The Power of Gradients in Inverse Dynamics Problems

Tao Du

MIT CSAIL

What is a dynamic system?

Robotics
ETH Zurich, ANYmal

Molecular Biology
SCLS

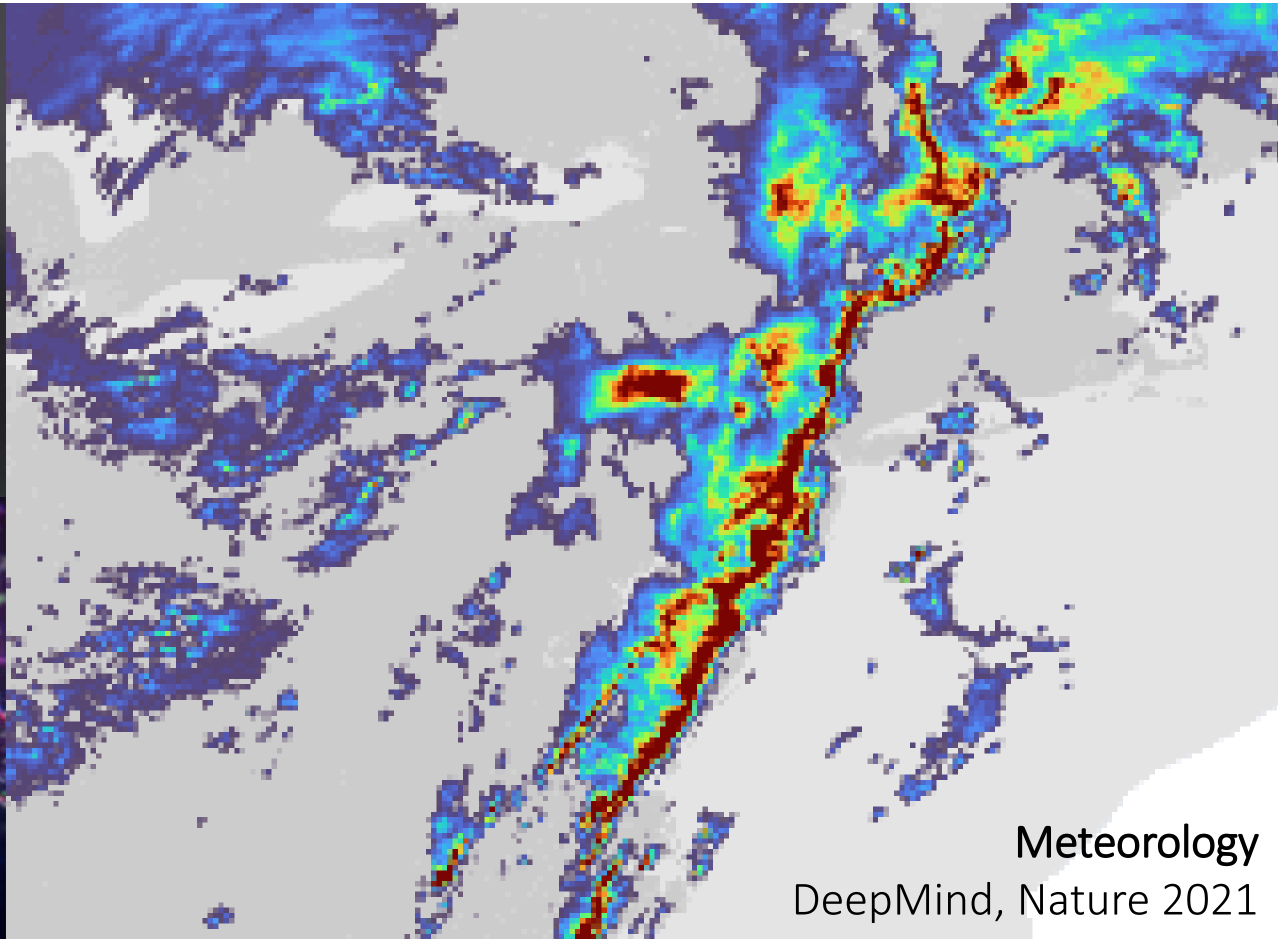

Meteorology
DeepMind, Nature 2021

What is a dynamic system?

"A dynamical system is particle or ensemble of particles whose state varies over time and thus obeys differential equations involving time derivatives."

---Nature Portfolio

What is a dynamic system?

States	Time derivatives
$s,$	$\frac{ds}{dt}, \frac{d^2s}{dt^2}, \dots$

What is a dynamic system?

Dynamic model

$$F\left(s, \quad \frac{ds}{dt}, \frac{d^2 s}{dt^2}, \cdots \right) = 0$$

What is a dynamic system?

Dynamic model

$$F\left(s, \frac{ds}{dt}, \frac{d^2s}{dt^2}, \cdots \right) = 0$$

Example

Rigid-body systems: Euler-Lagrange equation

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}}\right) - \frac{\partial T}{\partial q} - Q = 0$$

Deformable objects: continuum mechanics

$$\nabla \cdot \sigma + f = 0$$

Fluid systems: Navier-Stokes equation

$$\frac{du}{dt} + (u \cdot \nabla)u - \nu\nabla^2 u = -\frac{1}{\rho}\nabla p + g$$

The input and the output

Input

Parameters

Dynamic model

$$F(\quad s, \qquad \frac{ds}{dt}, \frac{d^2 s}{dt^2}, \dots \quad) = 0$$

Example

Rigid-body systems: Euler-Lagrange equation

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}}\right) - \frac{\partial T}{\partial q} - Q = 0$$

Deformable objects: continuum mechanics

$$\nabla \cdot \sigma + f = 0$$

Fluid systems: Navier-Stokes equation

$$\frac{du}{dt} + (u \cdot \nabla)u - \nu \nabla^2 u = -\frac{1}{\rho}\nabla p + g$$

The input and the output

Input

Parameters

Dynamic model

$$F\left(s, \quad \frac{ds}{dt}, \frac{d^2 s}{dt^2}, \dots \quad \right) = 0$$

Example

Intrinsic parameters

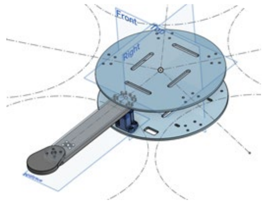

Extrinsic parameters

Example

Rigid-body systems: Euler-Lagrange equation

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}}\right) - \frac{\partial T}{\partial q} - Q = 0$$

Deformable objects: continuum mechanics

$$\nabla \cdot \sigma + f = 0$$

Fluid systems: Navier-Stokes equation

$$\frac{du}{dt} + (u \cdot \nabla)u - \nu \nabla^2 u = -\frac{1}{\rho}\nabla p + g$$

The input and the output

Input

Parameters

Dynamic model

$$F(\ s,\ \frac{ds}{dt}, \frac{d^2s}{dt^2}, \cdots\ ) = 0$$

Output

State sequences

Example

Intrinsic parameters

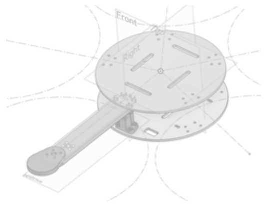

Extrinsic parameters

Example

Rigid-body systems: Euler-Lagrange equation

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}}\right) - \frac{\partial T}{\partial q} - Q = 0$$

Deformable objects: continuum mechanics

$$\nabla \cdot \sigma + f = 0$$

Fluid systems: Navier-Stokes equation

$$\frac{du}{dt} + (u \cdot \nabla)u - \nu \nabla^2 u = -\frac{1}{\rho}\nabla p + g$$

Example

States from simulation

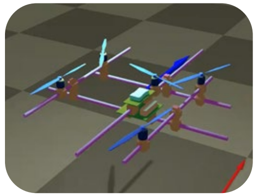

States from experiments

The forward dynamics problem

Given the model and parameters of a dynamic system, compute its state sequence.

The forward dynamics problem

Given the model and parameters of a dynamic system, compute its state sequence.

Parametrization

Initializing parameters

The forward dynamics problem

Given the model and parameters of a dynamic system, compute its state sequence.

Parametrization

Initializing parameters

Modeling

Deriving governing equations

The forward dynamics problem

Given the model and parameters of a dynamic system, compute its state sequence.

Parametrization

Initializing parameters

Modeling

Deriving governing equations

Evaluation

Computing performance metrics

The forward dynamics problem

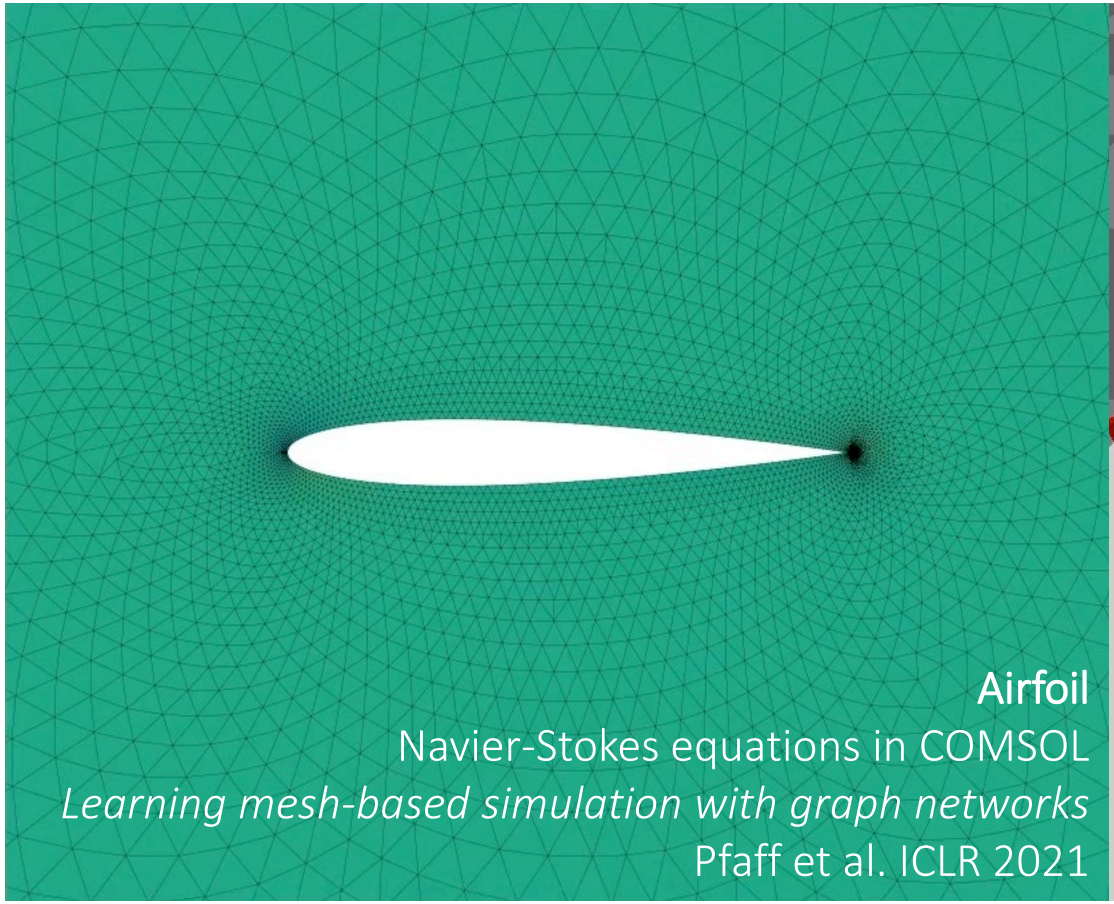

Airfoil
Navier-Stokes equations in COMSOL
Learning mesh-based simulation with graph networks
Pfaff et al. ICLR 2021

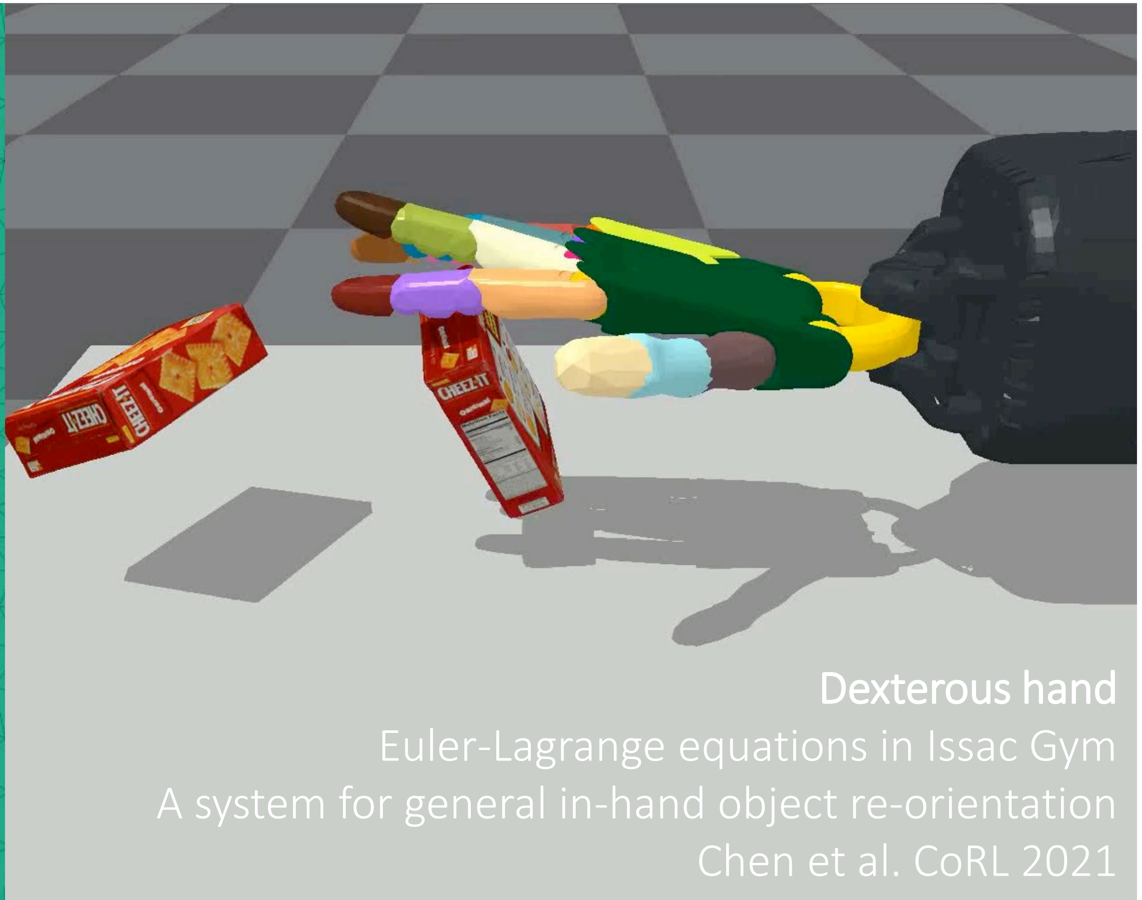

Dexterous hand
Euler-Lagrange equations in Issac Gym
A system for general in-hand object re-orientation
Chen et al. CoRL 2021

The inverse dynamics problem

Given the state sequence of a dynamic system, infer its model and parameters.

The inverse dynamics problem

Given the state sequence of a dynamic system, infer its model and parameters.

Parametrization

Initializing parameters

Modeling

Deriving governing equations

Evaluation

Computing performance metrics

The inverse dynamics problem

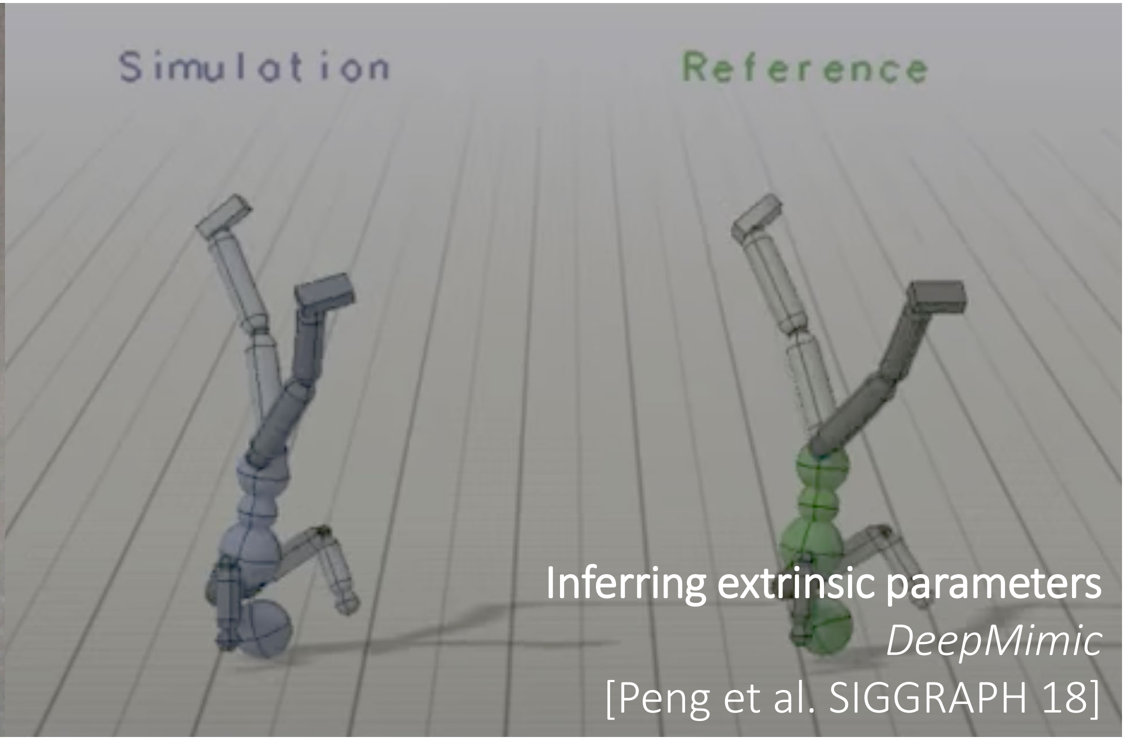

Our topic today: the gradient methodology

Parametrization

Initializing parameters

Modeling

Deriving governing equations

Evaluation

Computing performance metrics

Our topic today: the gradient methodology

∇ Parametrization

Initializing parameters

Modeling

Deriving governing equations

Evaluation

Computing performance metrics

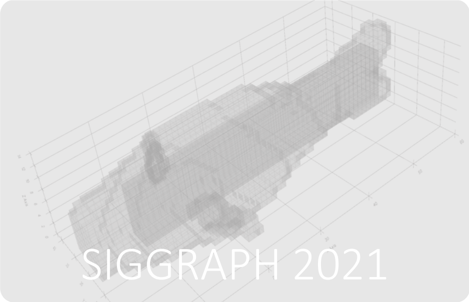

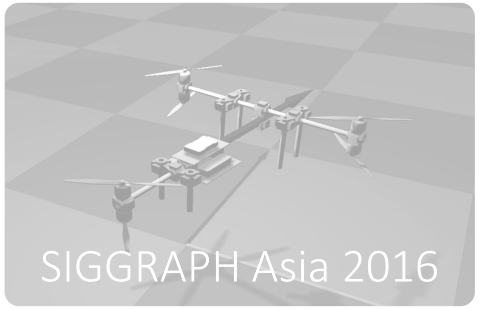

Our topic today: the gradient methodology

∇ Parametrization

Initializing parameters

∇ Modeling

Deriving governing equations

Evaluation

Computing performance metrics

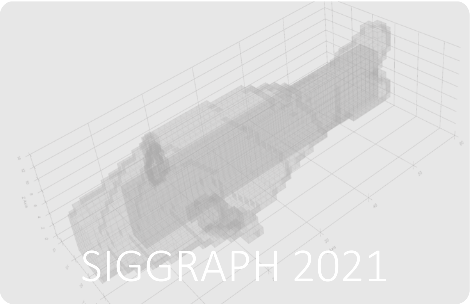
SIGGRAPH 2021

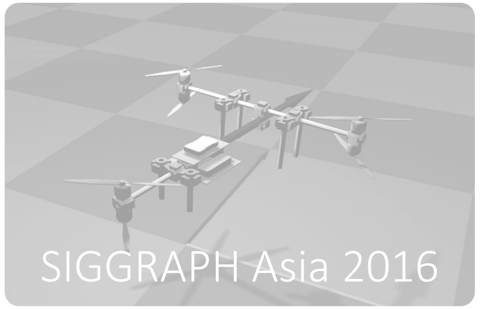
SIGGRAPH Asia 2016

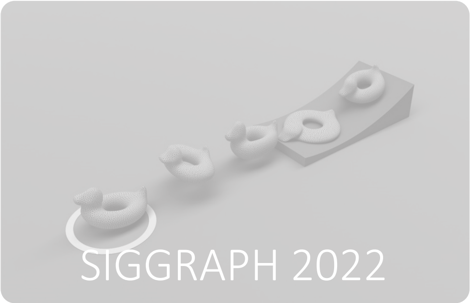
SIGGRAPH 2022


SIGGRAPH 2022

Our topic today: the gradient methodology

∇ Parametrization

Initializing parameters

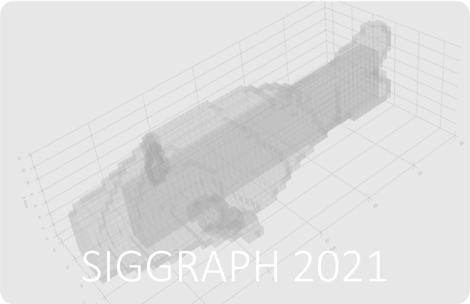

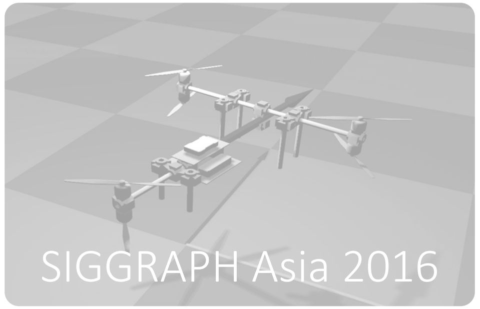

∇ Modeling

Deriving governing equations

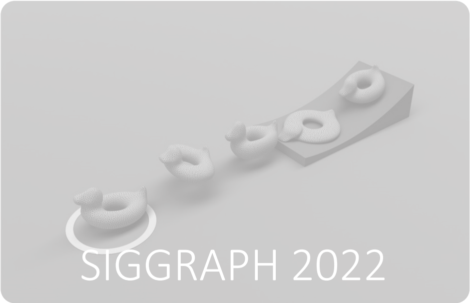

∇ Evaluation

Computing performance metrics

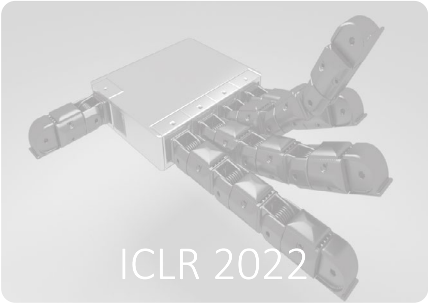

Our topic today: the gradient methodology

∇ Parametrization

Initializing parameters

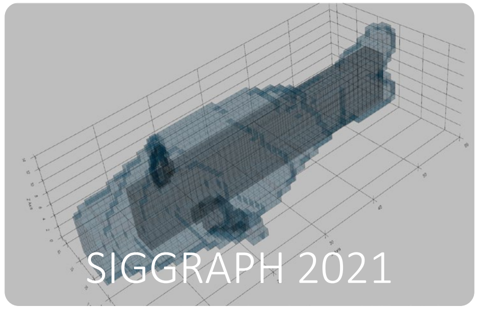

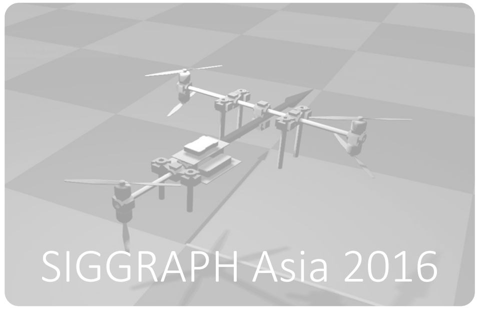

∇ Modeling

Deriving governing equations

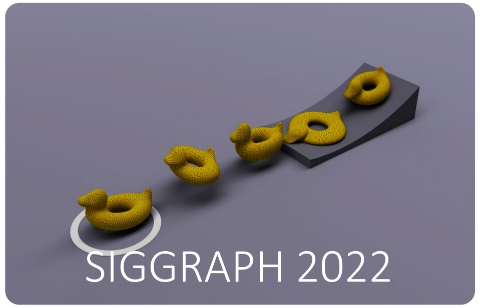

∇ Evaluation

Computing performance metrics

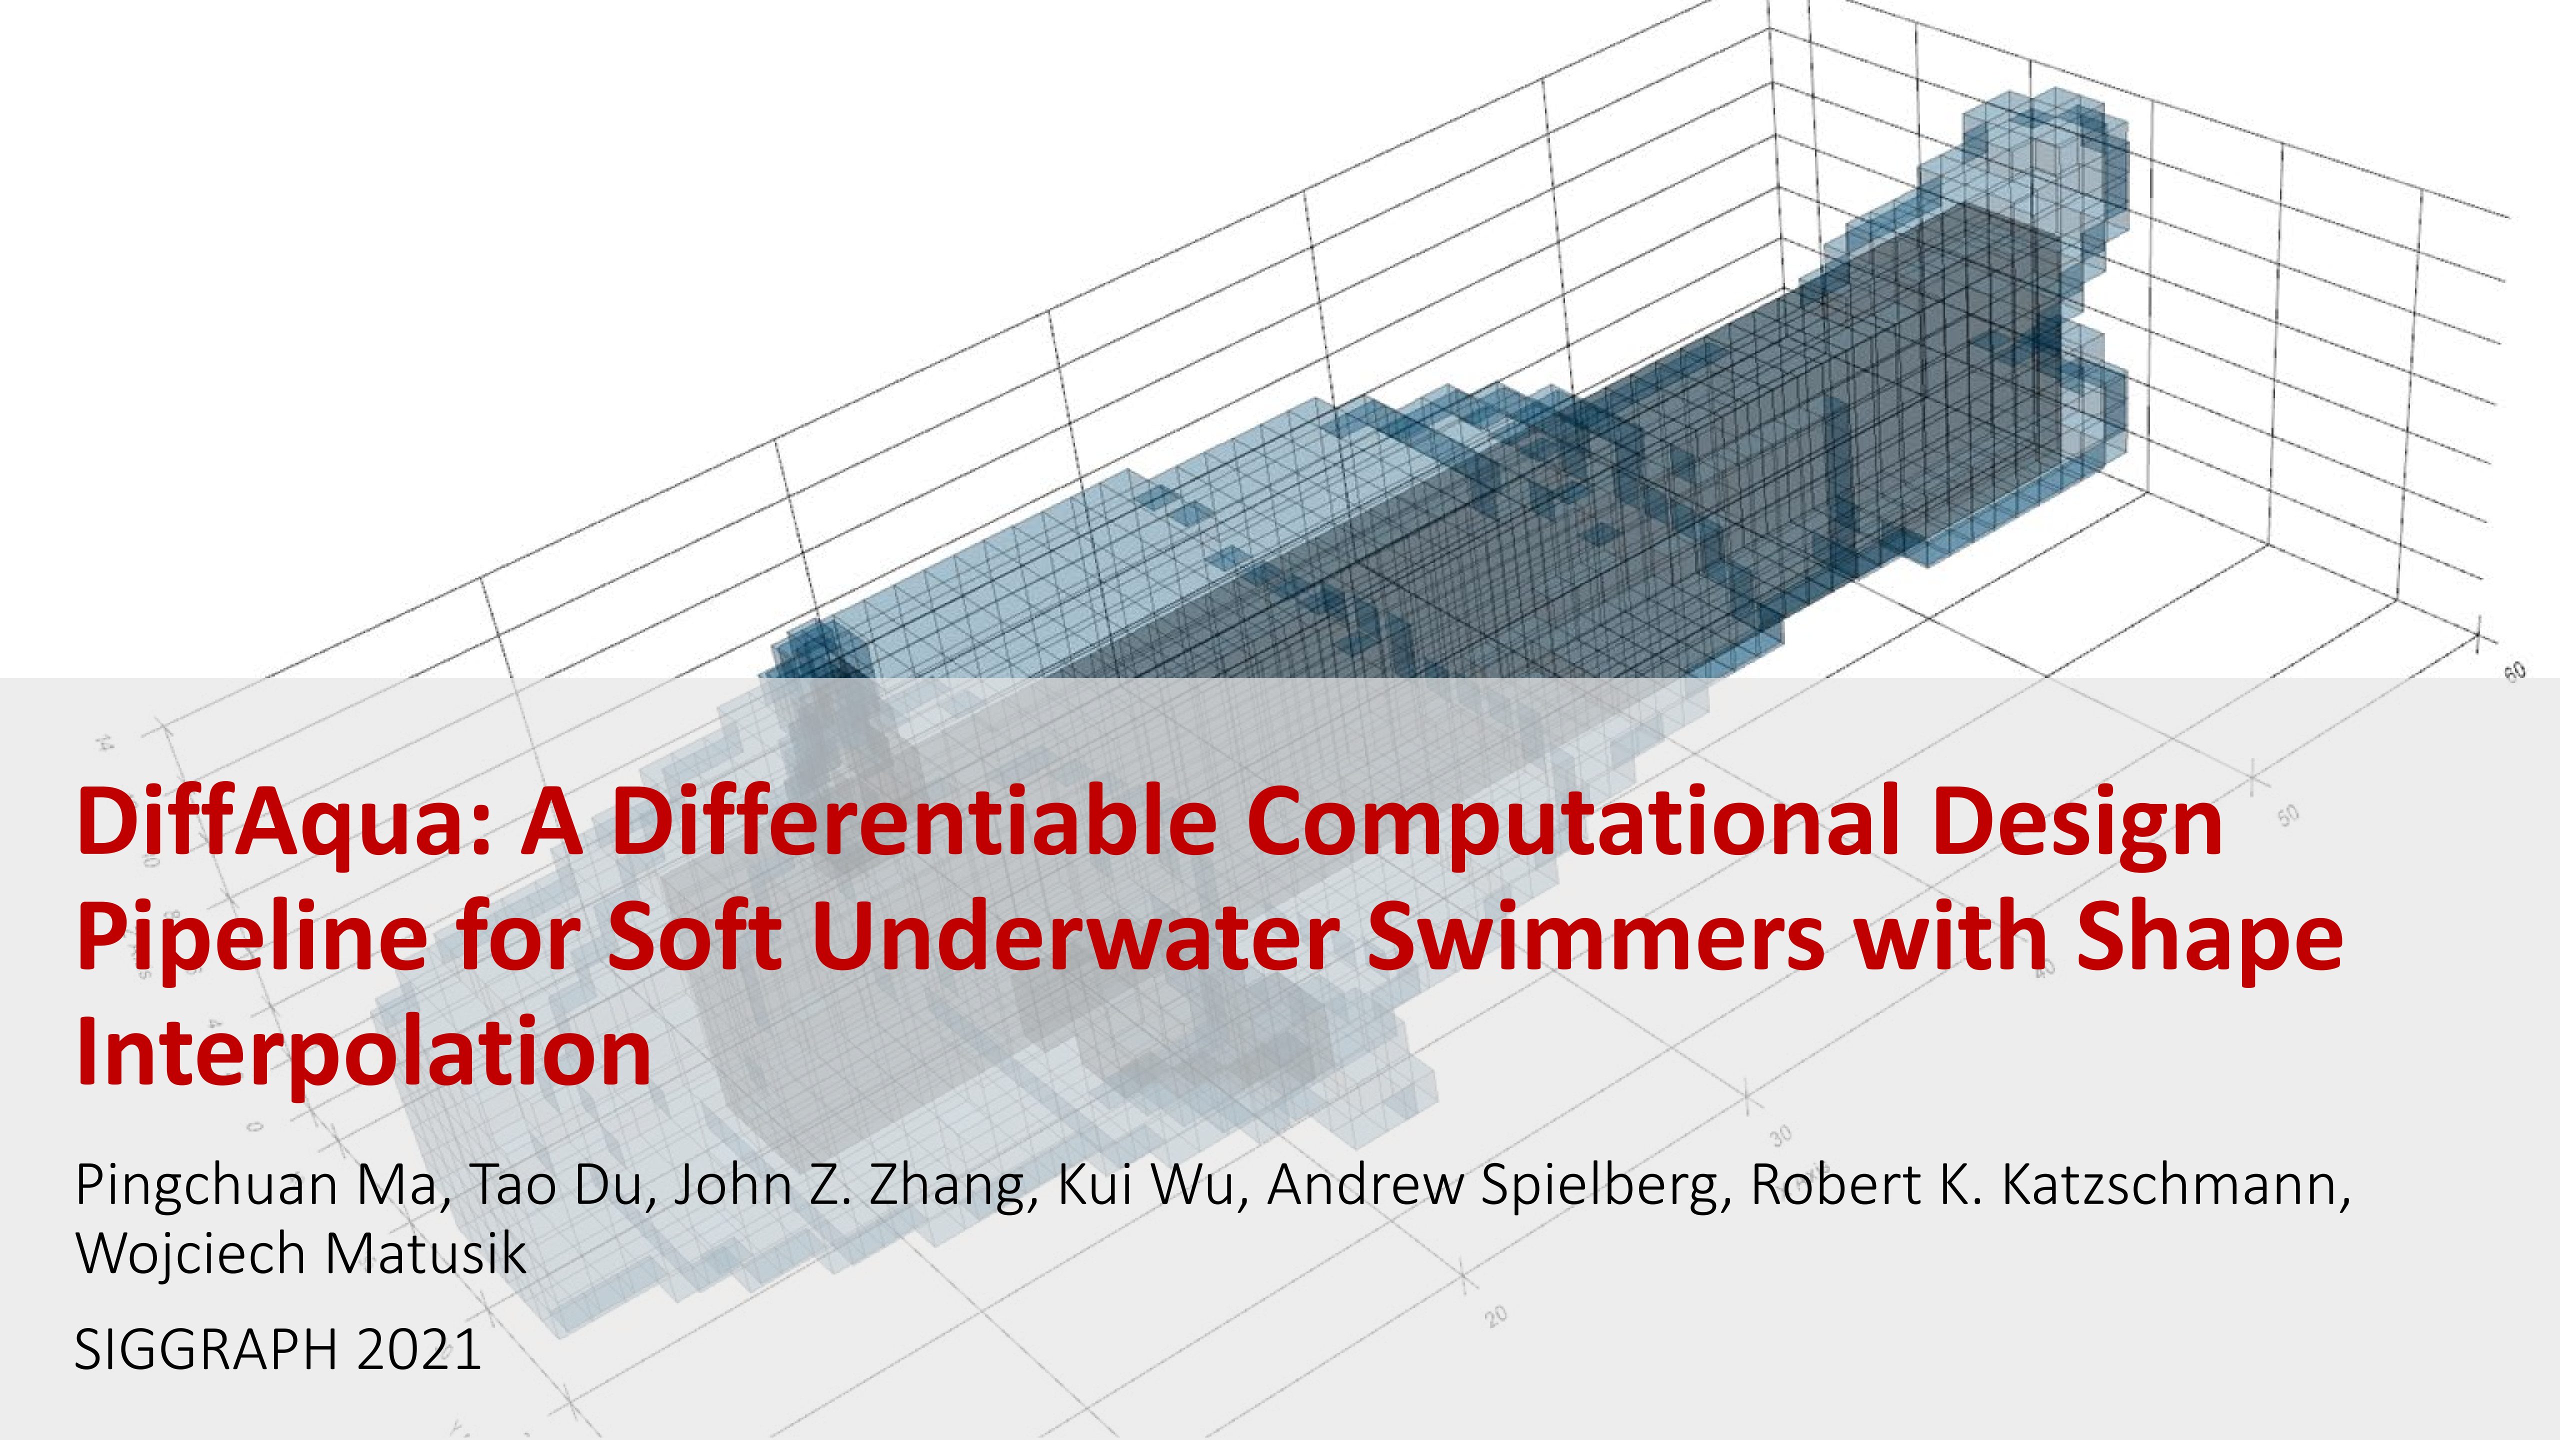

DiffAqua: A Differentiable Computational Design Pipeline for Soft Underwater Swimmers with Shape Interpolation

Pingchuan Ma, Tao Du, John Z. Zhang, Kui Wu, Andrew Spielberg, Robert K. Katzschmann, Wojciech Matusik

SIGGRAPH 2021

Problem statement

Find the optimal *shape* and *control* of soft robotic fishes to achieve *extremal* performance for underwater tasks.

Applications of soft robotic fish

Why is it an inverse dynamics problem

Parametrization

Modeling

Evaluation

Desired performance

e.g. flow-resistance

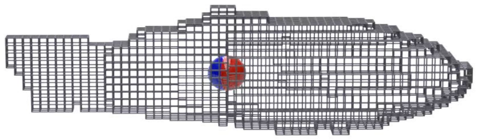

Why is it an inverse dynamics problem

Parametrization	Modeling	Evaluation
	Known dynamic model	Desired performance
	Continuum mechanics	e.g. flow-resistance
	$\nabla \cdot \sigma + f = 0$	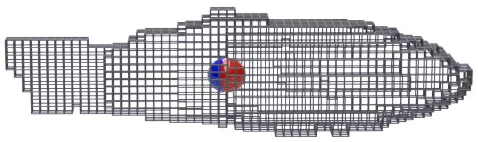

Why is it an inverse dynamics problem

Parametrization	Modeling	Evaluation
Shape and control to be determined	Known dynamic model	Desired performance
	Continuum mechanics	e.g. flow-resistance

$$\nabla \cdot \sigma + f = 0$$

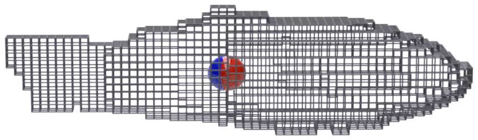

The challenges

Fishes are **soft**: many degrees of freedom are needed.

Fishes are **diverse**: it’s difficult to find one compact representation for all.

Parametrization is the key

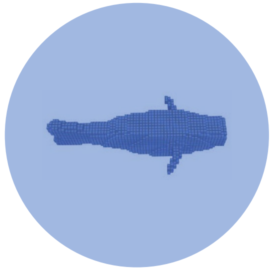

40k DoFs
3 fins

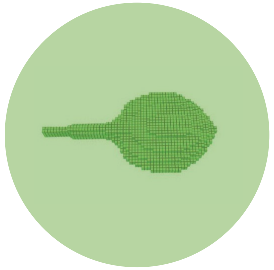

40k DoFs
1 fin

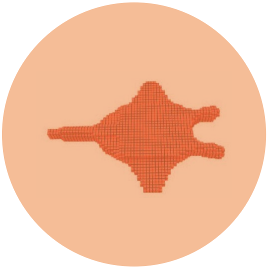

40k DoFs
5 fin

Our approach: Wasserstein gradients

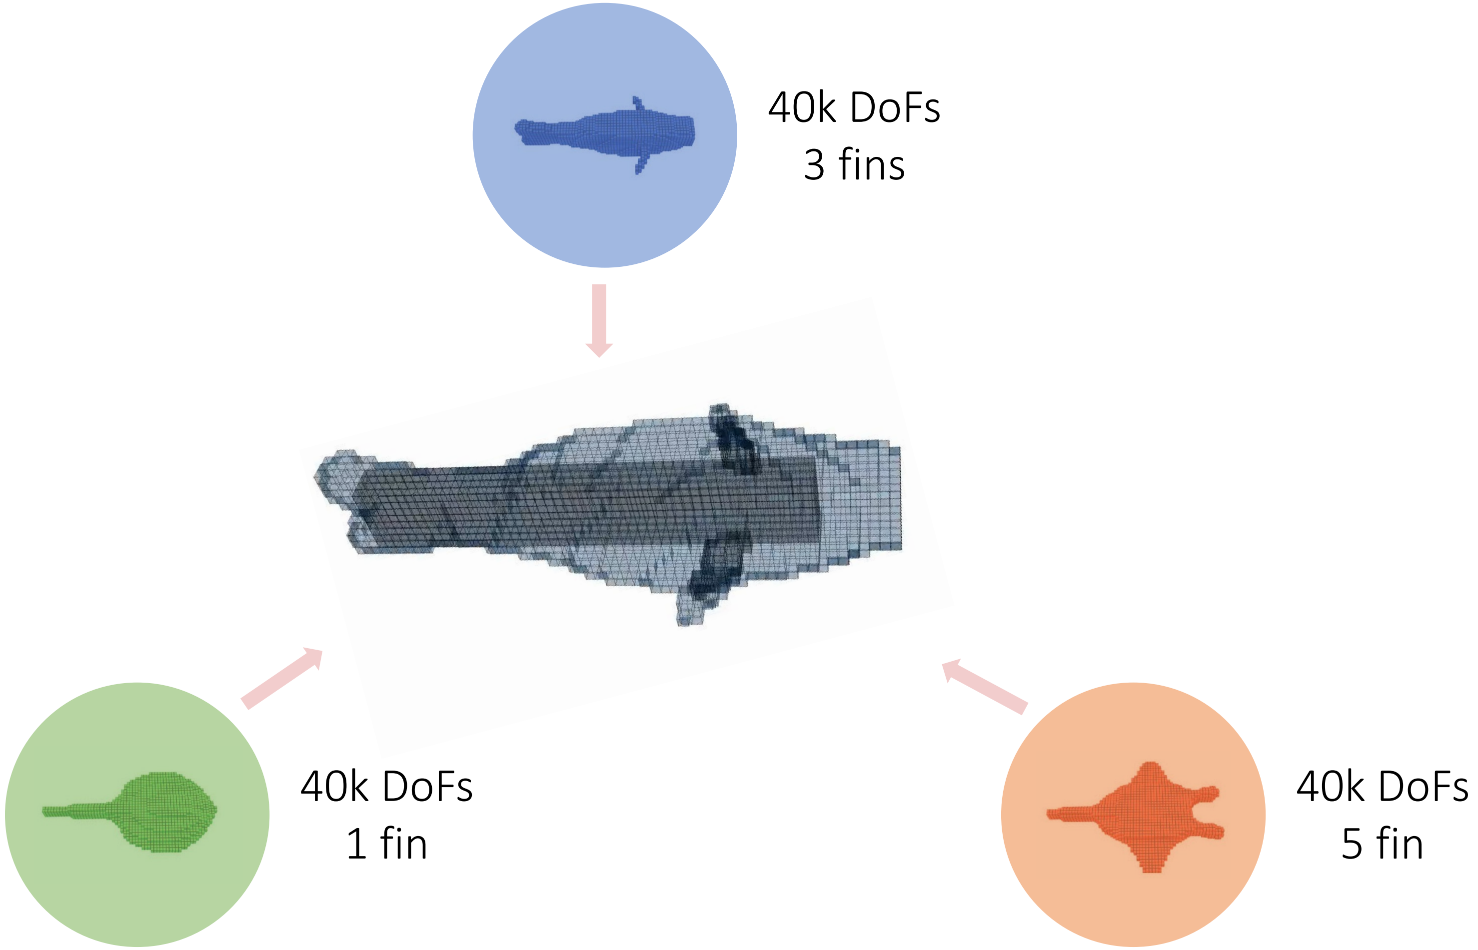

Our approach: Wasserstein gradients

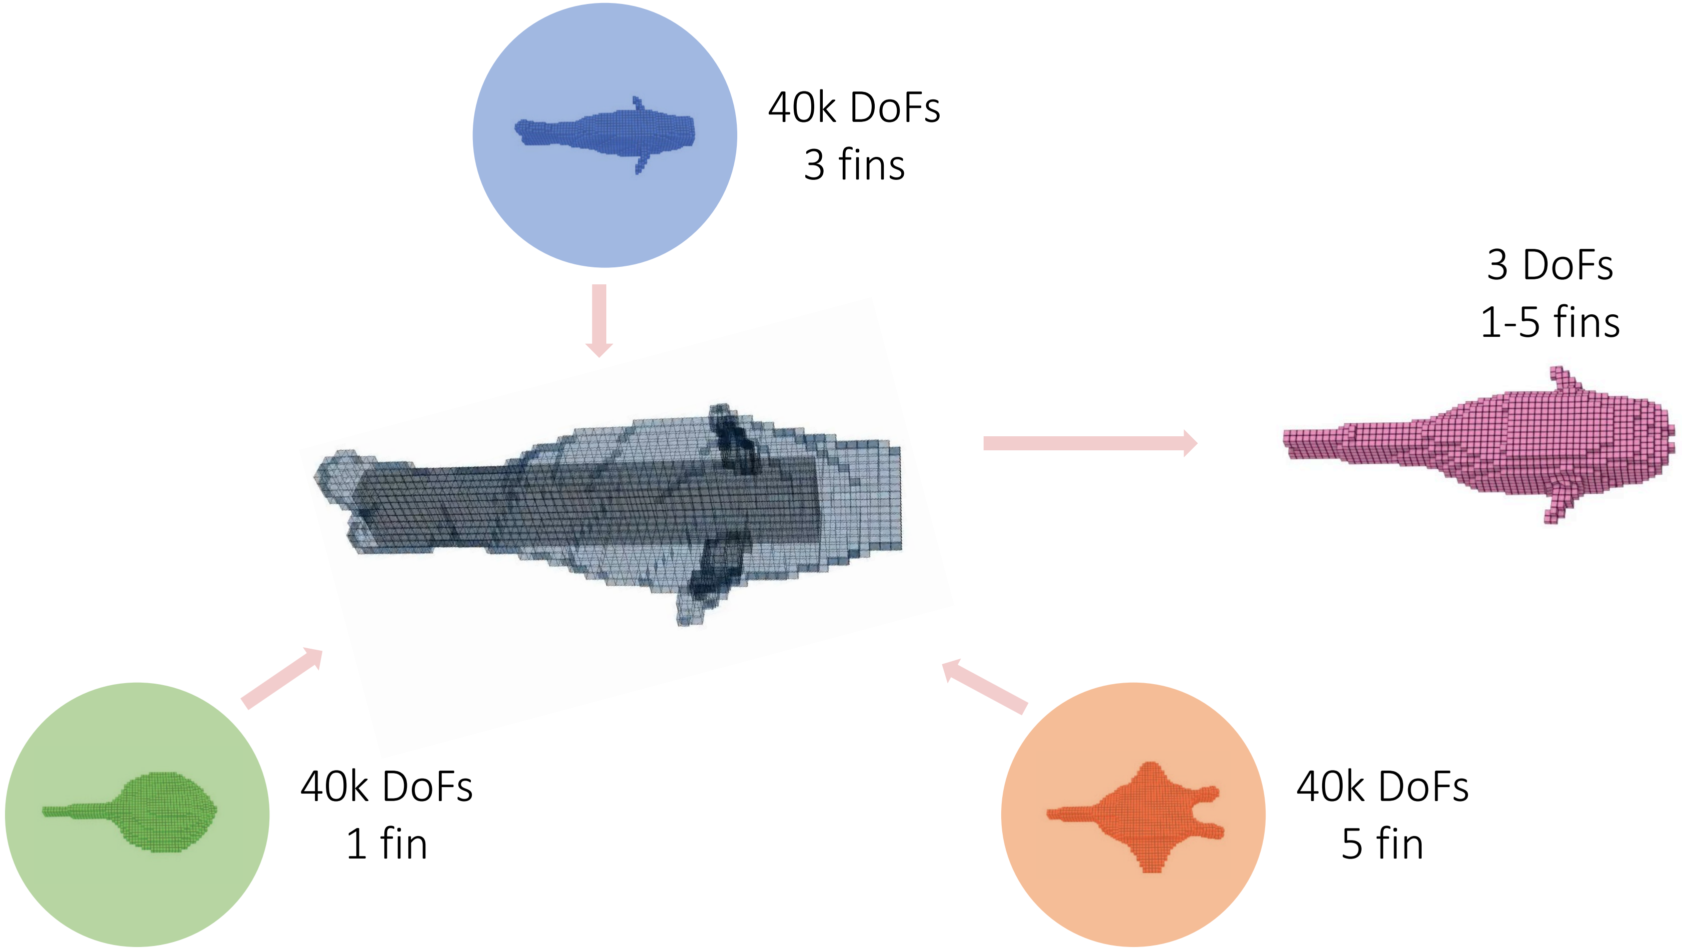

Our approach: Wasserstein gradients

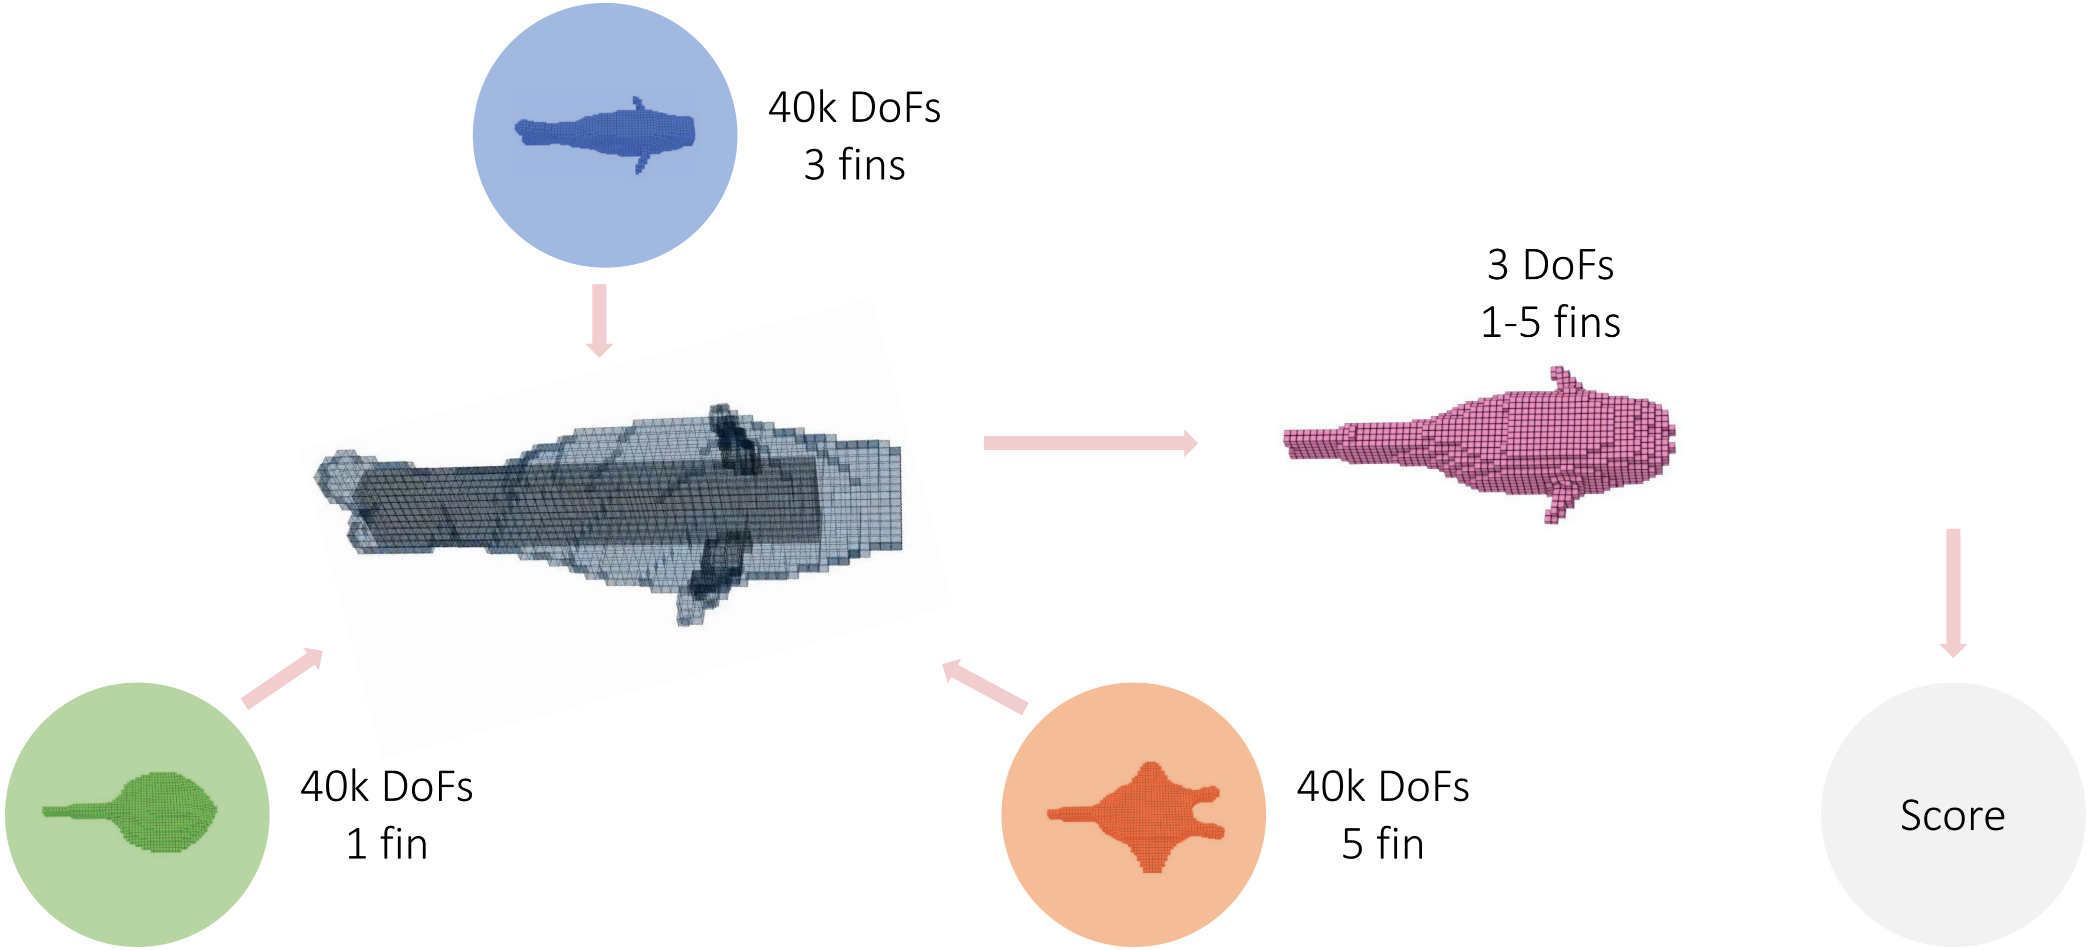

Our approach: Wasserstein gradients

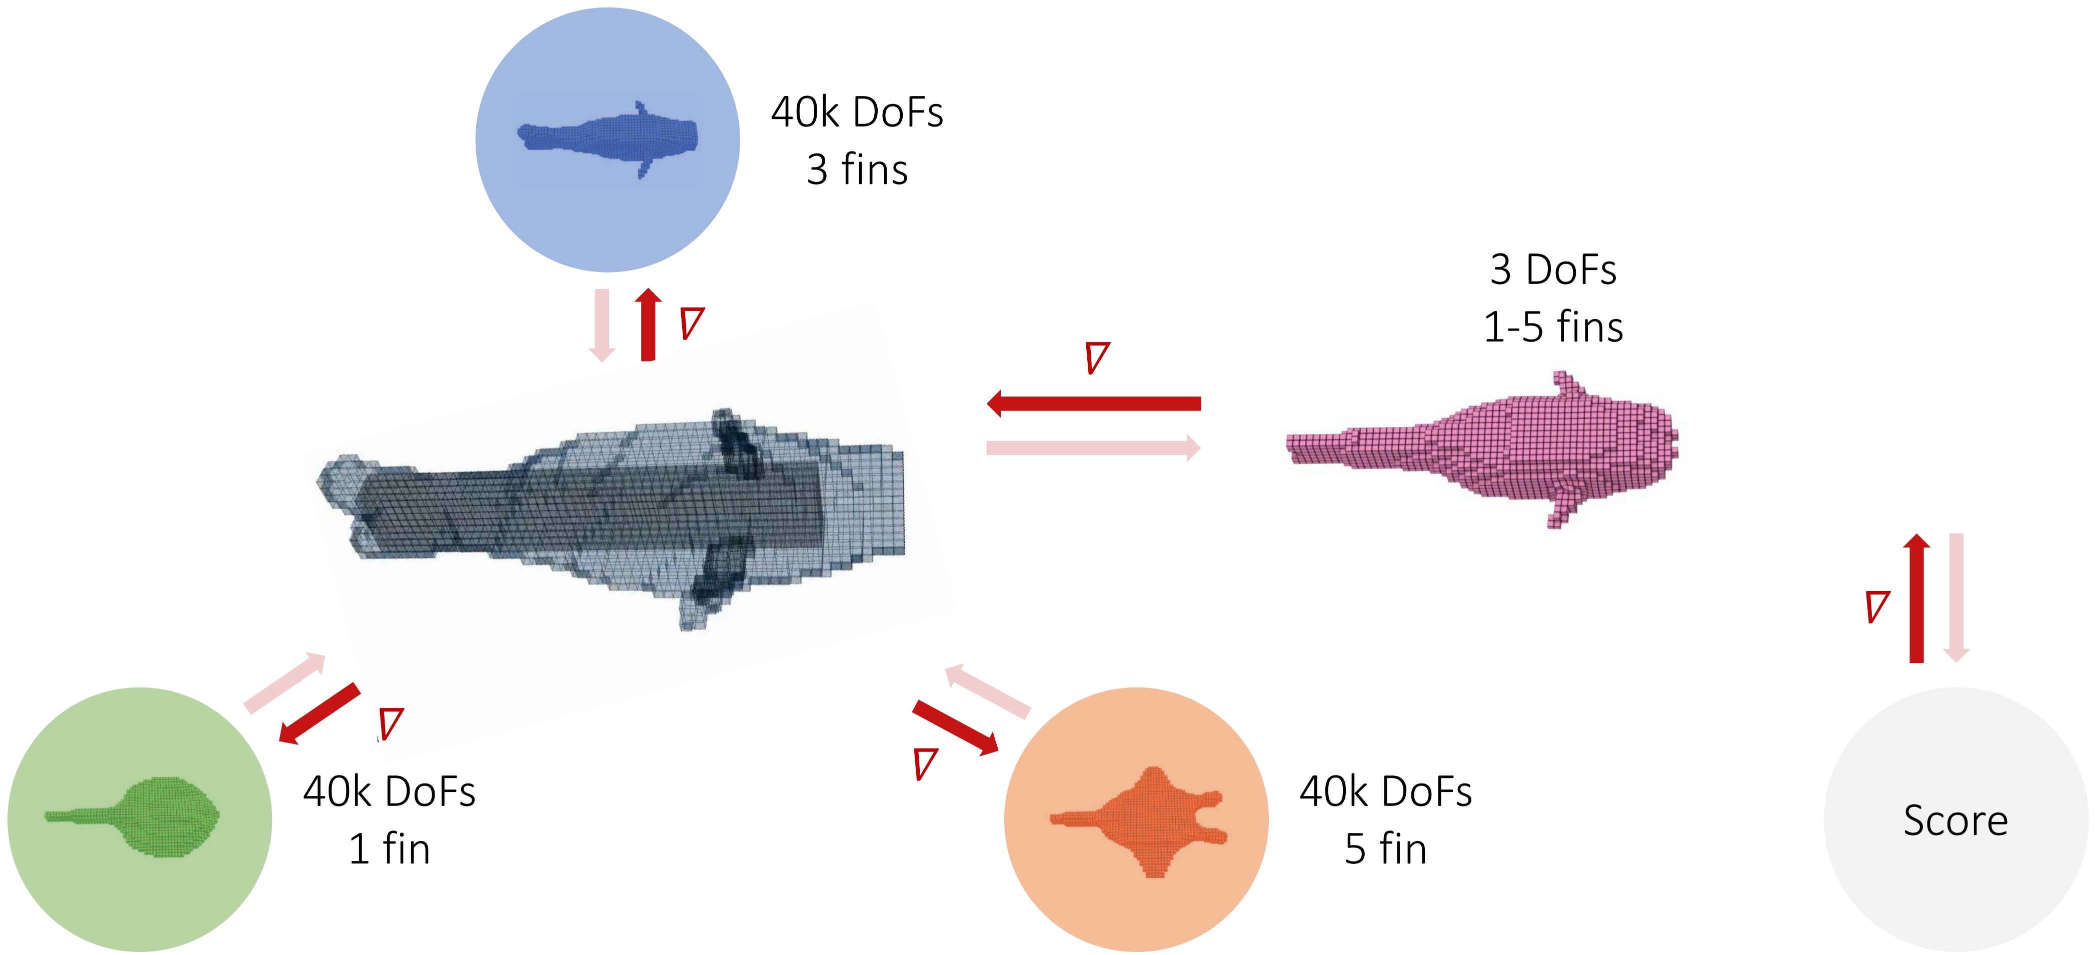

Example: speedy fish

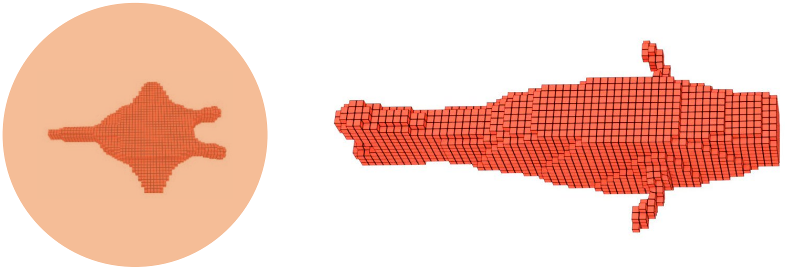

Optimized control only

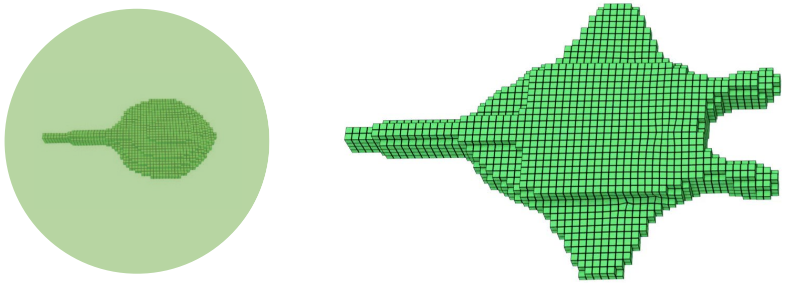

Optimized control only

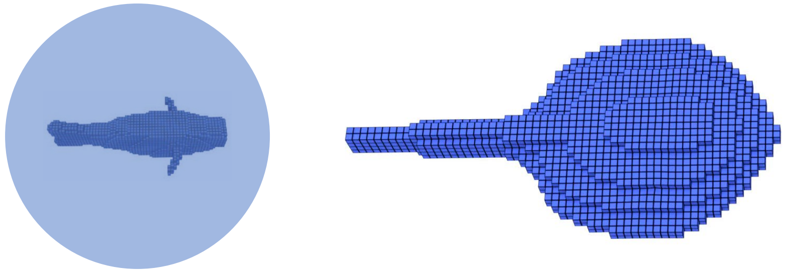

Optimized control only

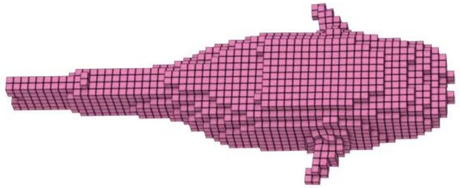

Optimized shape and control

Example: flow-resistant fish

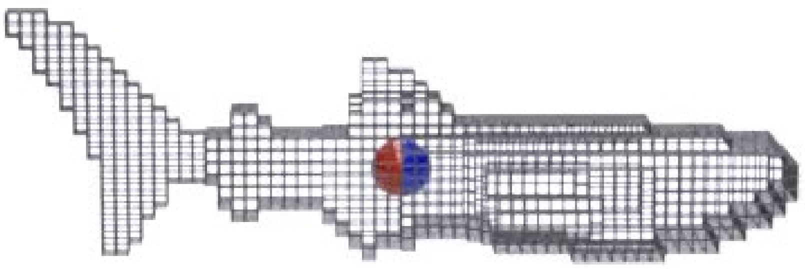

Unoptimized

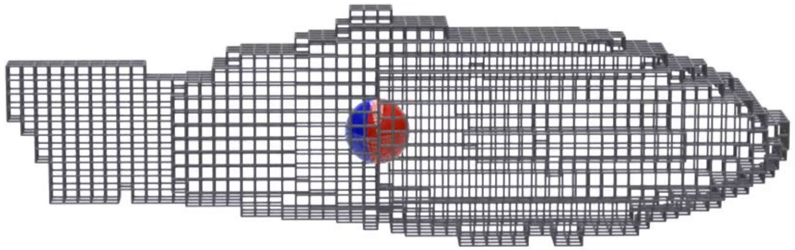

Optimized

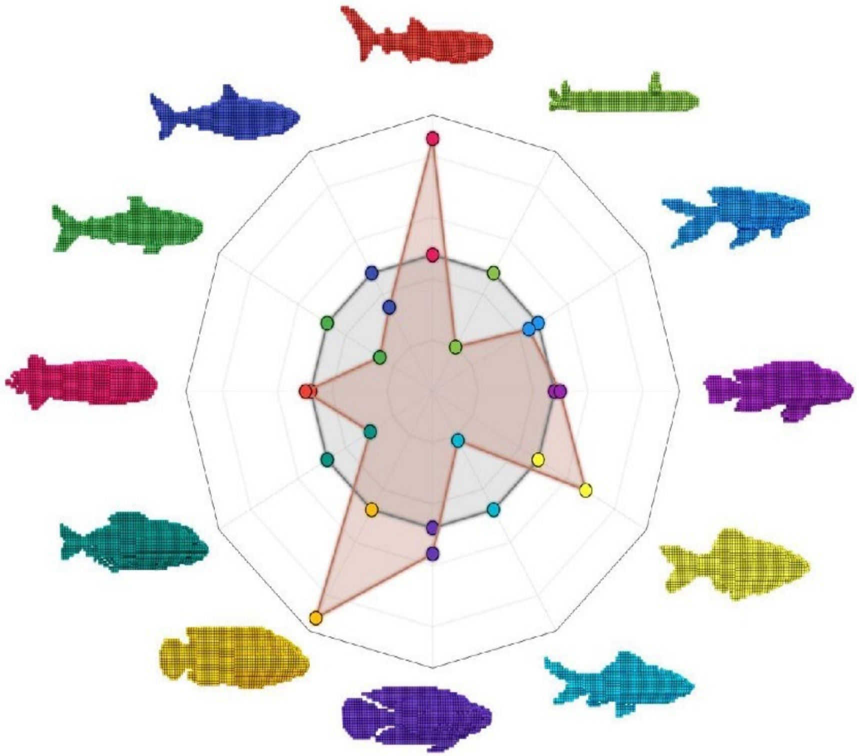

Wasserstein weights

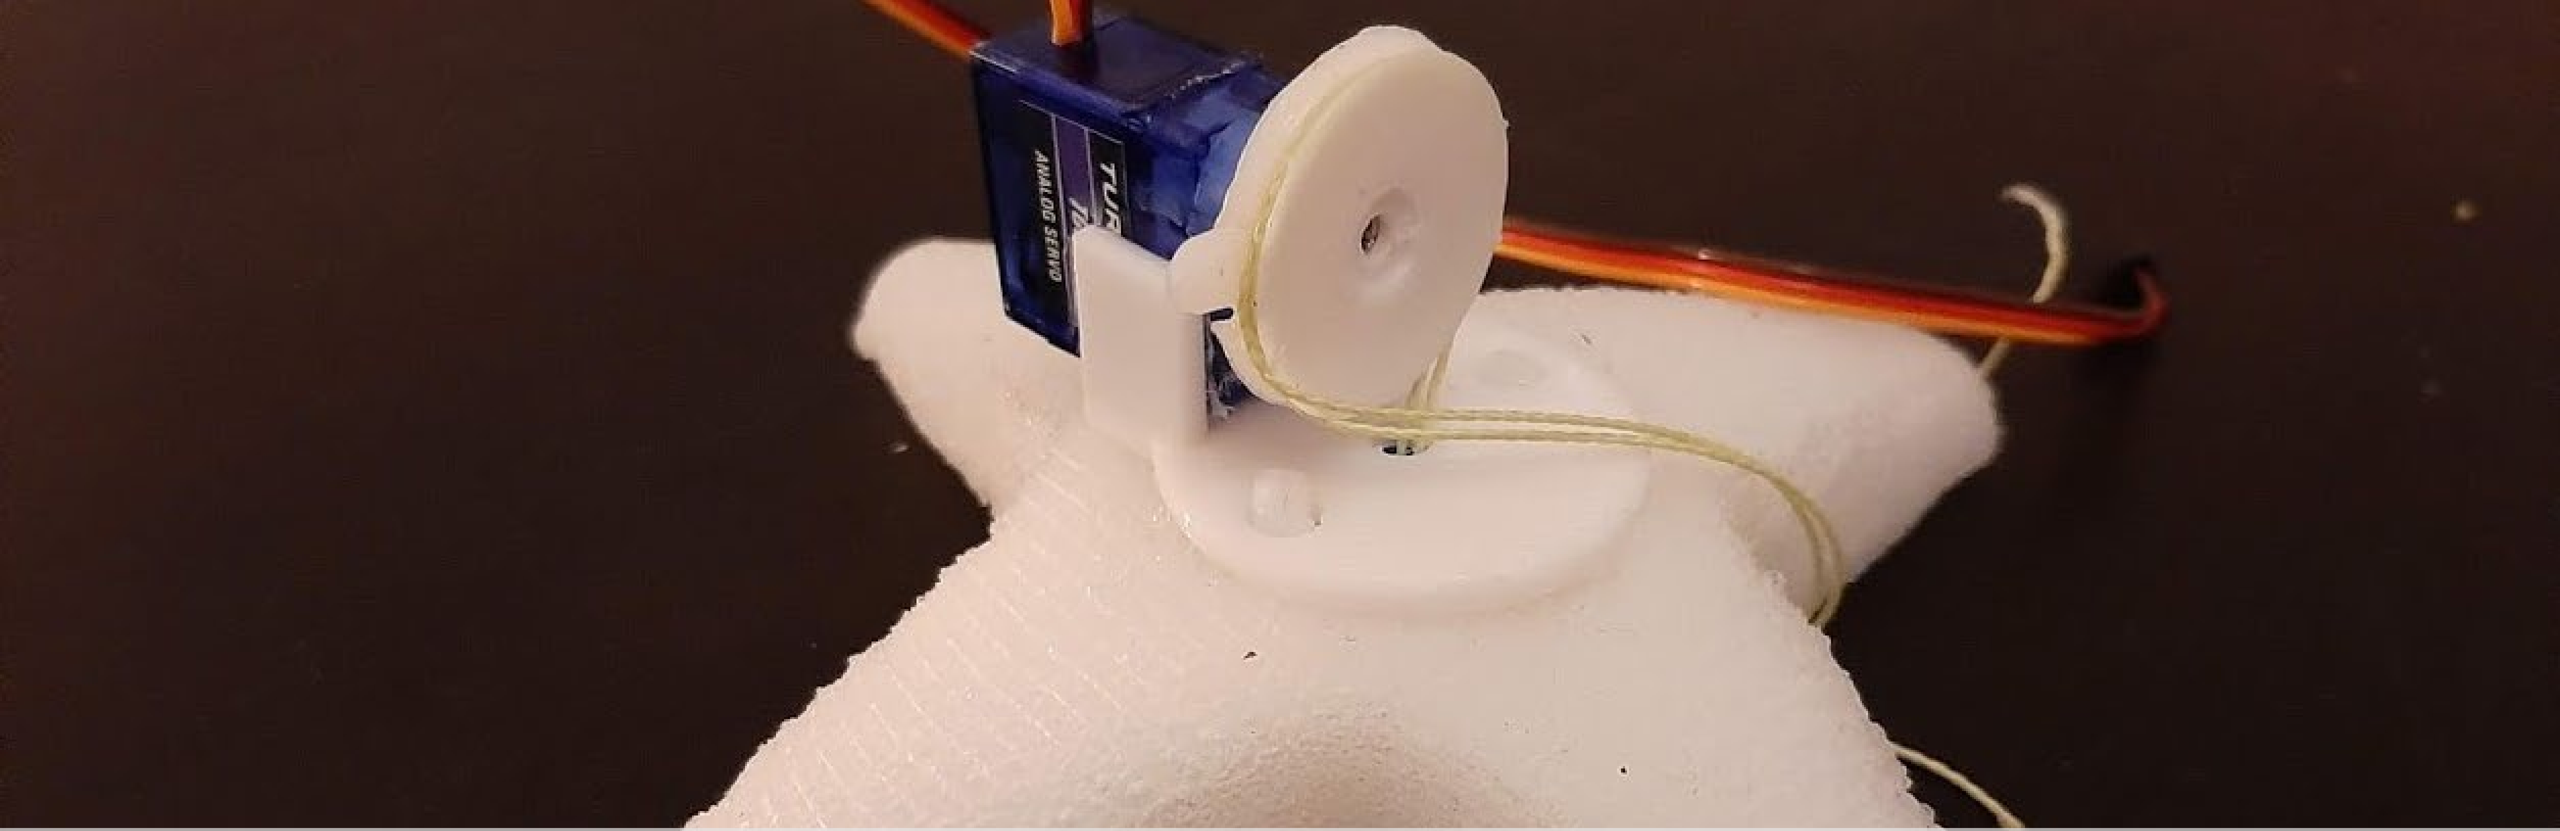

Underwater Soft Robot Modeling and Control with Differentiable Simulation

Tao Du*, Josie Hughes*, Sebastien Wah, Wojciech Matusik, Daniela Rus

IEEE RA-L/RoboSoft 2021

Summary

∇ Parametrization

Initializing parameters

∇ Modeling

Deriving governing equations

∇ Evaluation

Computing performance metrics

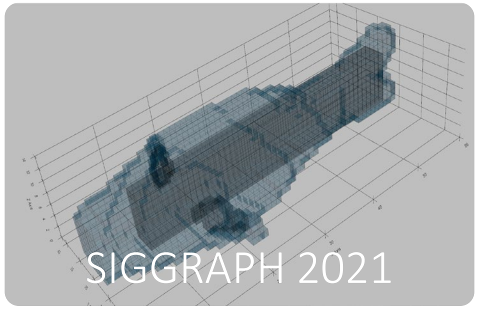

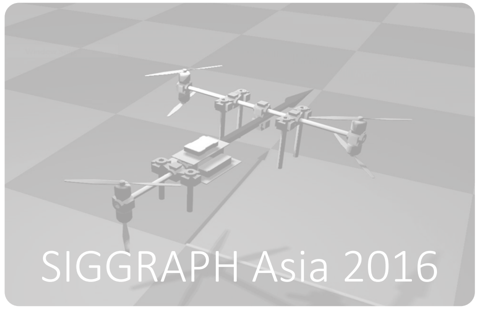

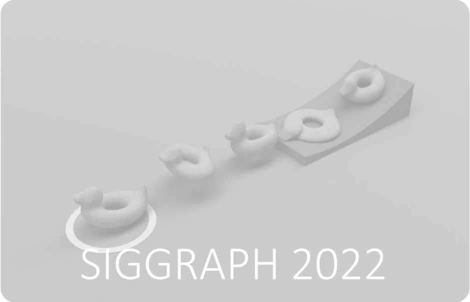

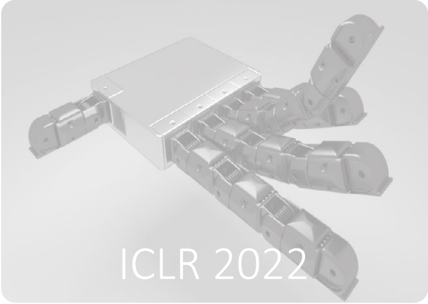

Summary

∇ Parametrization

Initializing parameters

∇ Modeling

Deriving governing equations

∇ Evaluation

Computing performance metrics

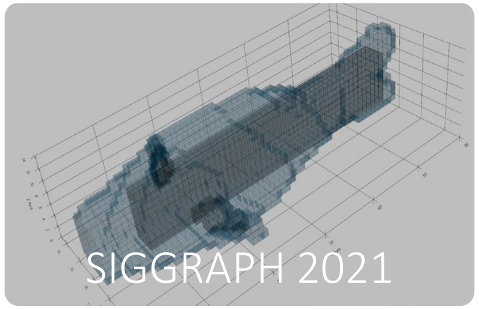

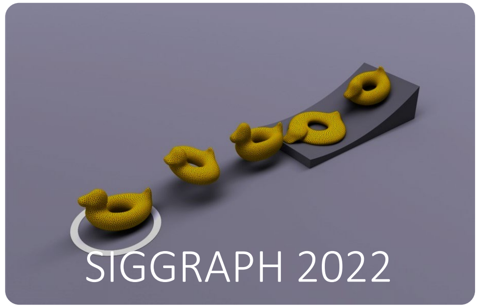

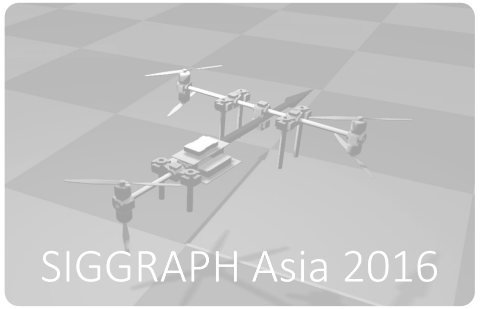

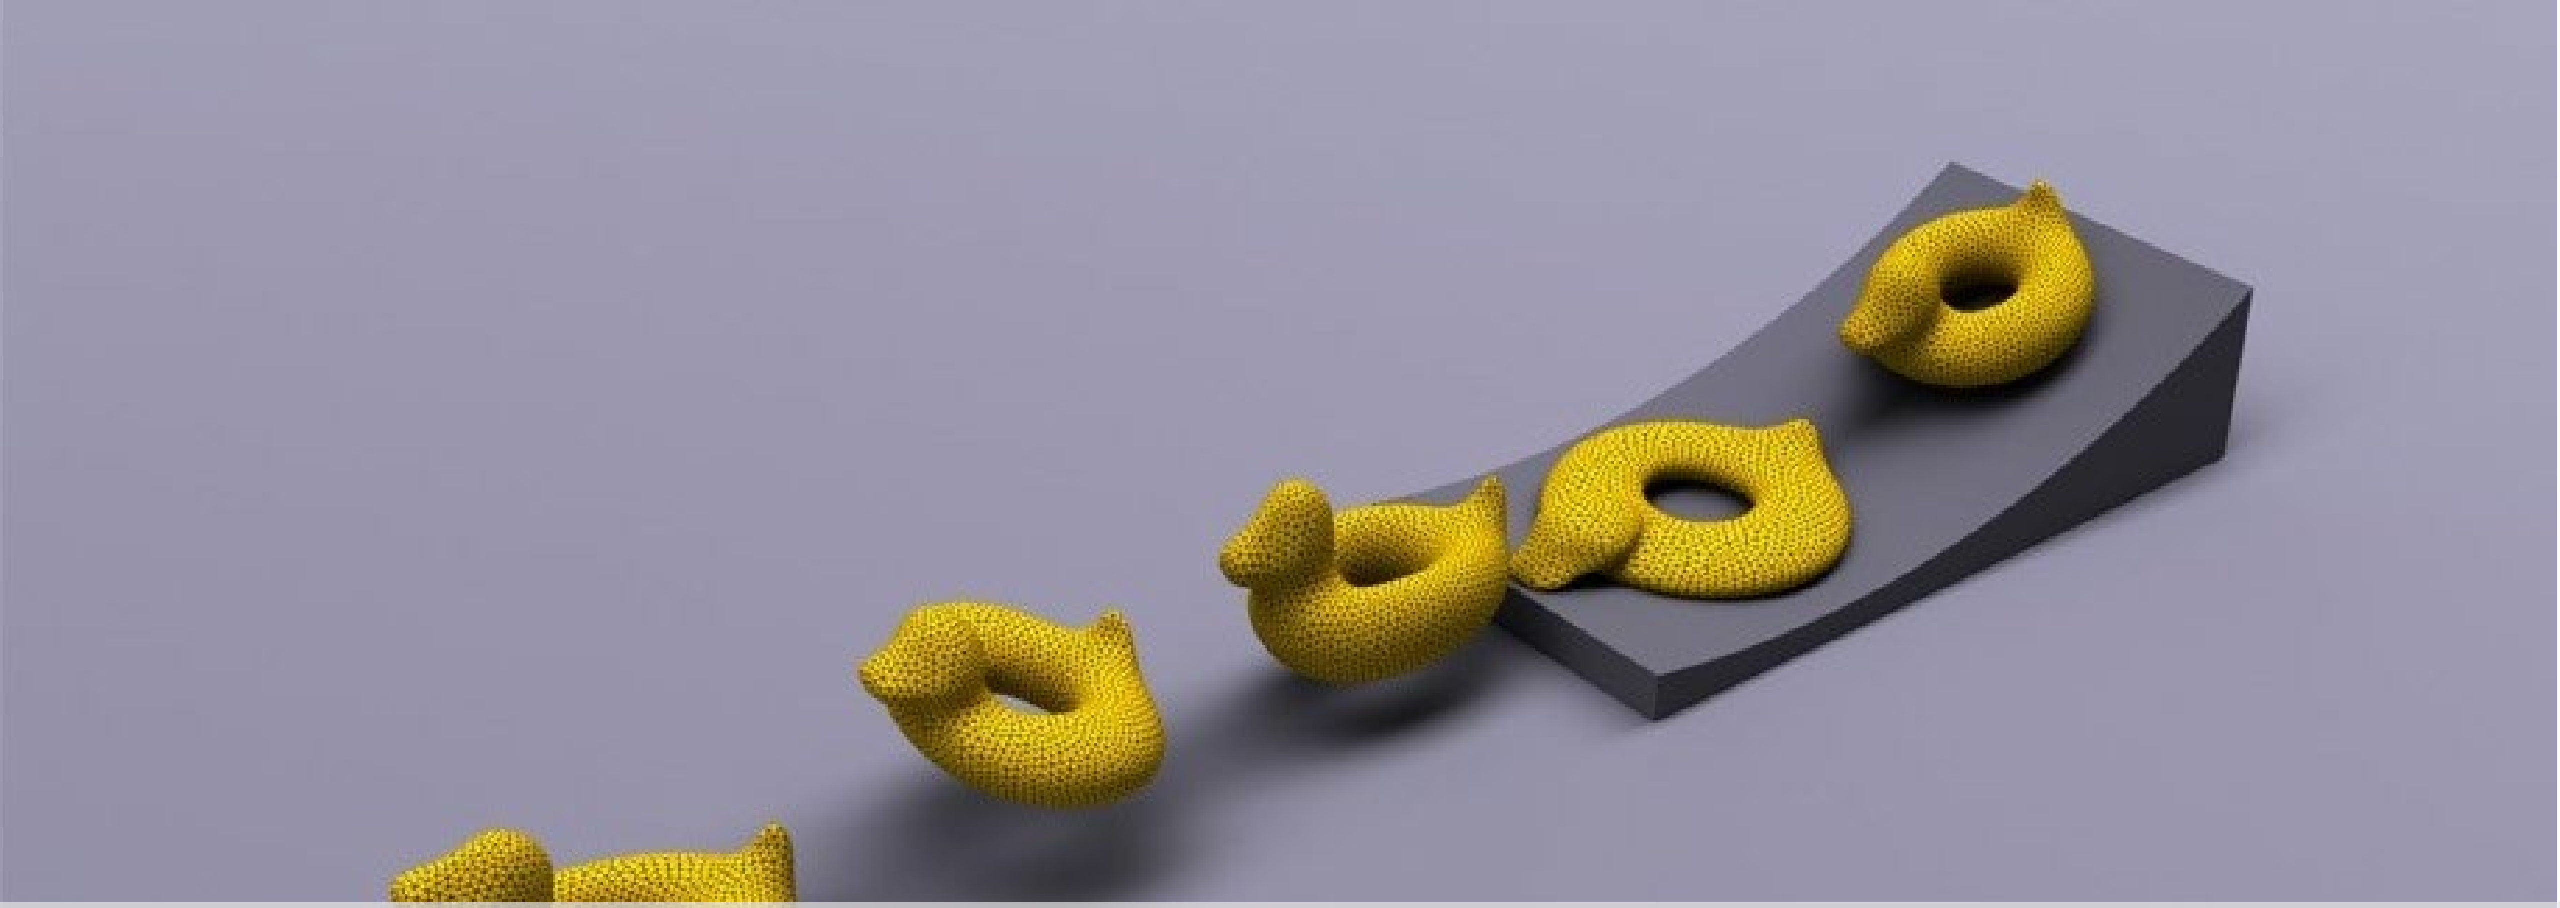

DiffPD: Differentiable Projective Dynamics

Tao Du, Kui Wu, Pingchuan Ma, Sebastien Wah, Andrew Spielberg, Daniela Rus, Wojciech Matusik

ACM Transactions on Graphics (SIGGRAPH 2022)

Feature highlights

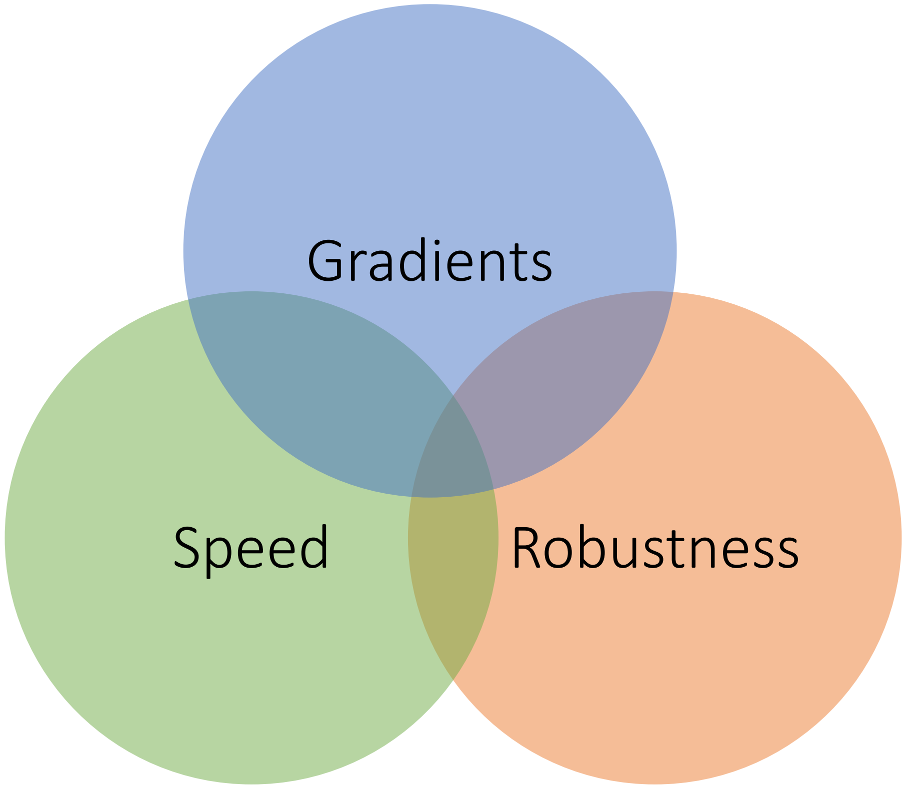

Feature highlights

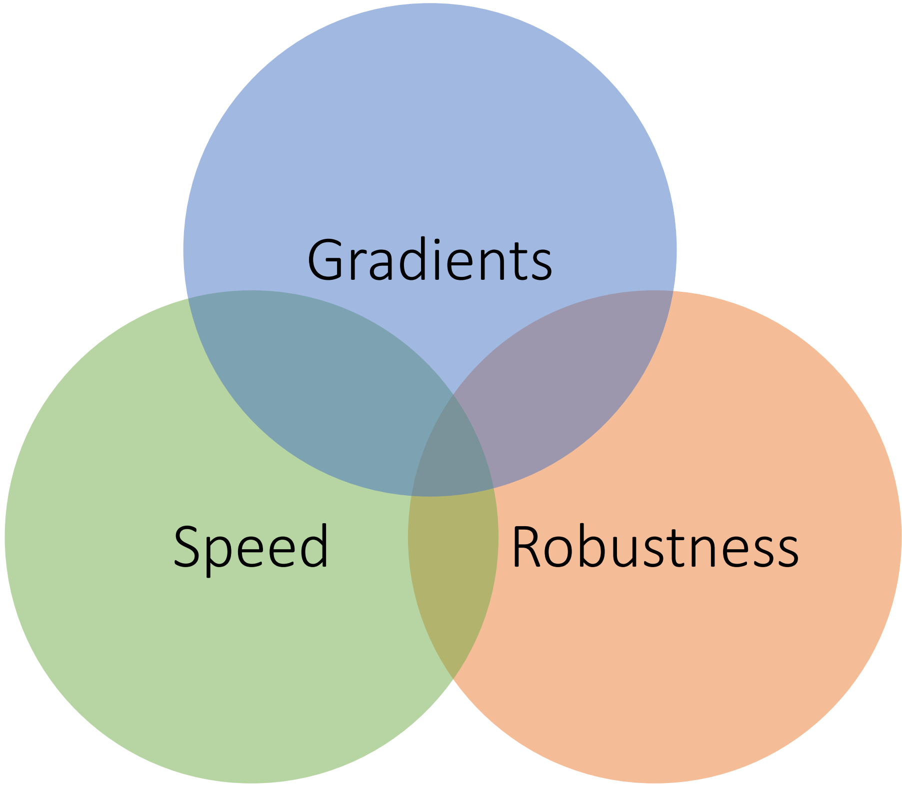

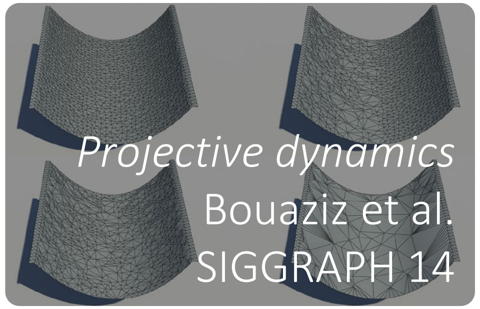

Applications: system identification

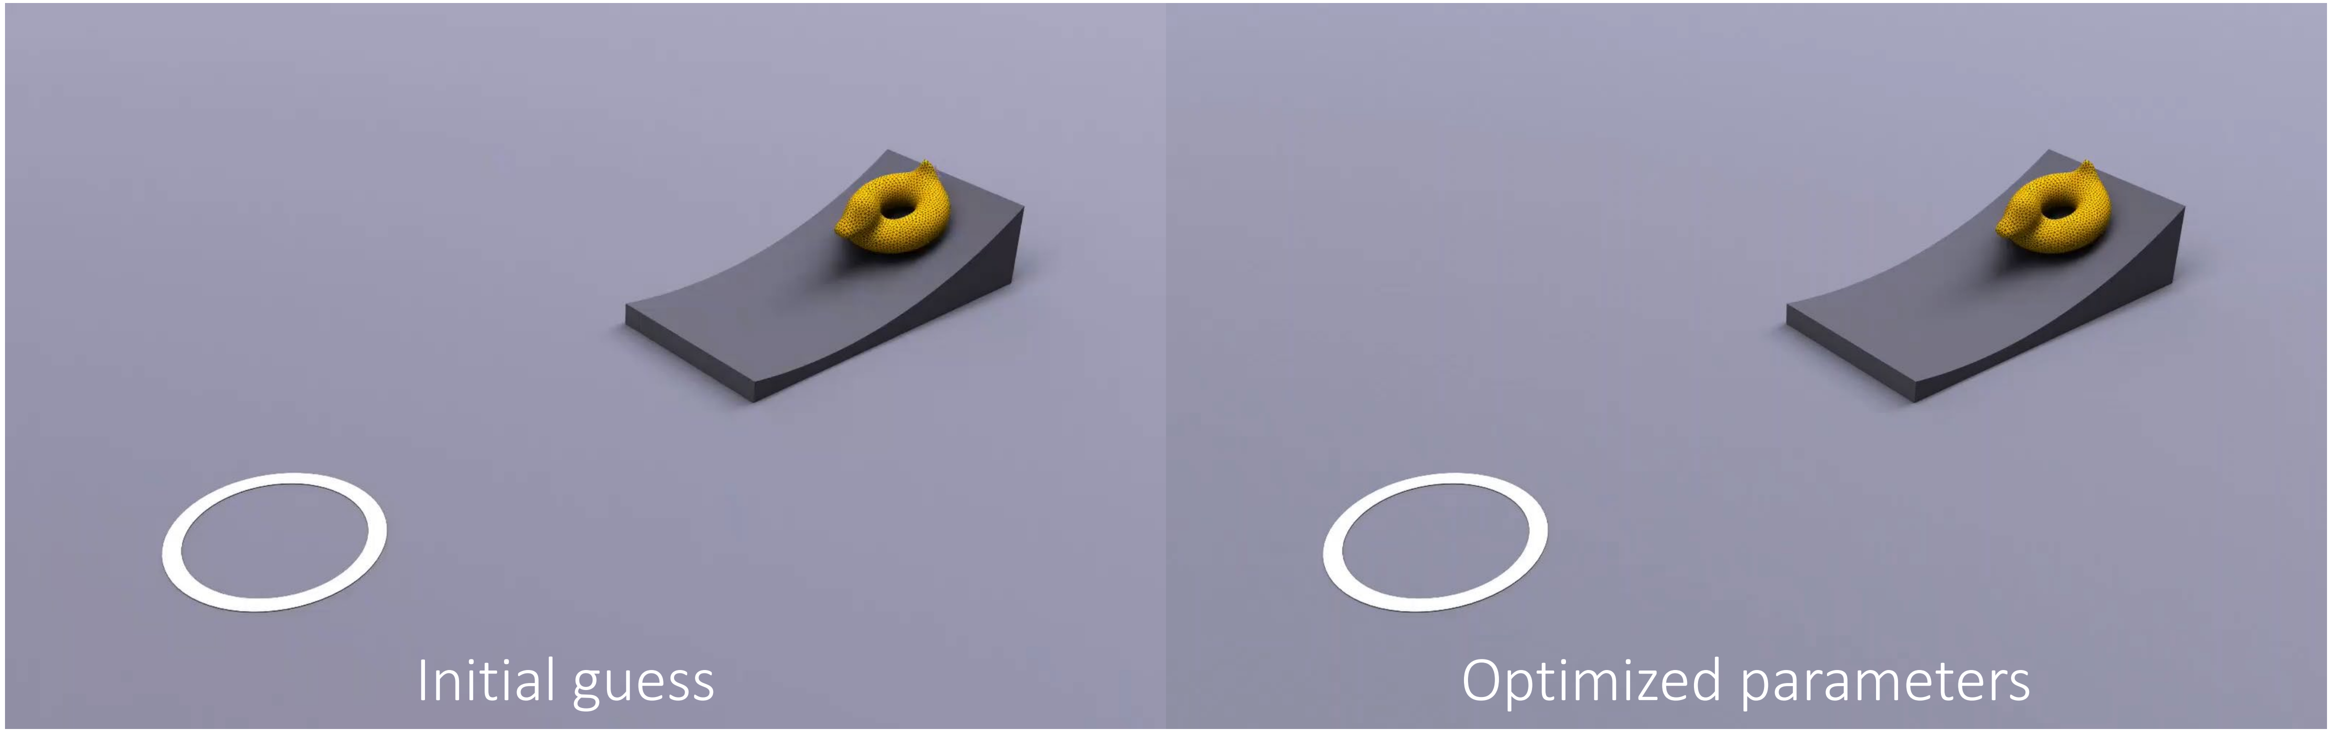

Applications: trajectory optimization

Initial guess

Optimized parameters

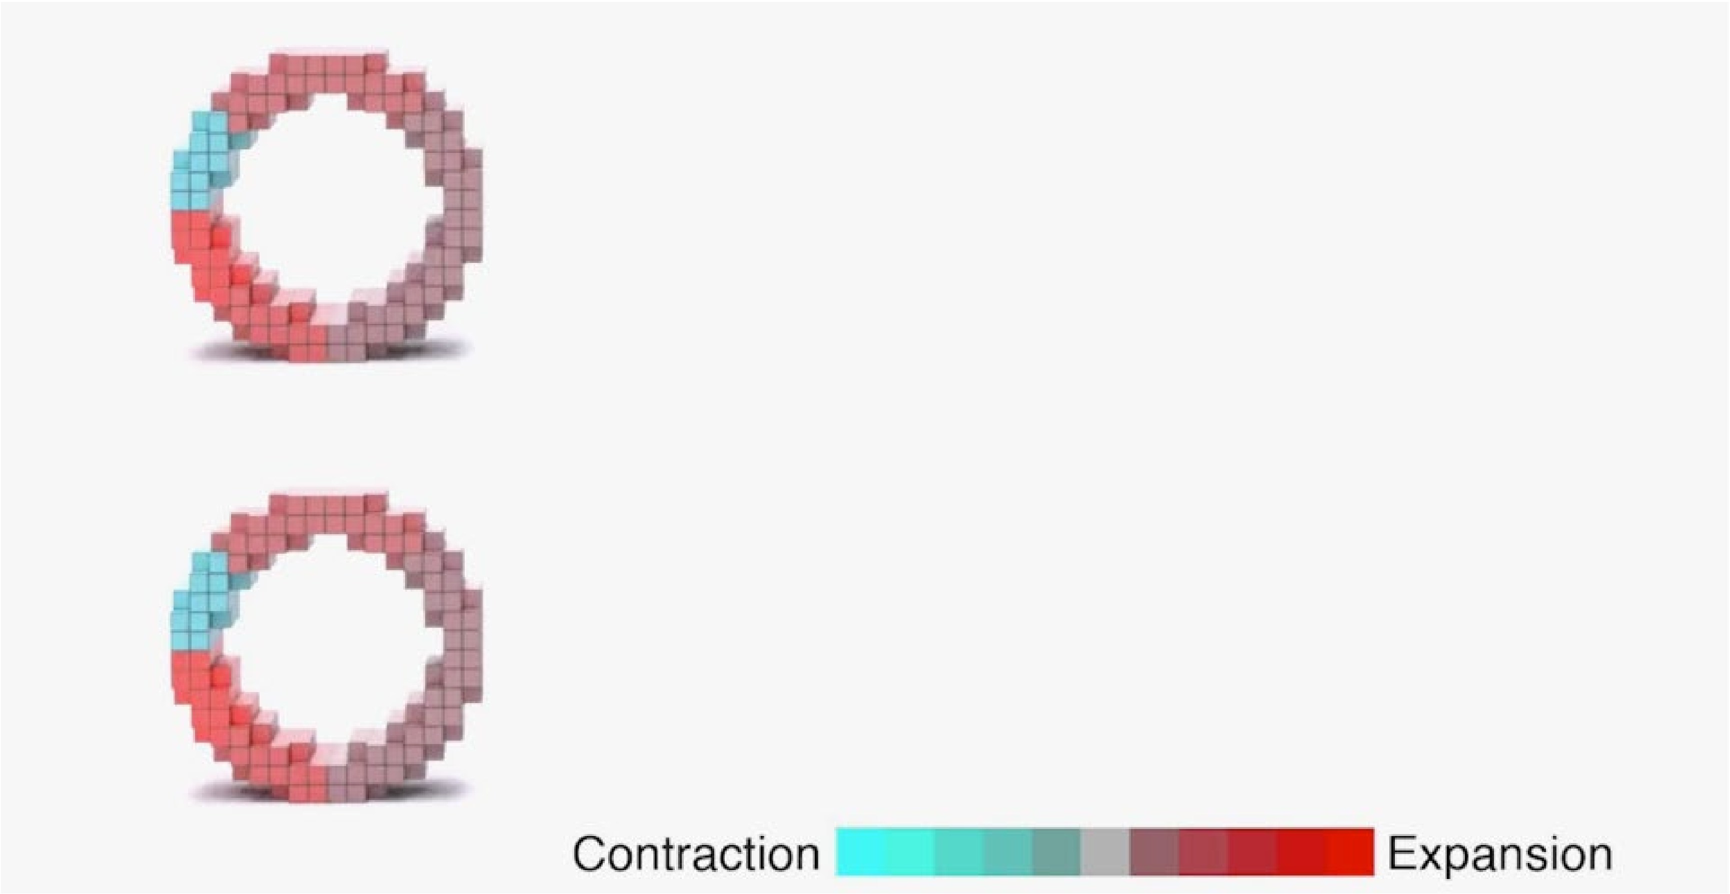

Applications: real-to-sim transfer

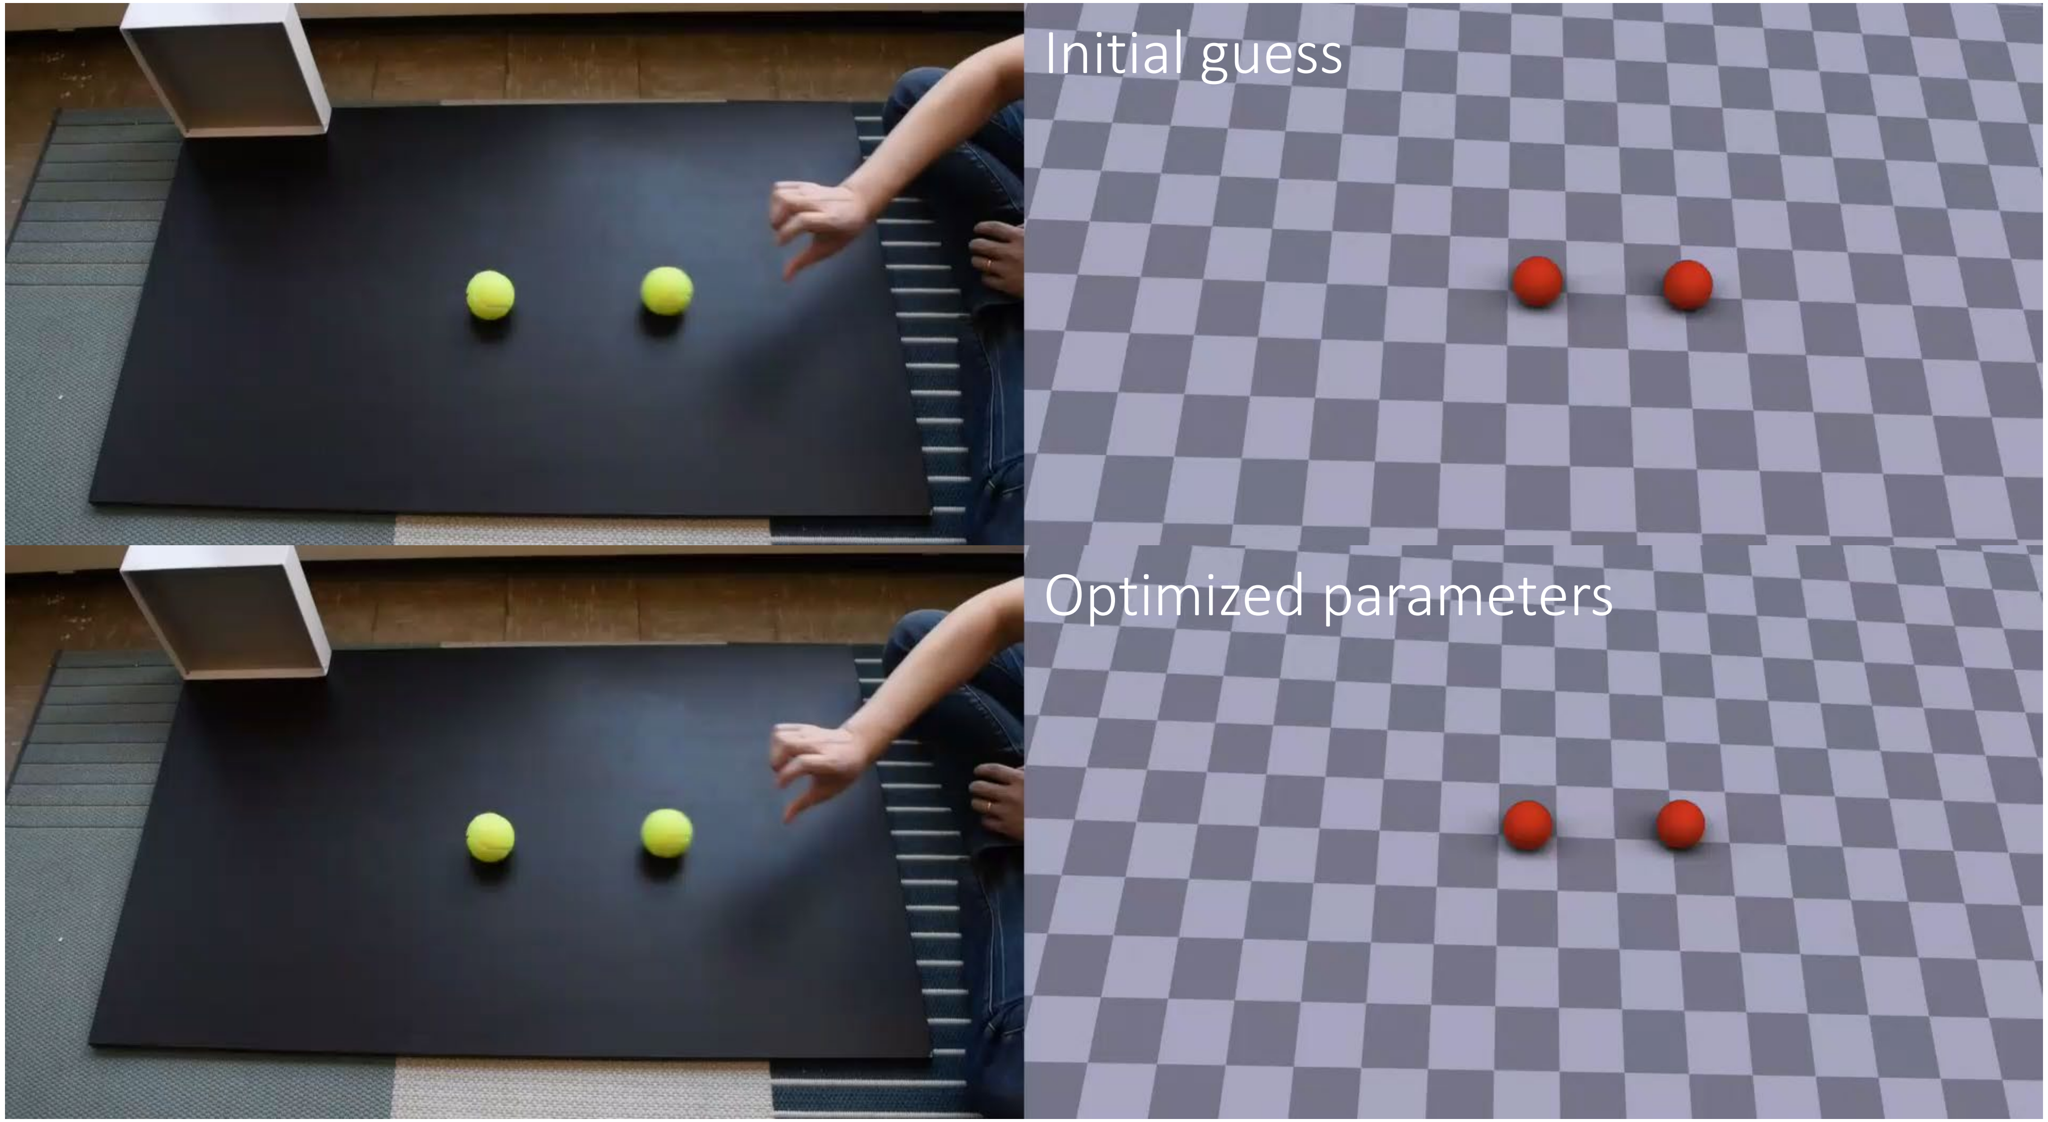

Summary

∇ Parametrization

Initializing parameters

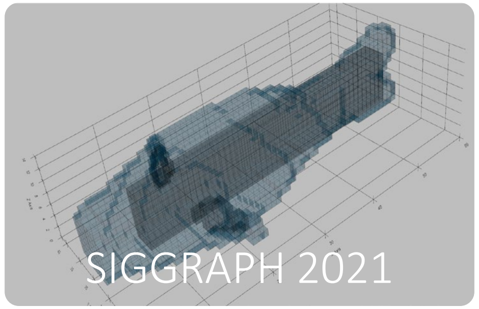

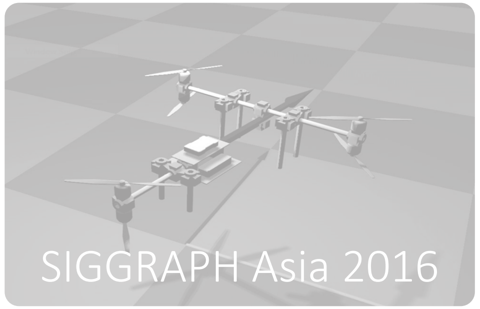

∇ Modeling

Deriving governing equations

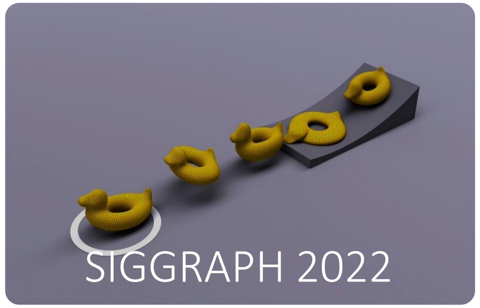

∇ Evaluation

Computing performance metrics

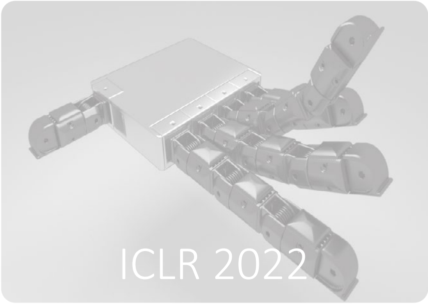

Summary

∇ Parametrization

Initializing parameters

∇ Modeling

Deriving governing equations

∇ Evaluation

Computing performance metrics

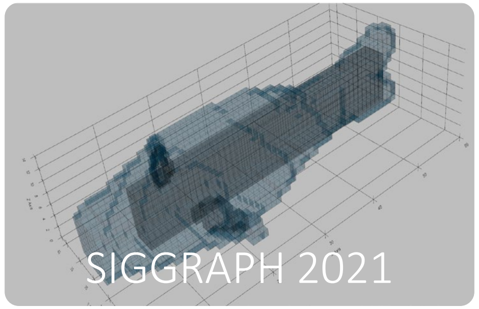

SIGGRAPH 2021

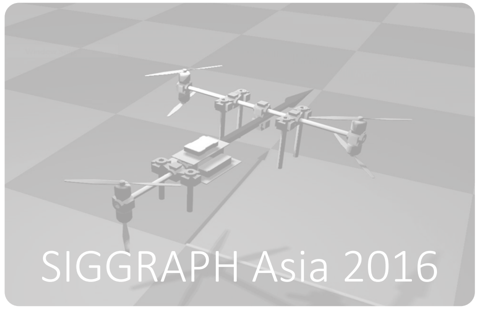

SIGGRAPH Asia 2016

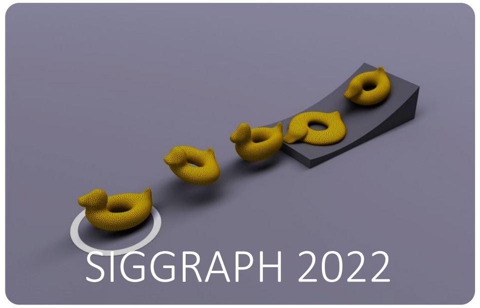

SIGGRAPH 2022

SIGGRAPH 2022

ICLR 2022

RA-L 2022

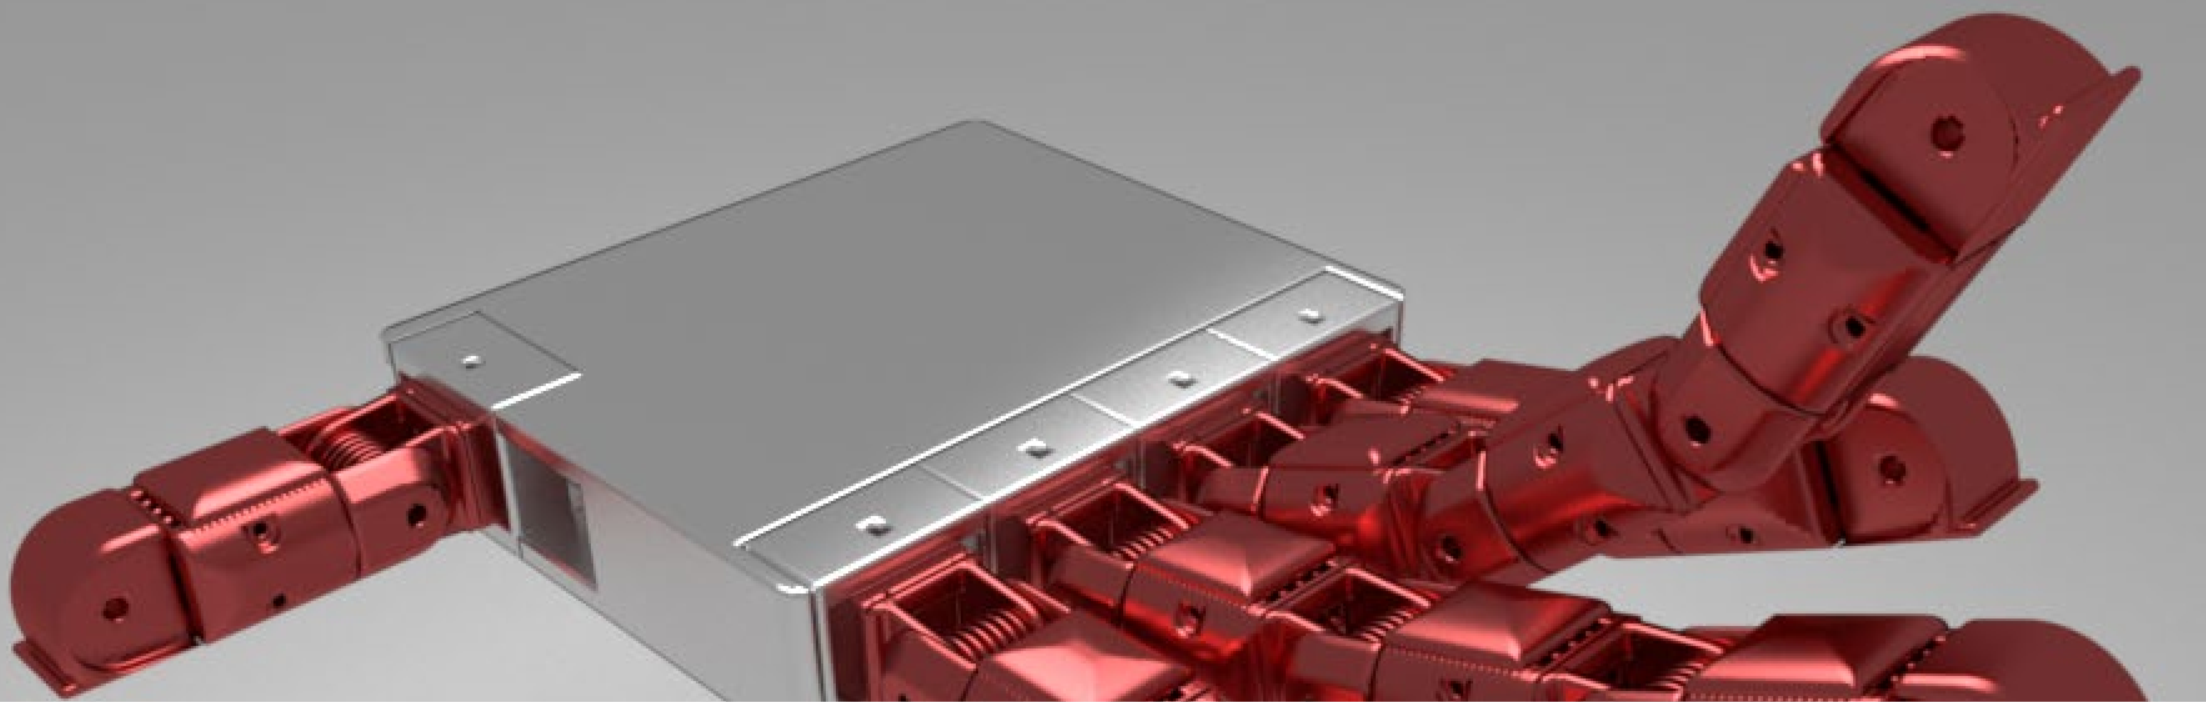

RISP: Rendering-Invariant State Predictor with Differentiable Simulation and Rendering for Cross-Domain Parameter Estimation

Pingchuan Ma*, Tao Du*, Joshua B. Tenenbaum, Wojciech Matusik, Chuang Gan

ICLR 2022 (oral)

Problem statement

Build a digital twin of a robot from its video of motion sequences.

Why is it an inverse dynamics problem

Parametrization

Modeling

Evaluation

Desired performance

Match motions from videos

Why is it an inverse dynamics problem

Parametrization

Modeling

Evaluation

Known dynamic model

Euler-Lagrange dynamics

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}}\right) - \frac{\partial T}{\partial q} - Q = 0$$

Desired performance

Match motions from videos

Why is it an inverse dynamics problem

Parametrization	Modeling	Evaluation
Control sequence to be determined	Known dynamic model	Desired performance
	Euler-Lagrange dynamics	Match motions from videos

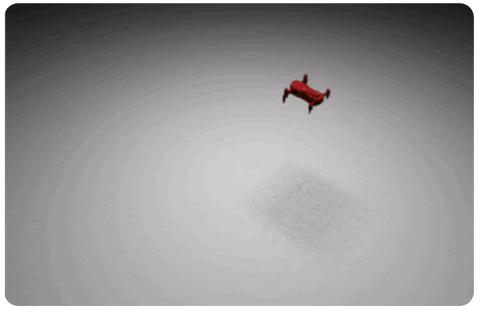

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}}\right) - \frac{\partial T}{\partial q} - Q = 0$$

The challenge

The unknown visual appearance parameters shadows the dynamics information.

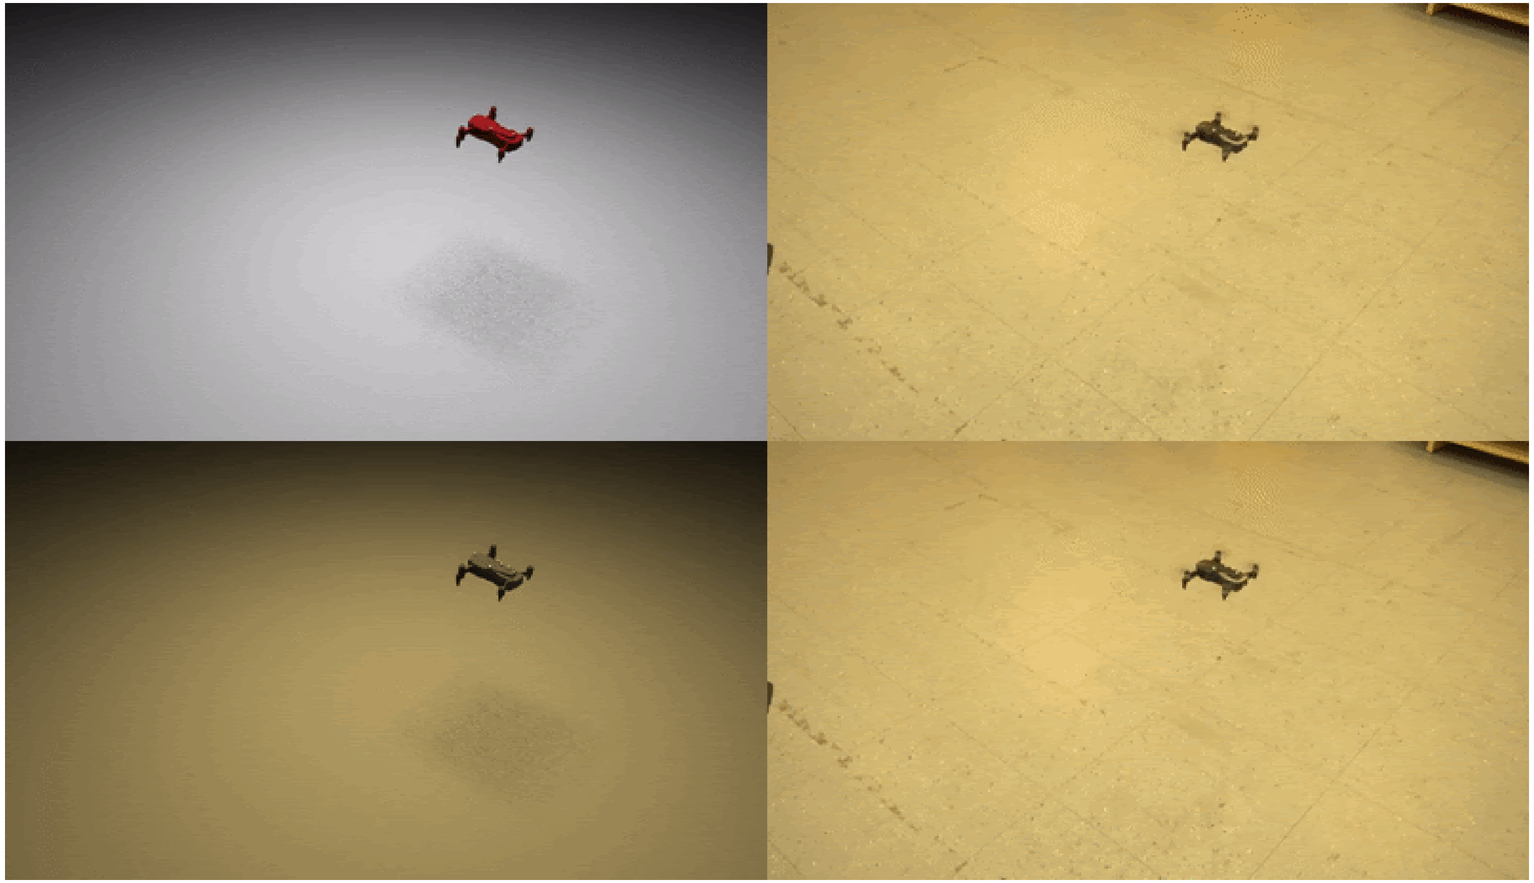

The first idea: a state-prediction network

Video input

The first idea: a state-prediction network

Video input

Neural network

The first idea: a state-prediction network

Video input

Neural network

Prediction states

The first idea: a state-prediction network

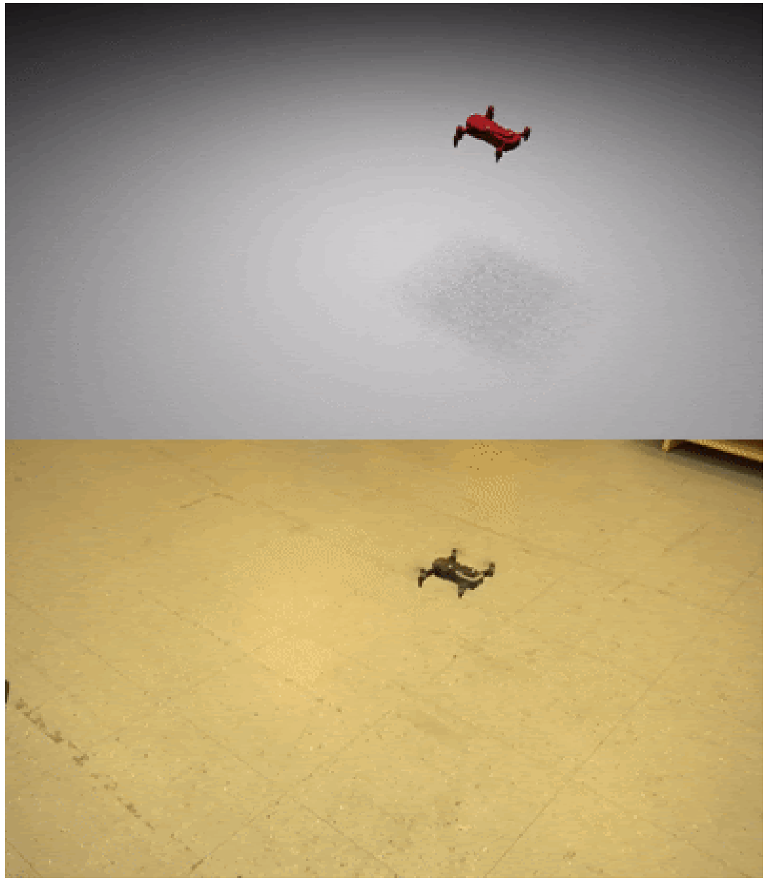

Neural network

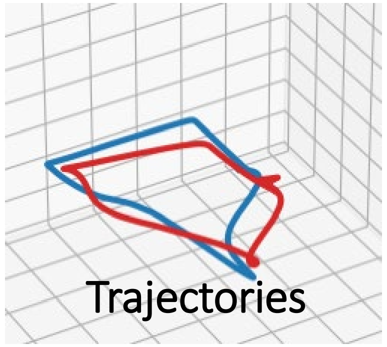

The state-of-the-art approach

Train the network using **domain randomization**.

Domain randomization failed here…

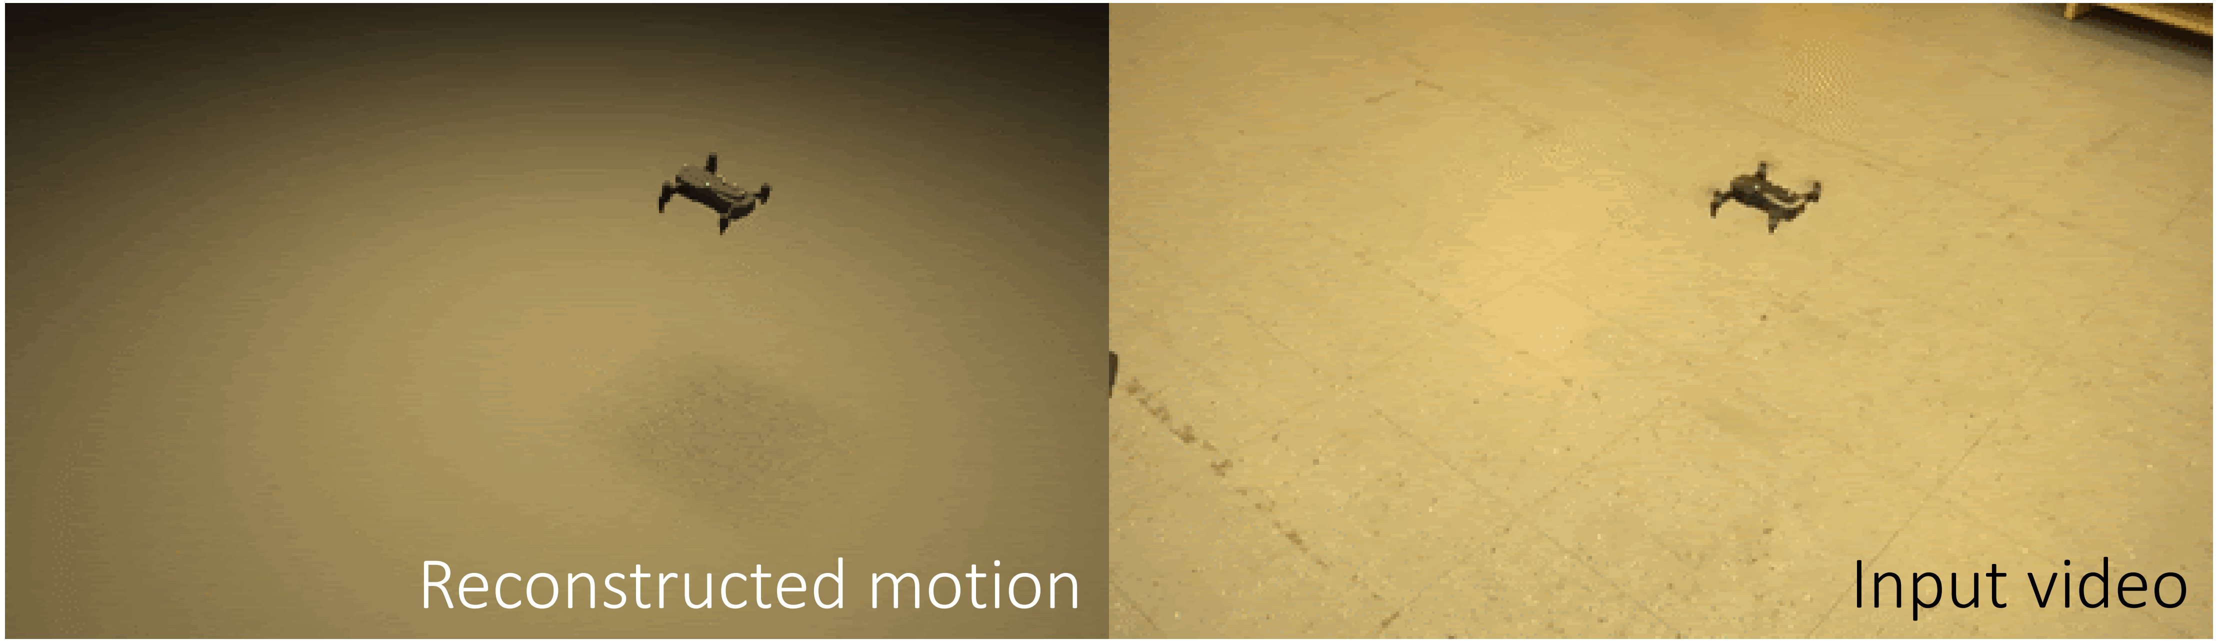

Why did domain randomization fail?

Video input

Neural network

Prediction
states

Why did domain randomization fail?

Dynamic model
states
actions

Visual appearance
lighting
material
texture

Video input

Neural network

Prediction
states

Why did domain randomization fail?

The network needs to maintain invariance under different visual appearances.

Dynamic model
states
actions

Visual appearance
lighting
material
texture

Video input

Neural network

Prediction
states

The second idea: rendering-invariance

The network needs to maintain invariance under different visual appearances.

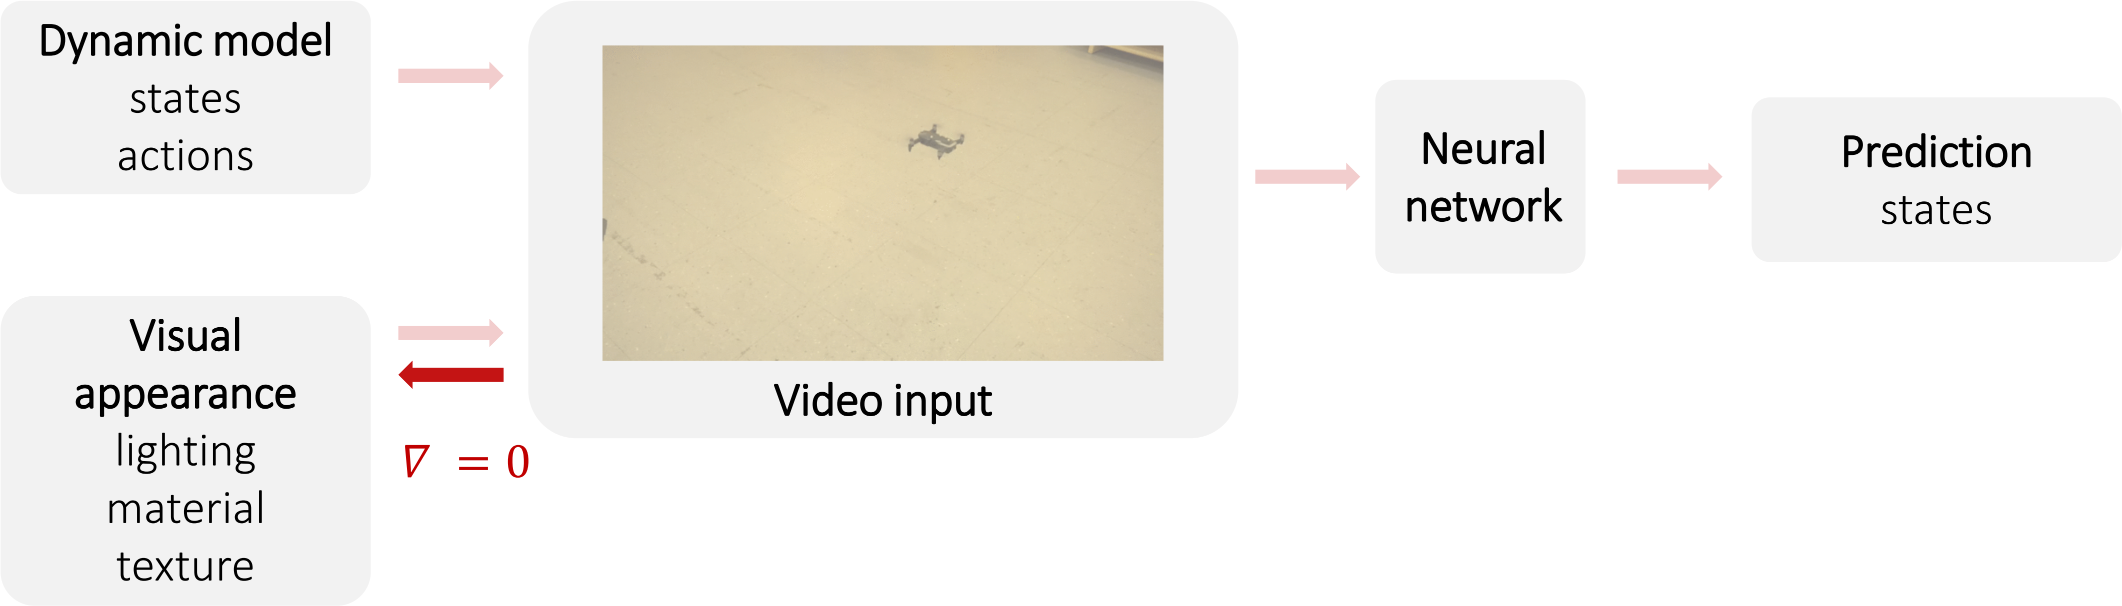

Results: quadrotors

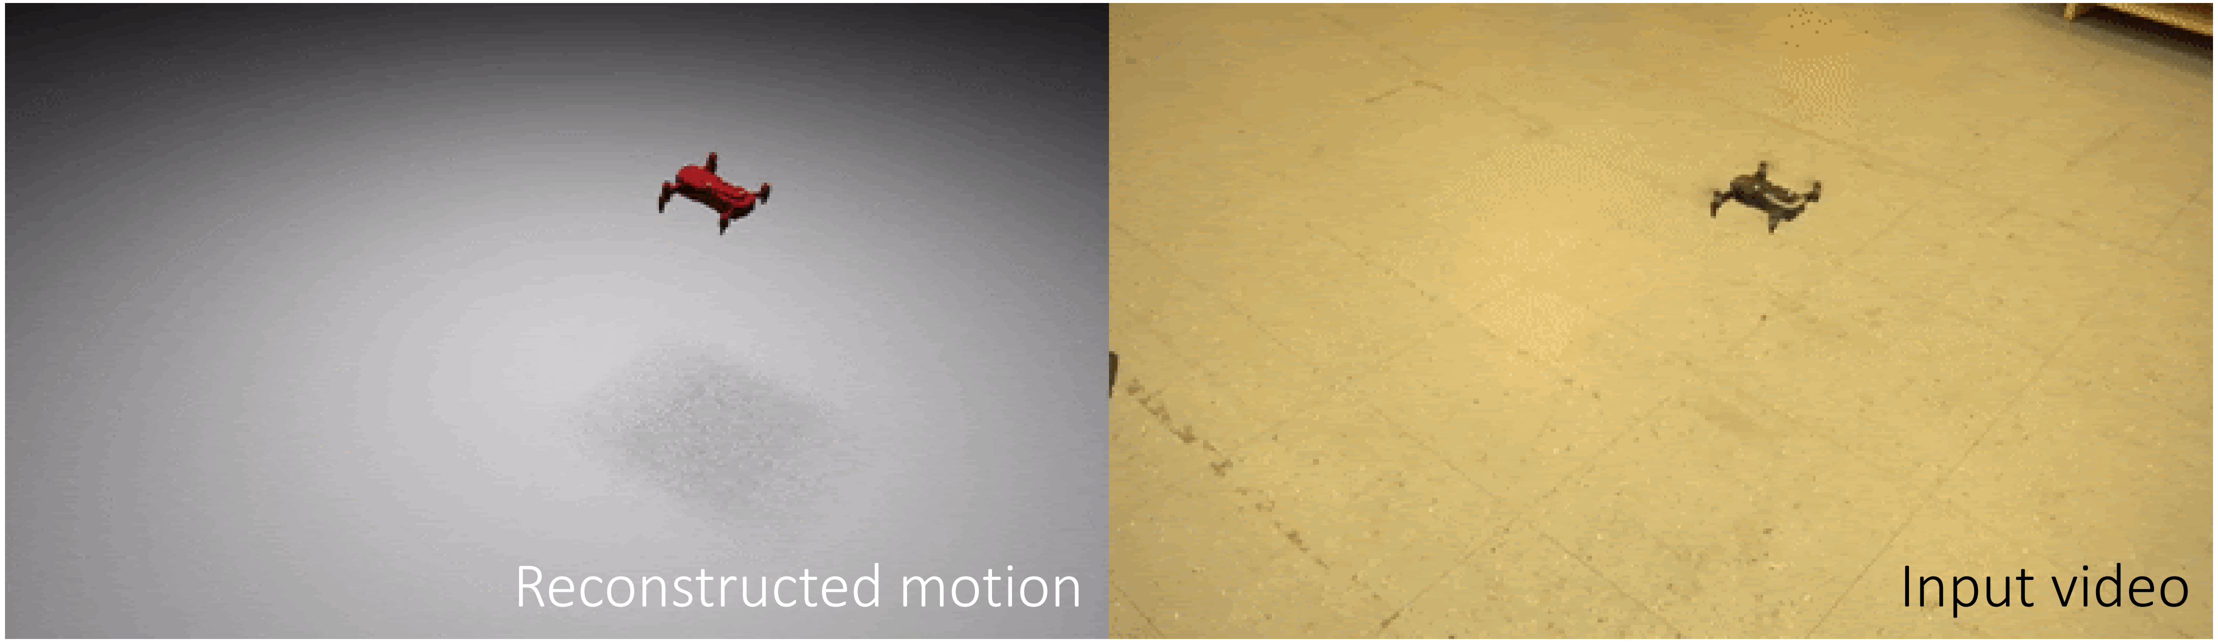

Note that the rendering configuration is intentionally made different.

Results: dexterous hand

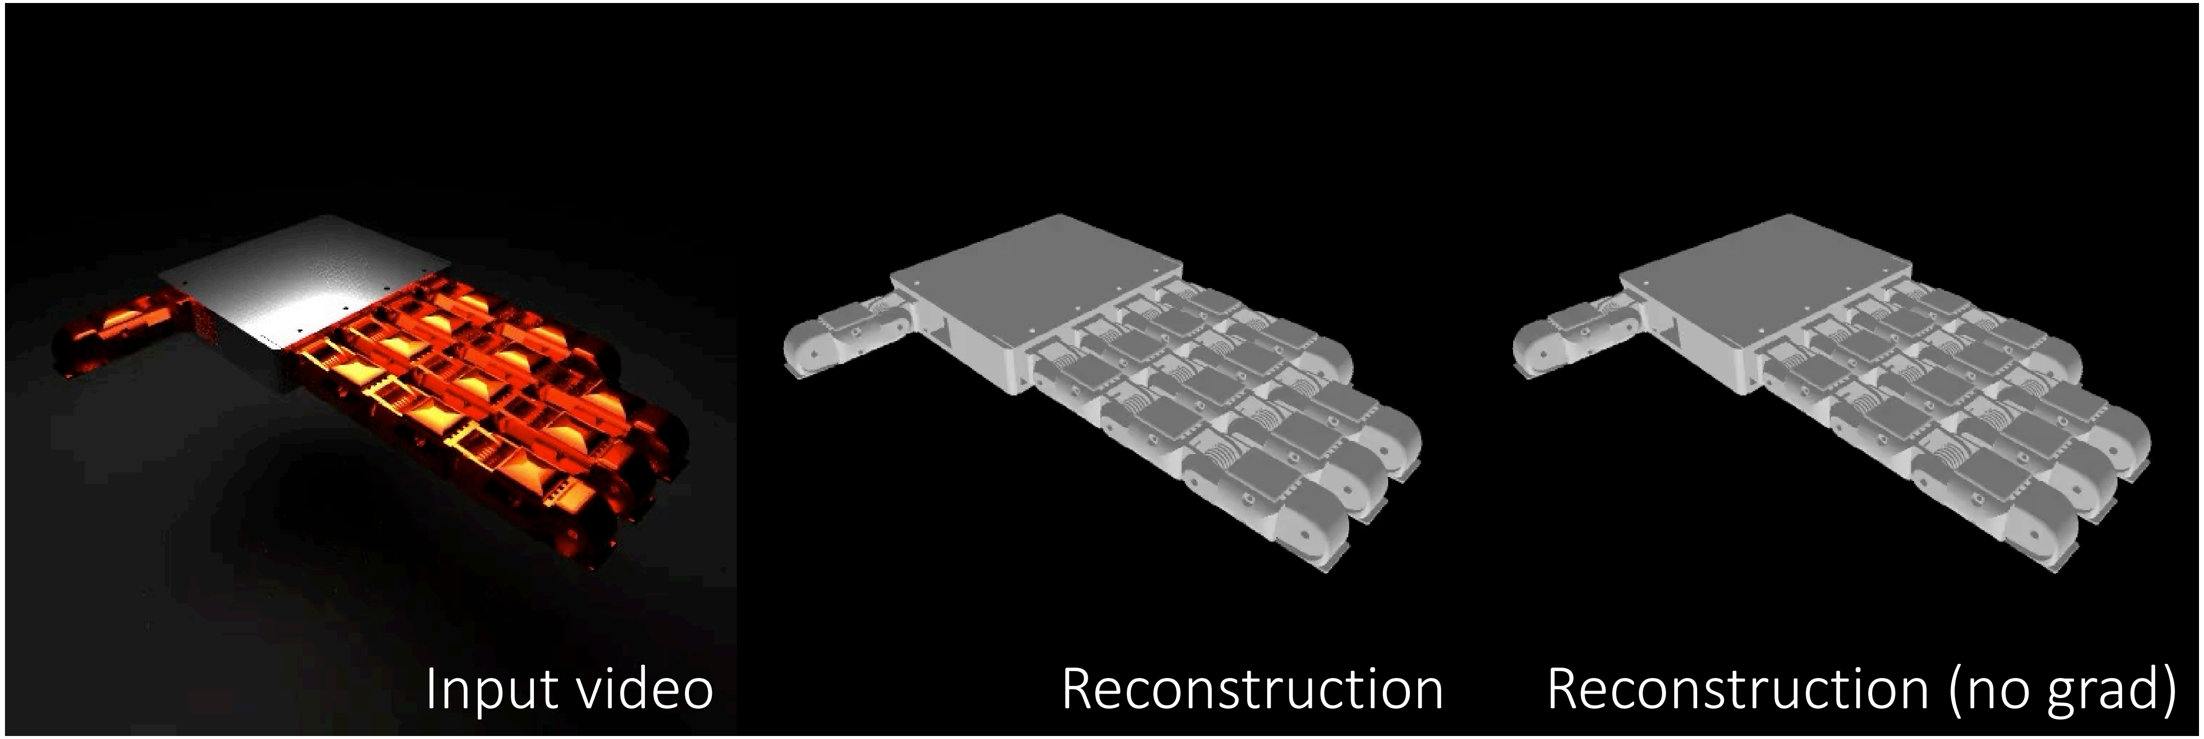

Note that the rendering configuration is intentionally made different.

Summary

∇ Parametrization	∇ Modeling	∇ Evaluation
Initializing parameters	Deriving governing equations	Computing performance metrics

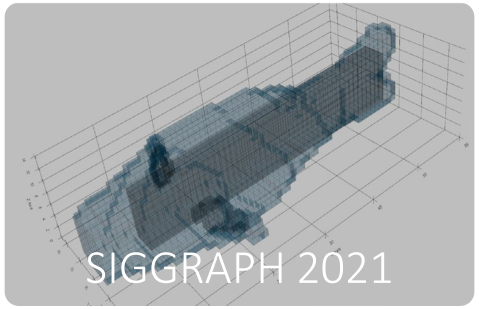

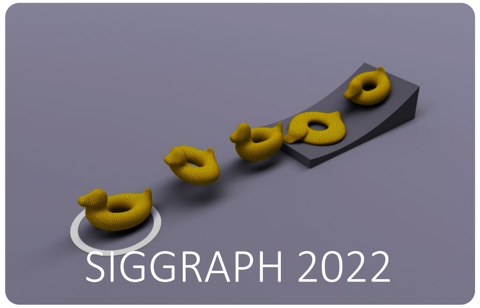

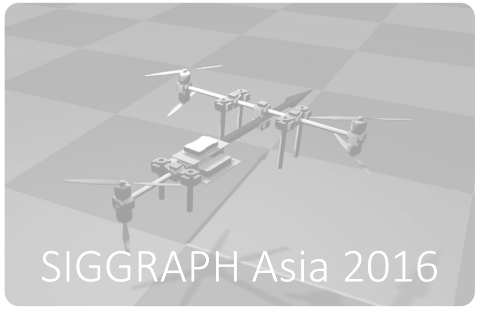

What is next?

Let me end the talk with what I consider one of the most inspiring inverse dynamics problems in history.

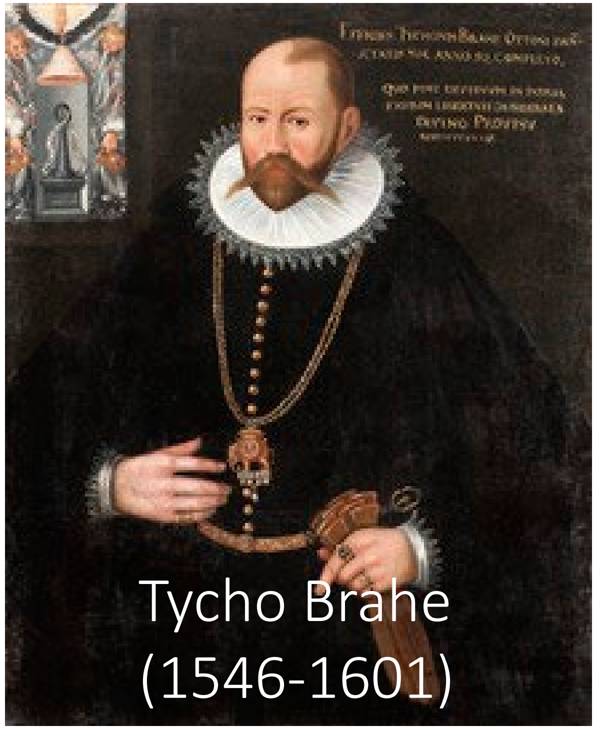
Tycho Brahe
(1546-1601)

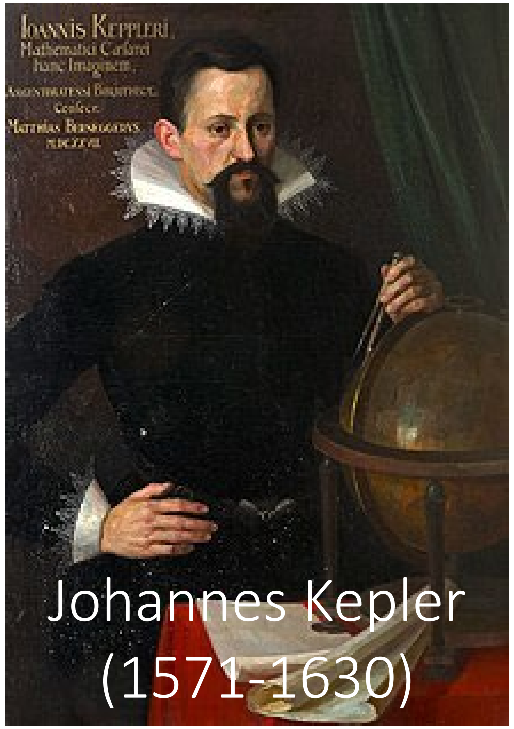
Johannes Kepler
(1571-1630)

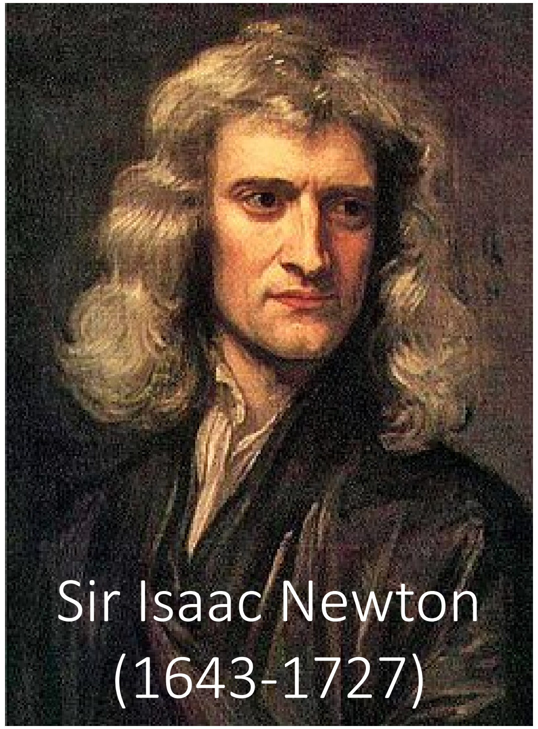
Sir Isaac Newton
(1643-1727)

What is next?

The most rewarding inverse problem is to discover scientific laws.

Acknowledgment

The papers covered in this talk were funded by the following sponsors:

Acknowledgment

Page 2: ANYmal: https://rsl.ethz.ch/robots-media/anymal.html.

Page 2: SCLS: http://www.scls.riken.jp/en/research/01_dynamics/index.html.

Page 2: Ravuri, S., Lenc, K., Willson, M. et al. *Skilful precipitation nowcasting using deep generative models of radar*. Nature 597, 672–677 (2021). https://doi.org/10.1038/s41586-021-03854-z.

Page 3: Natural Portfolio. https://www.nature.com/subjects/dynamical-systems.

Page 14: Pfaff et al. *Learning mesh-based simulation with graph networks*. ICLR 2021.

Page 14: Chen et al. *A system for general in-hand object re-orientation*. CoRL 2021.

Page 17: Hahn et al. *Real2Sim: visco-elastic parameter estimation from dynamic motion*. SIGGRAPH Asia 2019.

Page 17: Peng et al. *DeepMimic: Example-guided deep reinforcement learning of physics-based character skills*. SIGGRAPH 2018.

Page 25: Katzschmann et al. *Exploration of underwater life with an acoustically controlled soft robotic fish*. Science Robotics 2018.

Page 25: Project CETI. https://www.projectceti.org/.

Acknowledgment

Page 29: Video credit to Jie Xu.

Page 42: Hu et al. *ChainQueen: A real-time differentiable physical simulator for soft robotics*. ICRA 2019.

Page 42: Bouaziz et al. *Projective dynamics: Fusing constraint projections for fast simulation*. SIGGRPAH 2014.

Page 42: Geilinger et al. *ADD: Analytically differentiable dynamics for multi-body systems with frictional contact*. SIGGRAPH Asia 2020.

Page 67: Tycho Brahe. https://en.wikipedia.org/wiki/Tycho_Brahe.

Page 67: Johannes Kepler: https://en.wikipedia.org/wiki/Johannes_Kepler.

Page 67: Isaac Newton: https://en.wikipedia.org/wiki/Isaac_Newton.

Thank you!

Papers, code, and data are available at https://people.csail.mit.edu/taodu

Email: taodu@csail.mit.edu