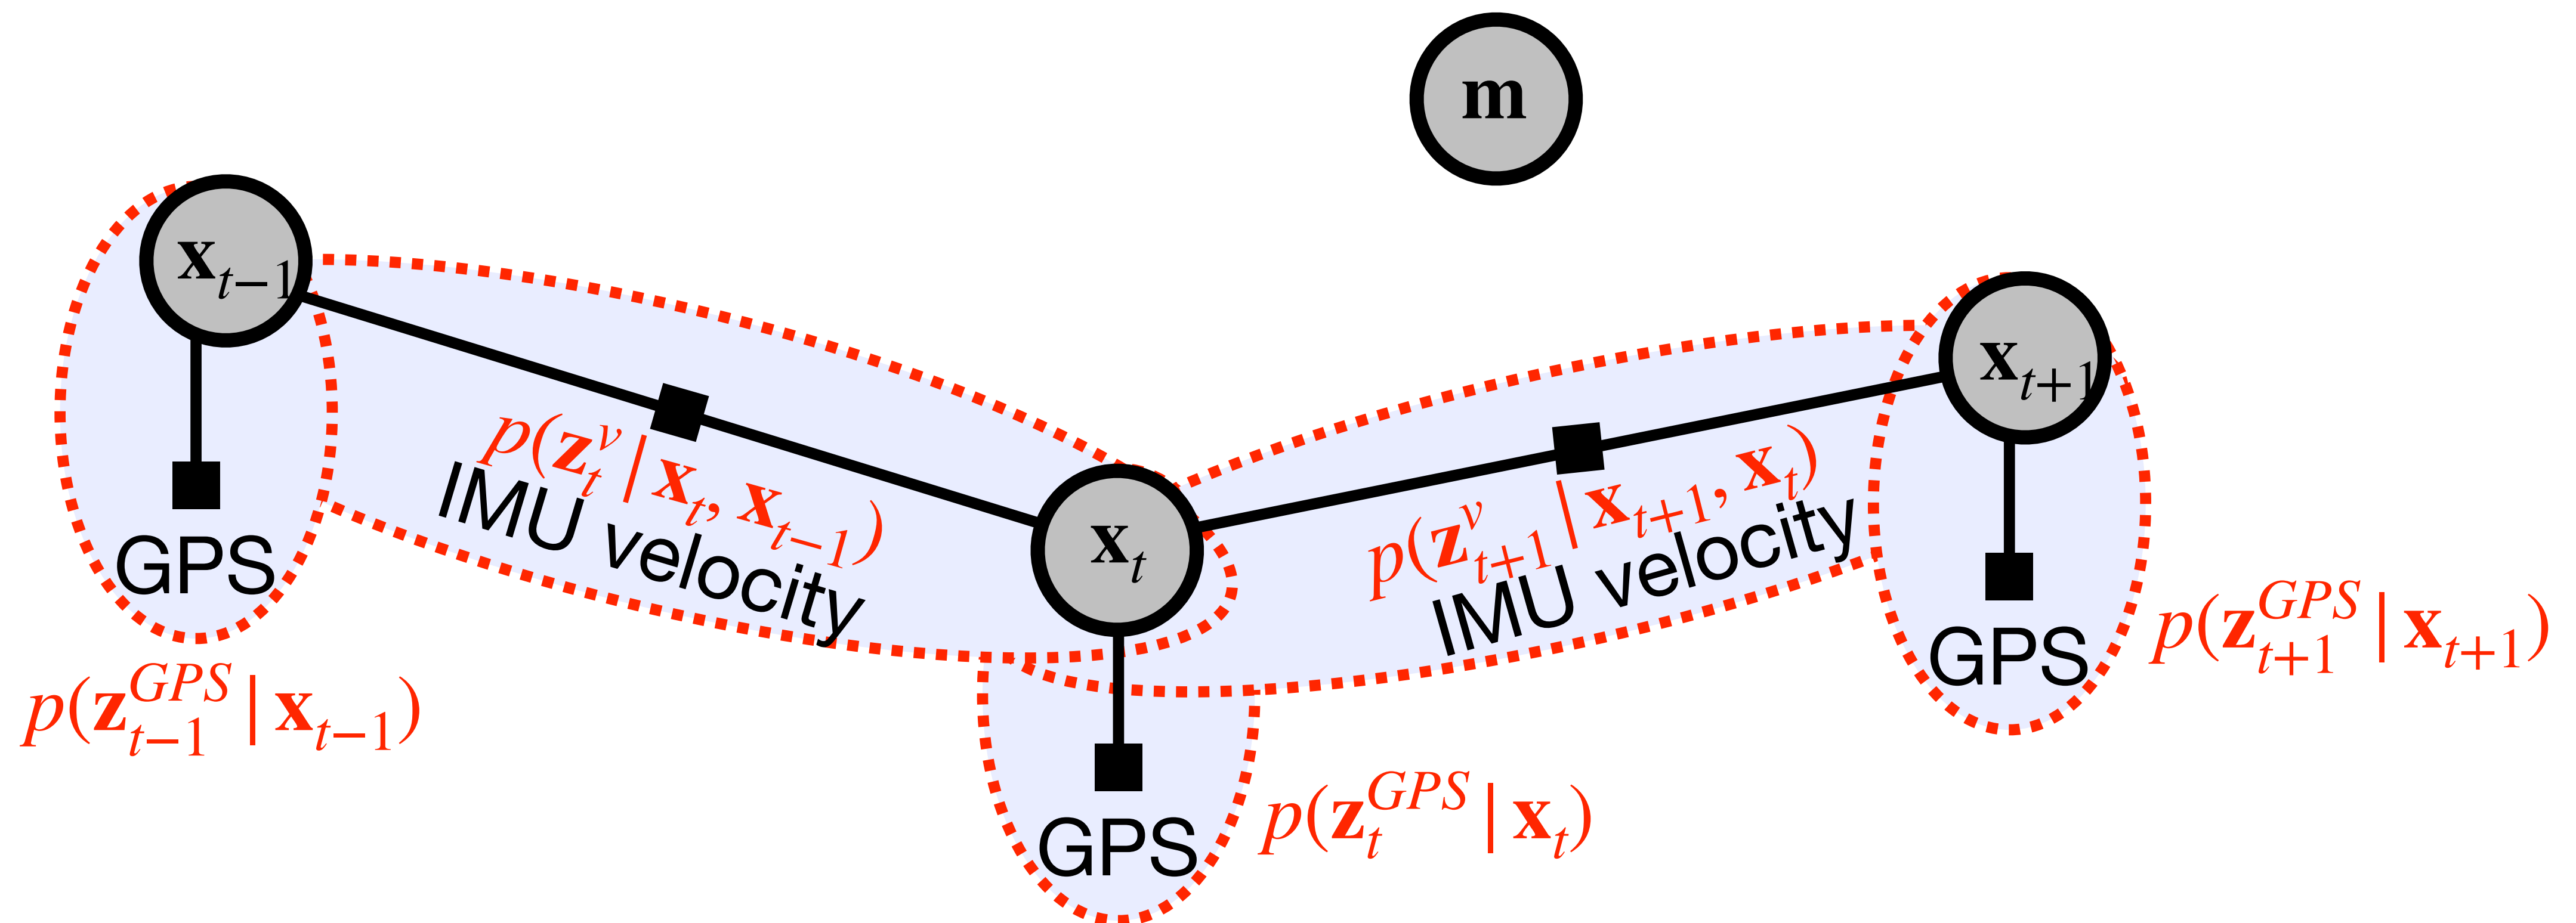




# Application II: robot-terrain-interaction-aware SLAM

$$\text{MAP: } \mathbf{x}^* = \arg \max_{\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m}} p(\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m} \mid \mathbf{z}_1 \dots \mathbf{z}_t, \mathbf{u}_1 \dots \mathbf{u}_t) = \arg \min_{\mathbf{x}_0, \dots, \mathbf{x}_t, \mathbf{m}} \sum_j f_j(\mathbf{x}, \mathbf{m}, \mathbf{z})^2 \Rightarrow \text{LM}$$

Factorgraph: graphical model of conditional independence among unknown variables


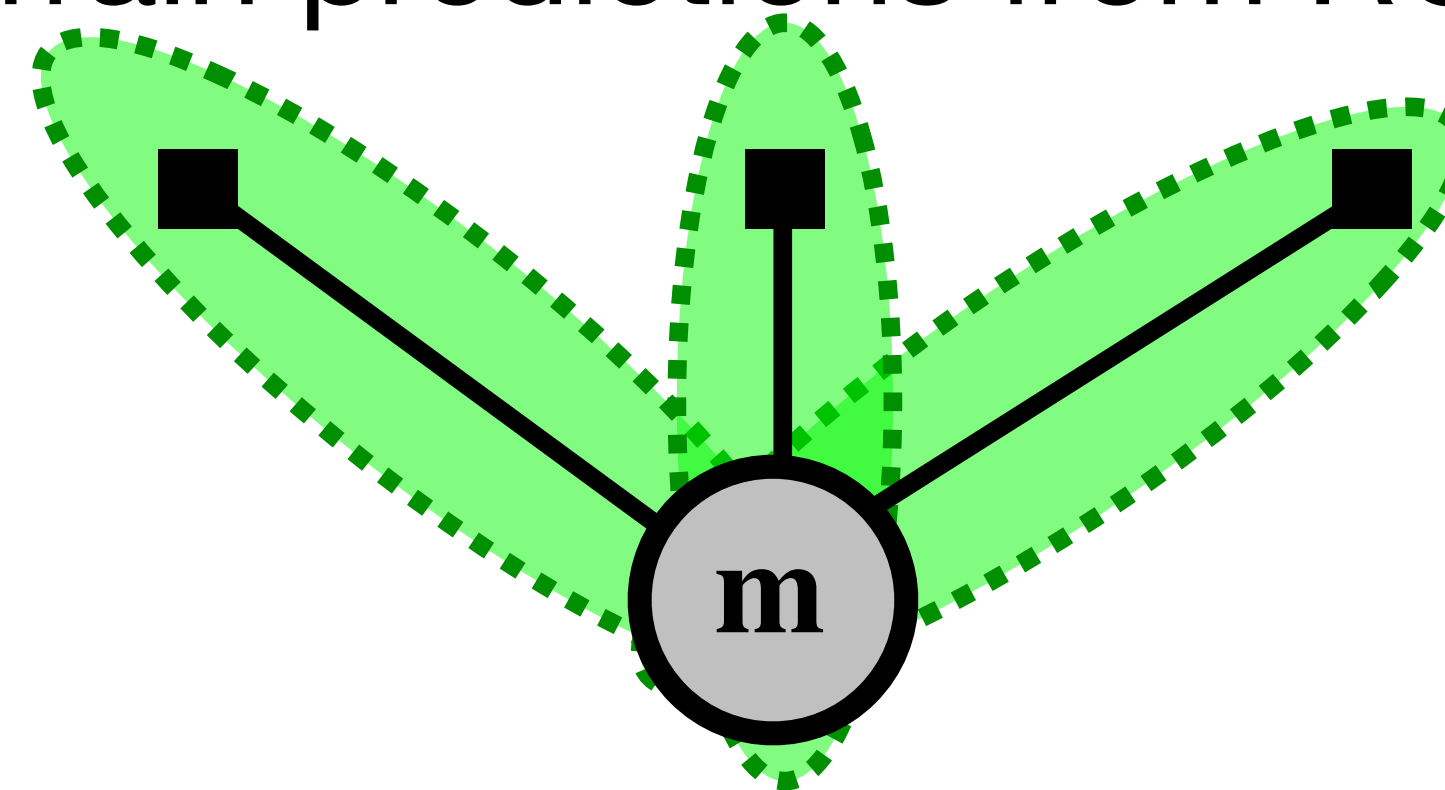




# Application II: robot-terrain-interaction-aware SLAM

MAP:  $\mathbf{x}^* = \arg \max_{\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m}} p(\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m} \mid \mathbf{z}_1 \dots \mathbf{z}_t, \mathbf{u}_1 \dots \mathbf{u}_t) = \arg \min_{\mathbf{x}_0, \dots, \mathbf{x}_t, \mathbf{m}} \sum_j f_j(\mathbf{x}, \mathbf{m}, \mathbf{z})^2 \Rightarrow \text{LM}$

terrain predictions from RGBD



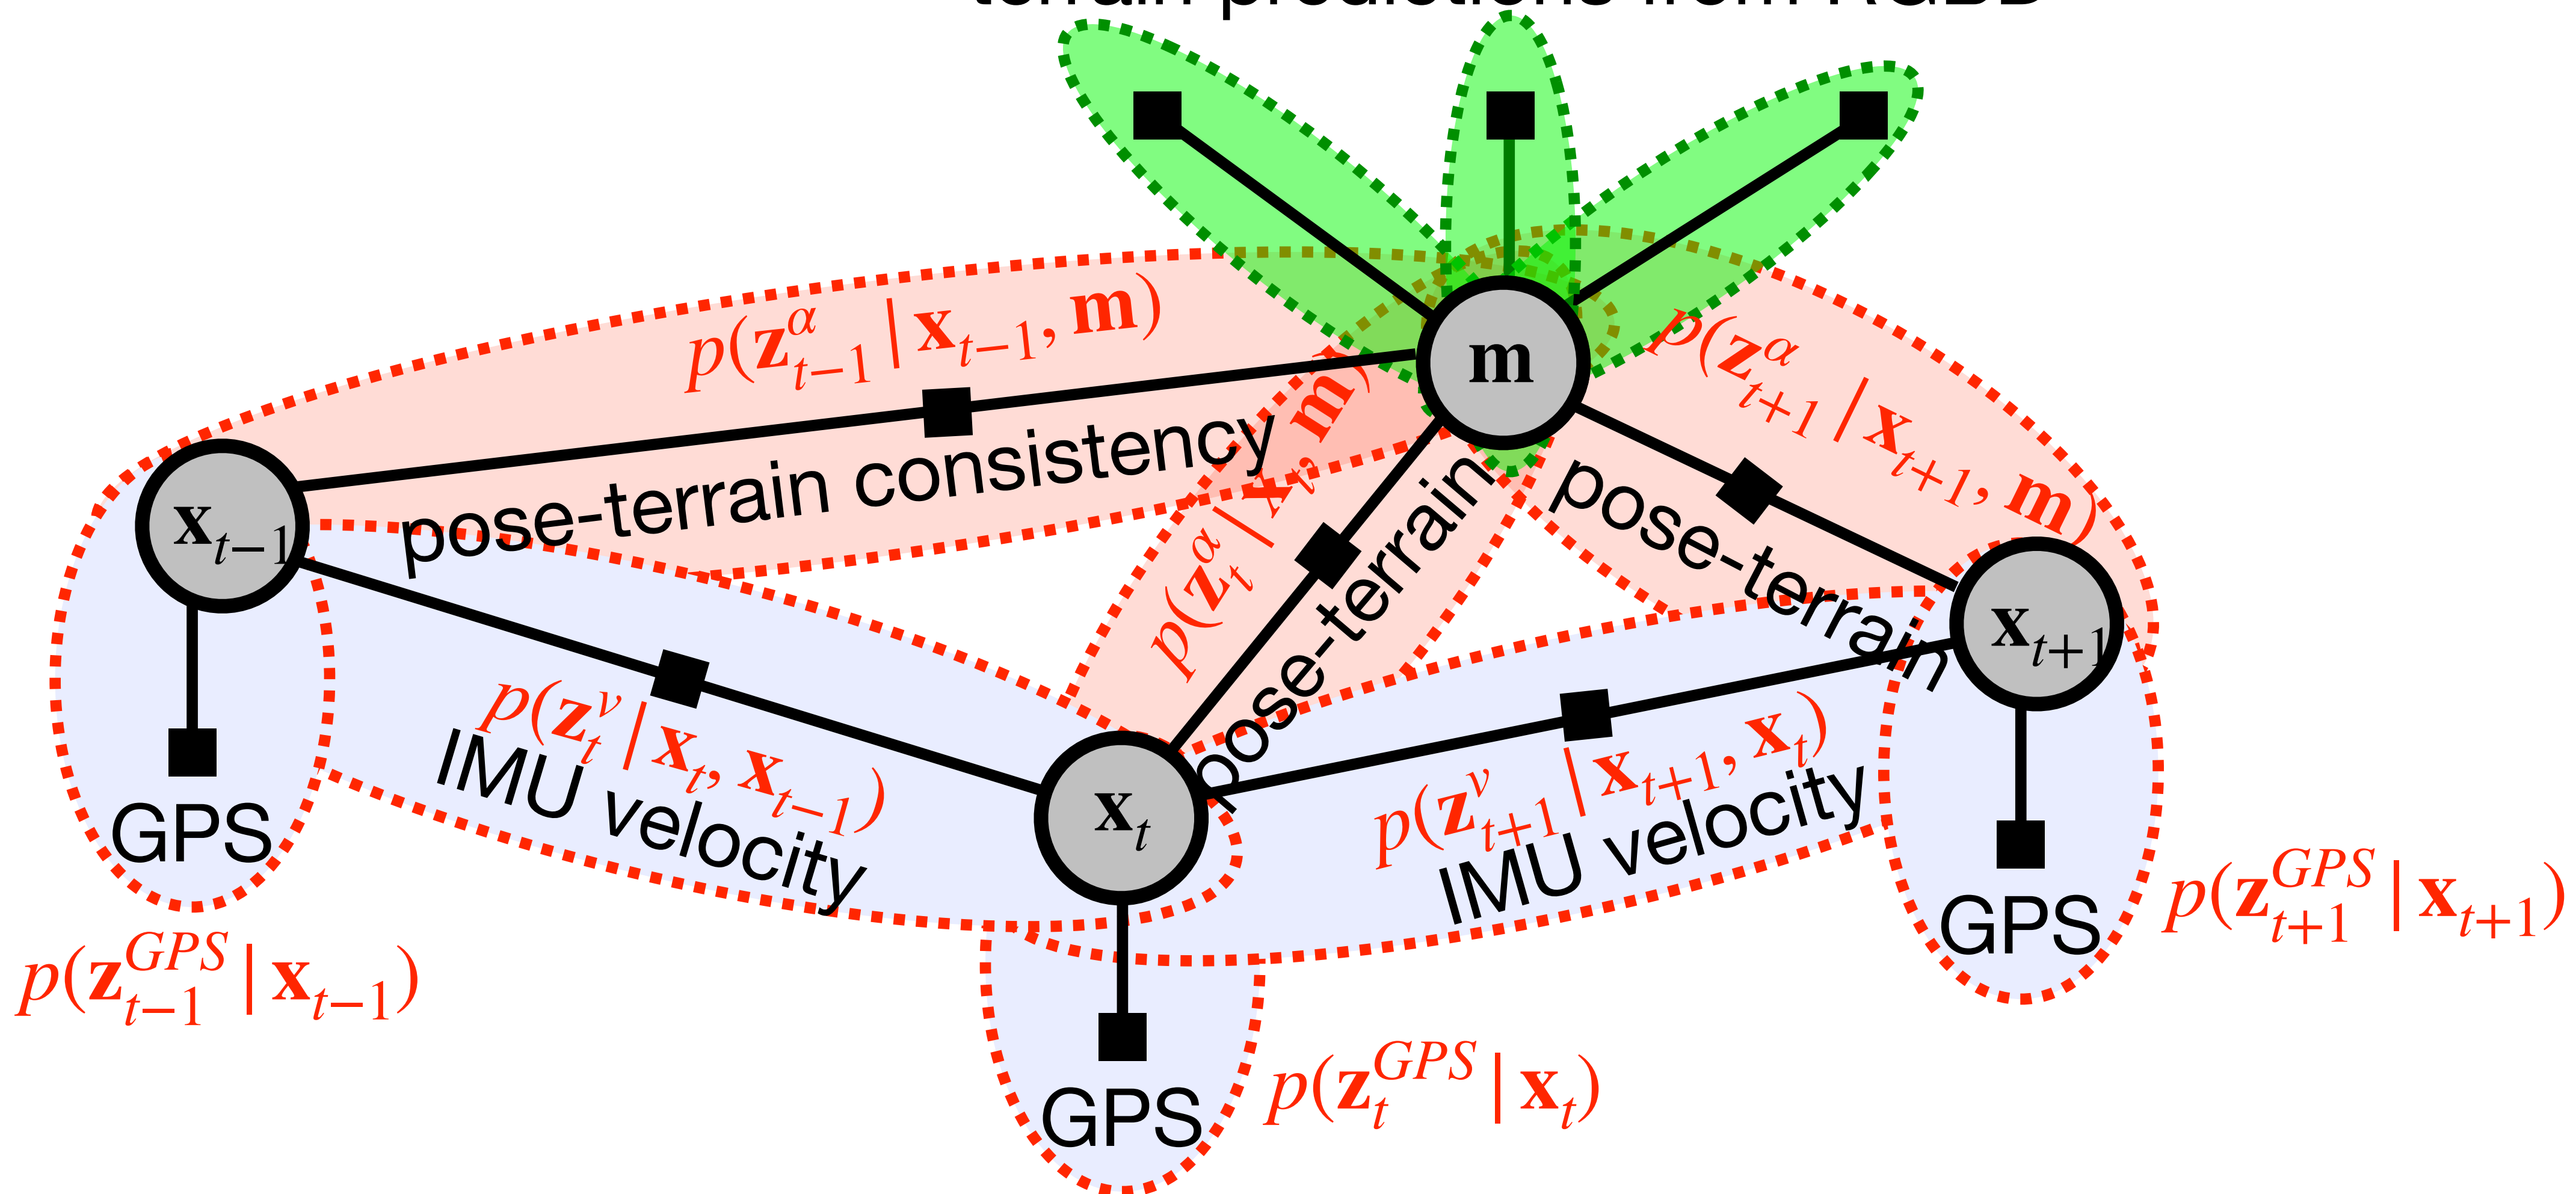


# Application II: robot-terrain-interaction-aware SLAM

MAP:  $\mathbf{x}^* = \arg \max_{\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m}} p(\mathbf{x}_0 \dots \mathbf{x}_t, \mathbf{m} \mid \mathbf{z}_1 \dots \mathbf{z}_t, \mathbf{u}_1 \dots \mathbf{u}_t)$  =  $\arg \min_{\mathbf{x}_0, \dots, \mathbf{x}_t, \mathbf{m}} \sum_j f_j(\mathbf{x}, \mathbf{m}, \mathbf{z})^2 \Rightarrow \text{LM}$

estimate this given this

terrain predictions from RGBD





# Application II: robot-terrain-interaction-aware SLAM

Given measurement:

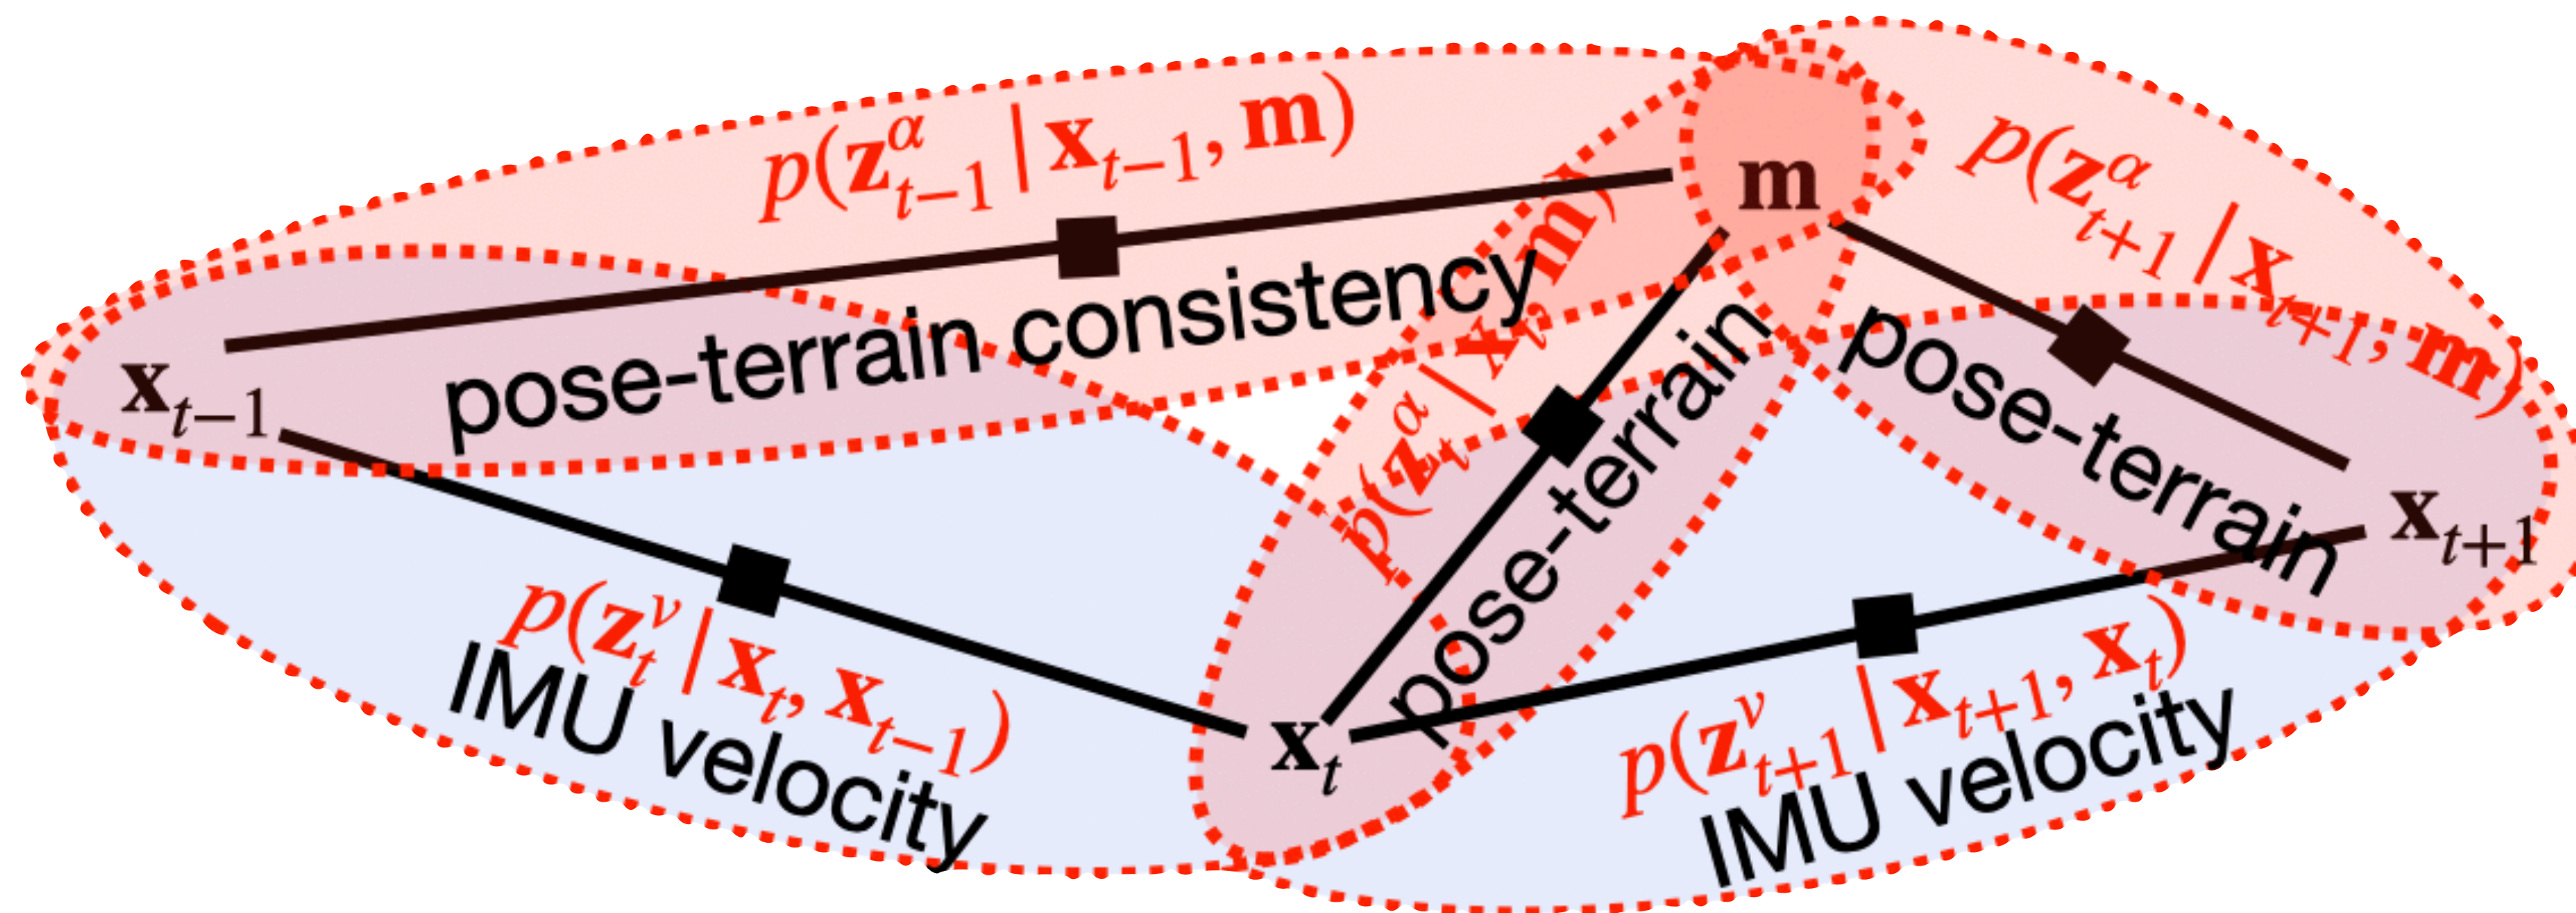
roll, pitch, yaw,  
track\_vel, flippers

$\mathbf{z}_1^\alpha, \mathbf{z}_1^\nu, \dots, \mathbf{z}_T^\alpha, \mathbf{z}_T^\nu$



Joint optimization of

$\mathbf{m}, \mathbf{x}_1, \dots, \mathbf{x}_T$ :




No GPS, no camera, no LIDAR  
(just angular velocities and tracks velocities)




Given measurement:  
roll, pitch, yaw,  
track\_vel, flippers


$L_{kkt} = 100.000000$



Joint optimization of  
 $\mathbf{m}, \mathbf{x}_1 \dots, \mathbf{x}_T$ :



Blind SLAM with **local** loop-closure






Lkkt=0.05023

Given measurement:  
roll, pitch, yaw,  
track\_vel, flippers

Joint optimization of  
 $\mathbf{m}, \mathbf{x}_1 \dots, \mathbf{x}_T$ :




0.0500

*y - rows*


2 25

Blind SLAM with **global** loop-closure

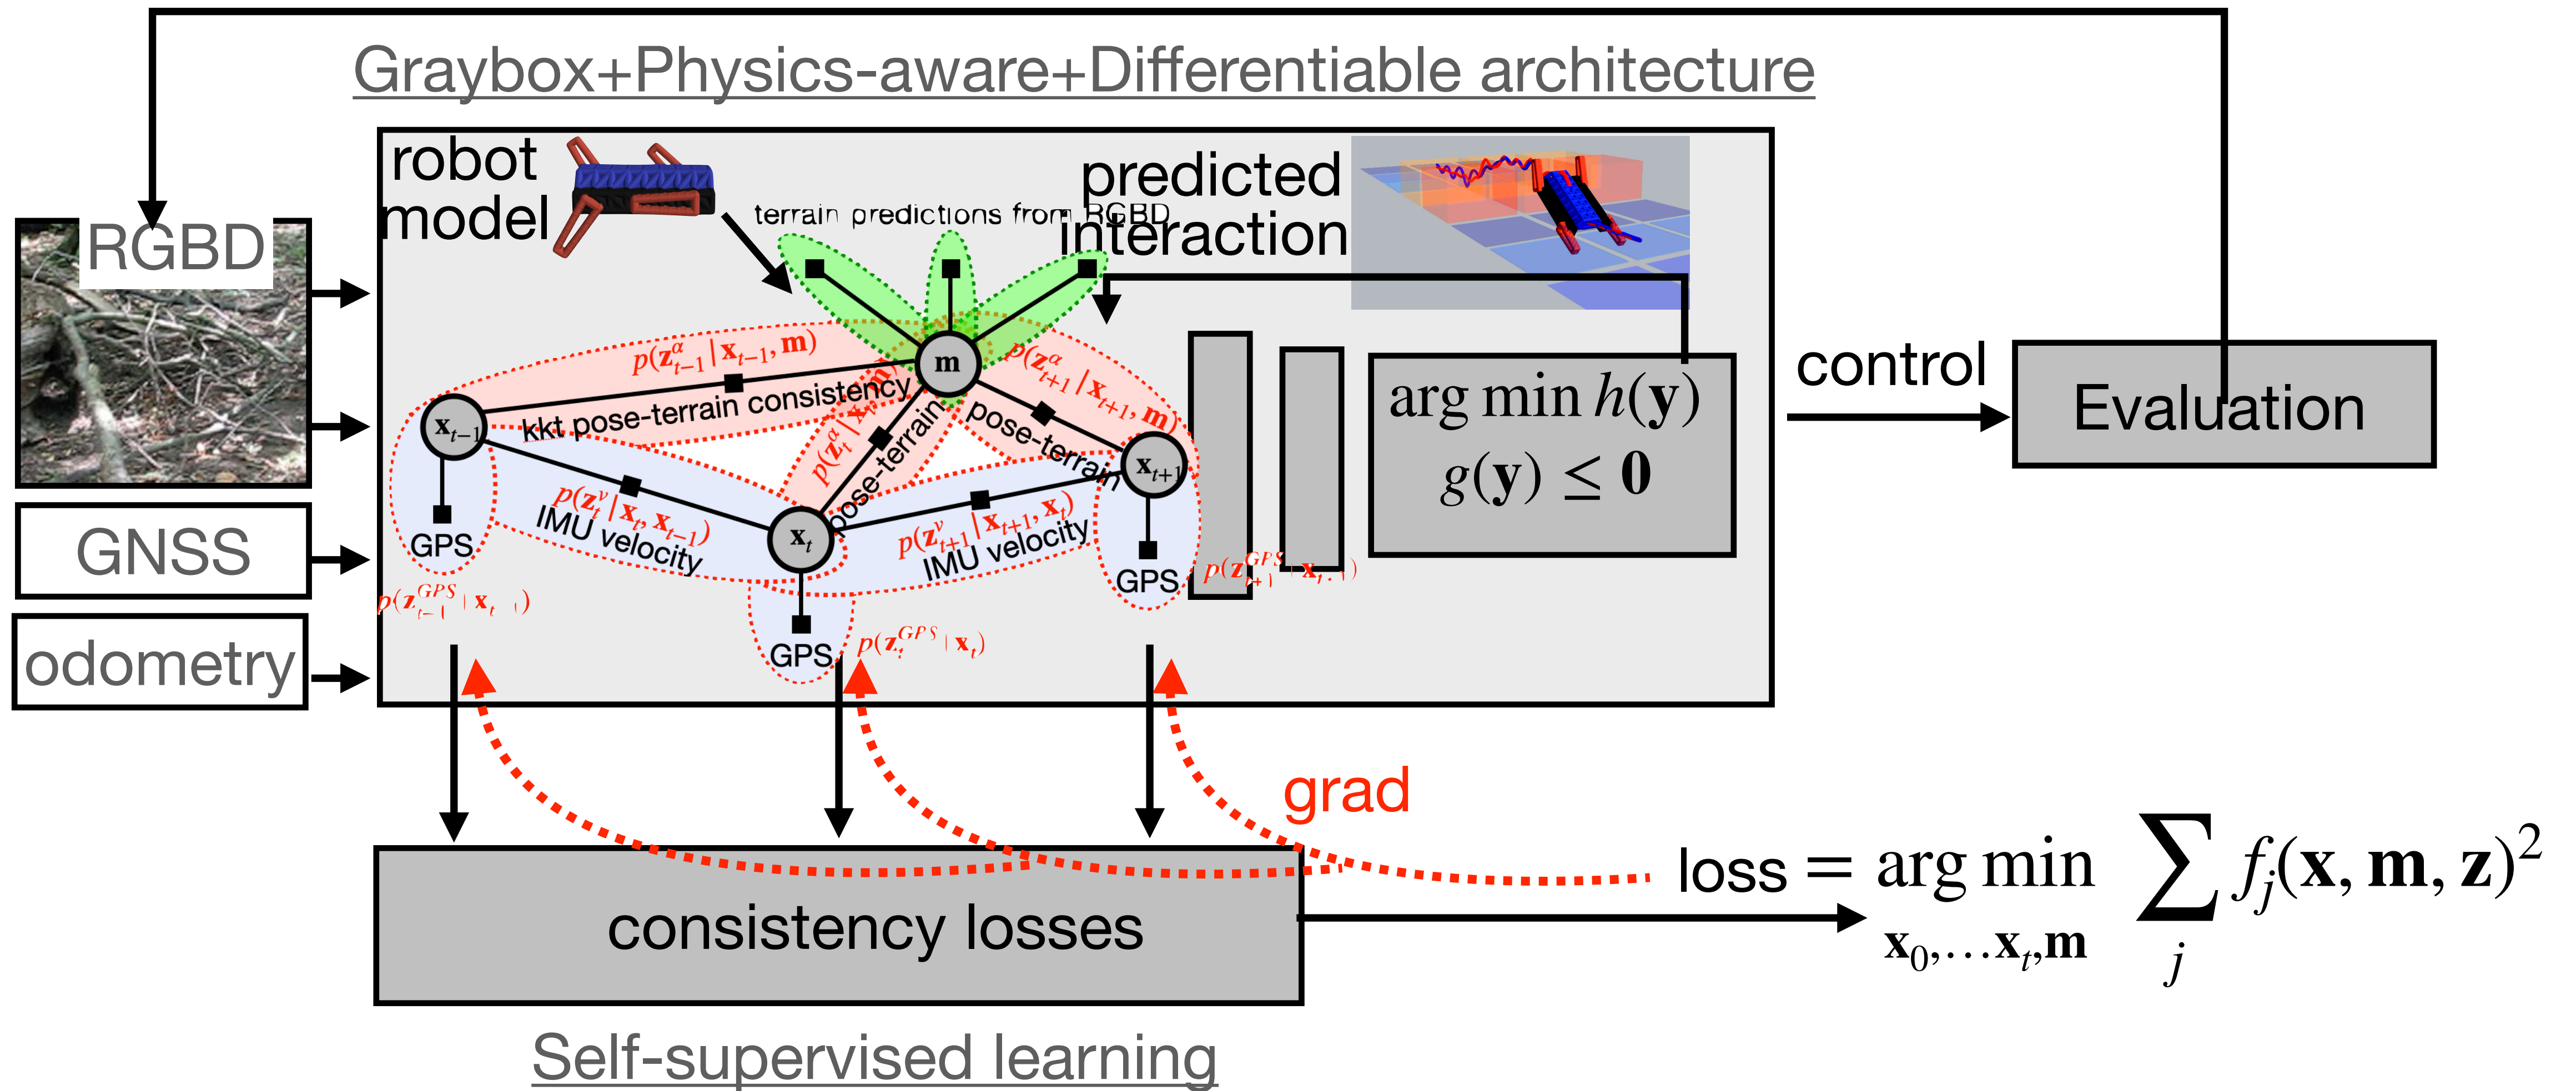




# Application III: self-supervision from factorgraph



# Application III: self-supervision from factorgraph





# The take-home message

- Differentiable graybox architecture is **less prone to overfitting**
- Differentiable graybox architecture **enable self-supervision**
- ARO course:
  - builds on top of fundamental principles from KyR theoretical courses (LAG, LGR, MA1/2, PST, OPT, ALP, KUI, ...)

1	Matematická analýza 1	Lineární algebra	Logika a grafy	Algoritmy a programování	Roboti	
2	Matematická analýza 2	Fyzika 1	Dif. rovnice a numerika	Programování v C	Kybernetika a umělá inteligence	
3	Komplexní analýza a transformace	Fyzika 2	Pravděpodobnost a statistika	Elektronické prvky a obvody	Signály, systémy a inference	
4	Prez. doved.	Automatické řízení	Logické systémy a procesory	Senzory a měření	Laboratoře	Povinně volitelný předmět
5	Optimalizace	Komunikace a distribuované systémy	Robotika	Projekt	Povinně volitelný předmět	
6	Bakalářská práce			Humanitní, umělecký a společenskovední seminář	Volitelný předmět	

- Build probabilistic model=> criterion => suggest solution => critically evaluate
- **Don't be satisfied with a solution** => always try to find less Ruby Goldberg's solution (more Interpretable/Explainable/Lepricon-free algorithm).