DiffPhyCon: A Generative Approach to Control Complex Physical Systems

[1]Long Wei* , [2]Peiyan Hu*#, [1]Ruiqi Feng*, [1]Haodong Feng, [3]Yixuan Du#, [1]Tao Zhang, [4]Rui Wang#, [5]Yue Wang, [2]Zhi-Ming Ma, [1]Tailin Wu

(* equal contributions; # intern at Westlake University)

[1]Westlake University, School of Engineering,

[2]Academy of Mathematics and Systems Science, Chinese Academy of Sciences,

[3]Jilin University, [4]Fudan University, [5]Microsoft Research Asia

Corresponding to: wutailin@westlake.edu.cn

Outline

- Introduction
- Approach
- Results

*

Problem Setup

- Complex physical system control task : for a control objective $\mathcal{J}$, find the optimal control signal $\mathbf{w}^*$ such that $\mathbf{w}^*$ and the resulted system states $\mathbf{u}$ minimize $\mathcal{J}$ under physical dynamics constraints $\mathcal{C}(\boldsymbol{u}, \boldsymbol{w}) = \mathbf{0}$:

$$\mathbf{w}^* = \mathrm{argmin}_{\mathbf{w}} \mathcal{J}(\mathbf{u}, \mathbf{w}),$$
$$\mathrm{s.t.}\ \mathcal{C}(\boldsymbol{u}, \boldsymbol{w}) = \mathbf{0}$$

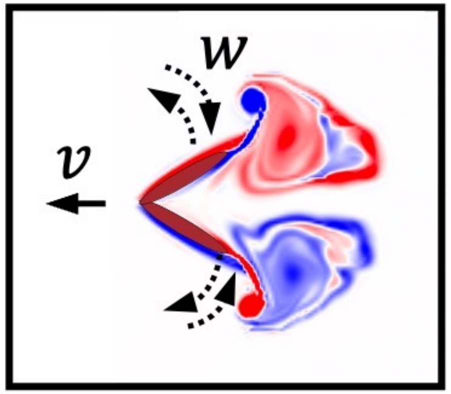

- E.g. how to control movement of wings of a jellyfish, such that it could achieve the highest speed in fluid, under the constraints of its boundary shapes and fluid dynamics

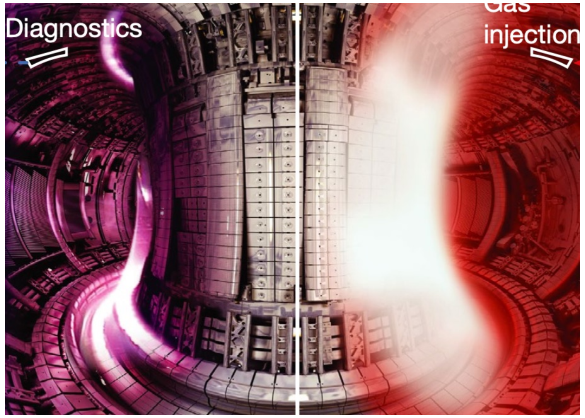

Fusion control

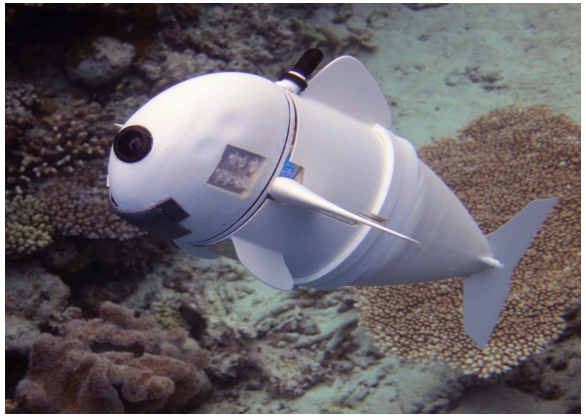
Underwater robot control


Rocket control

Key challenges

- Physical systems are typically high-dimensional, highly nonlinear
- Observed control signals are far from optimal solutions

Prior Works

- Classical numerical methods
- Pros: (1) first principle-based, (2) accurate, (3) with guaranteed error
- Cons: (1) computationally costly, (2) need rich expert knowledge, (3) weak at high-dimensional problems

- Recent deep learning-based and reinforcement methods
- Pros: (1) less engineering efforts, (2) offers speedup
- Cons: suffer from adversarial/myopic mode

Approach

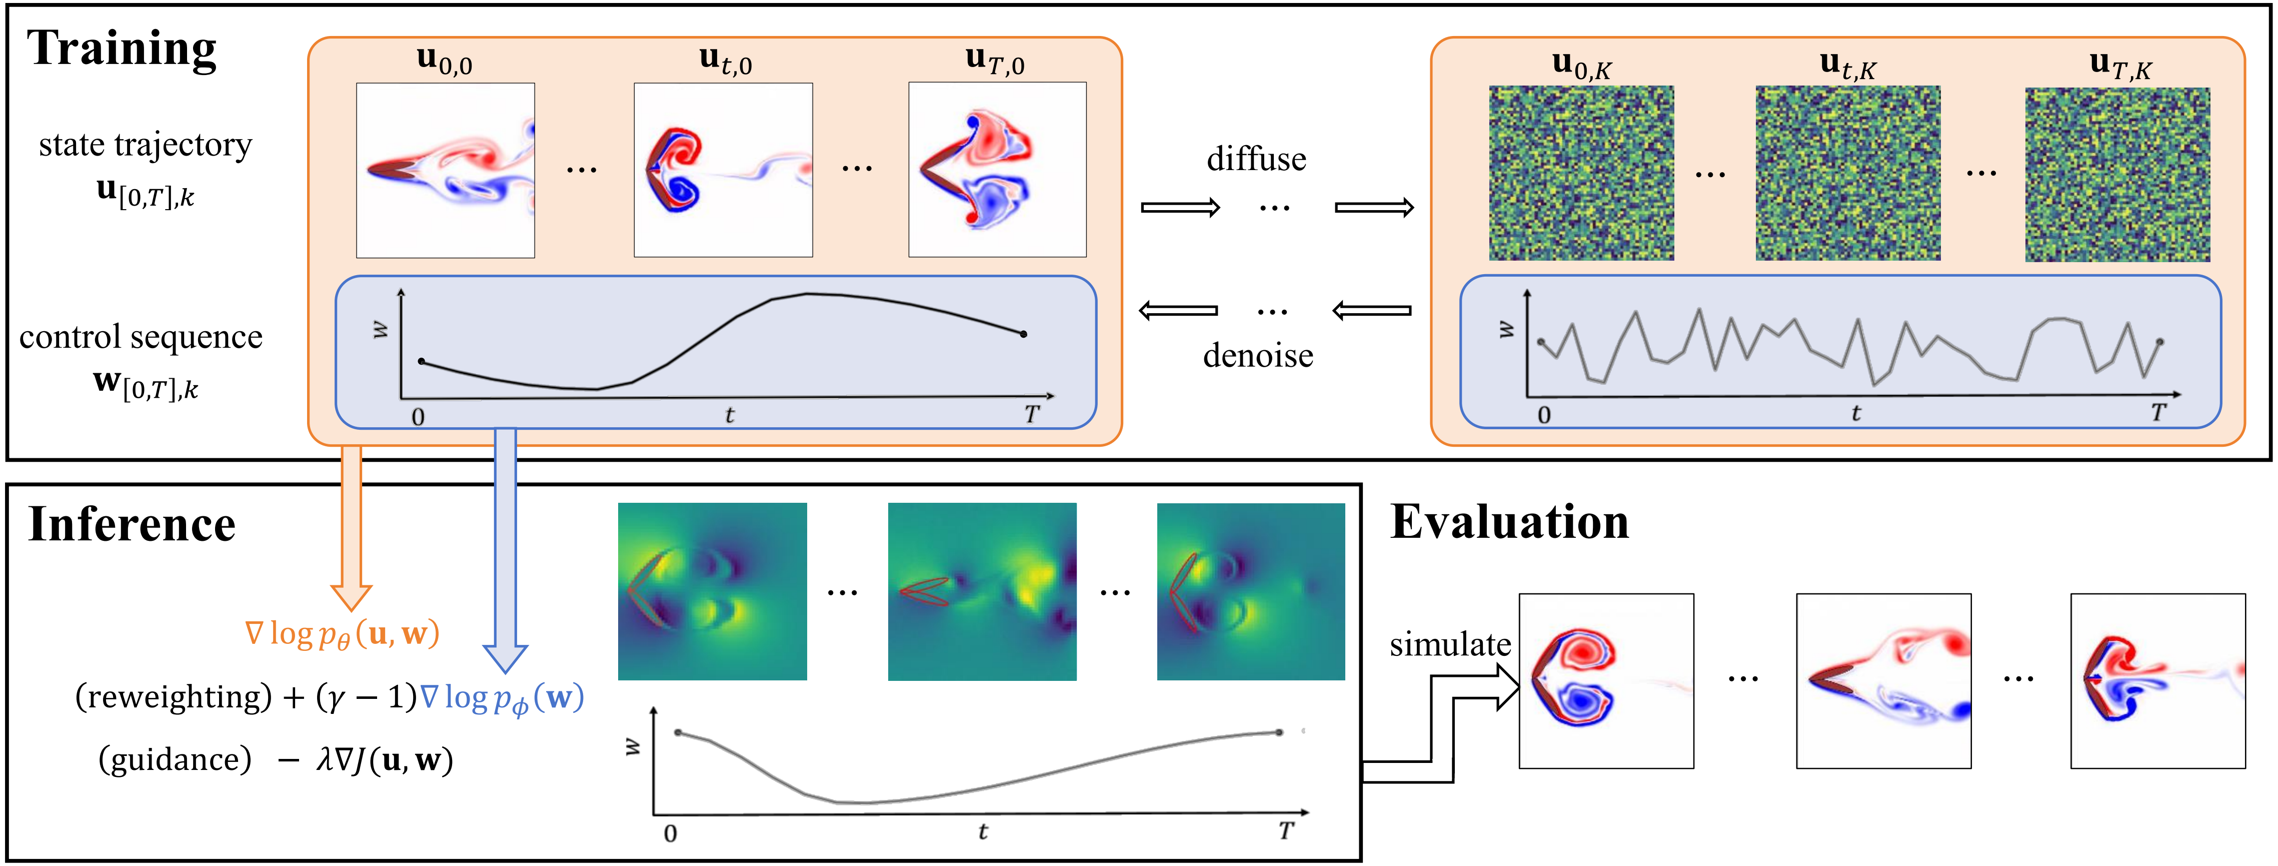

Training
state trajectory
$\mathbf{u}_{[0,T],k}$
control sequence
$\mathbf{w}_{[0,T],k}$
$\mathbf{u}_{0,0}$ $\mathbf{u}_{t,0}$ $\mathbf{u}_{T,0}$
$\mathbf{u}_{0,K}$ $\mathbf{u}_{t,K}$ $\mathbf{u}_{T,K}$
diffuse
denoise
Inference
$\nabla \log p_\theta(\mathbf{u}, \mathbf{w})$
(reweighting) $+ (\gamma - 1)\nabla \log p_\phi(\mathbf{w})$
(guidance) $- \lambda \nabla J(\mathbf{u}, \mathbf{w})$
Evaluation
simulate

Approach - EBM Perspective

- Reformulate the physical system control task as:

$$\mathbf{u}^*, \mathbf{w}^* = \mathrm{argmin}_{\mathbf{u},\mathbf{w}}[E_\theta(\mathbf{u},\mathbf{w}) + \lambda \cdot \mathcal{J}(\mathbf{u},\mathbf{w})]$$

E_θ: energy-based model (EBM); serves the purpose of a surrogate model in approximating PDE constraints

E_θ is learned by diffusion models $\boldsymbol{\epsilon}_\theta$: $\nabla_\theta \boldsymbol{E}_\theta \approx \boldsymbol{\epsilon}_\theta$

$$p(\mathbf{u},\mathbf{w}) \propto \exp(-E(\mathbf{u},\mathbf{w}))$$

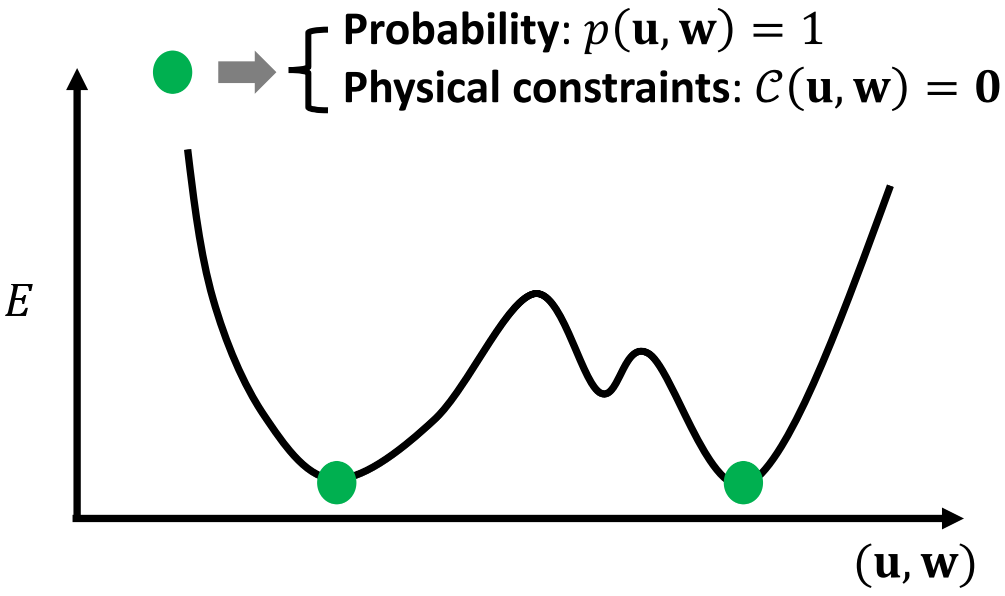

Approach - EBM Perspective

- Reformulate the physical system control task as:

$$\mathbf{u}^*, \mathbf{w}^* = \mathrm{argmin}_{\mathbf{u},\mathbf{w}}[E_\theta(\mathbf{u}, \mathbf{w}) + \lambda \cdot \mathcal{J}(\mathbf{u}, \mathbf{w})]$$

E_θ: energy-based model (EBM); serves the purpose of a surrogate model in approximating PDE constraints

E_θ is learned by diffusion models $\boldsymbol{\epsilon}_\theta$: $\nabla_\theta \boldsymbol{E}_\theta \approx \boldsymbol{\epsilon}_\theta$

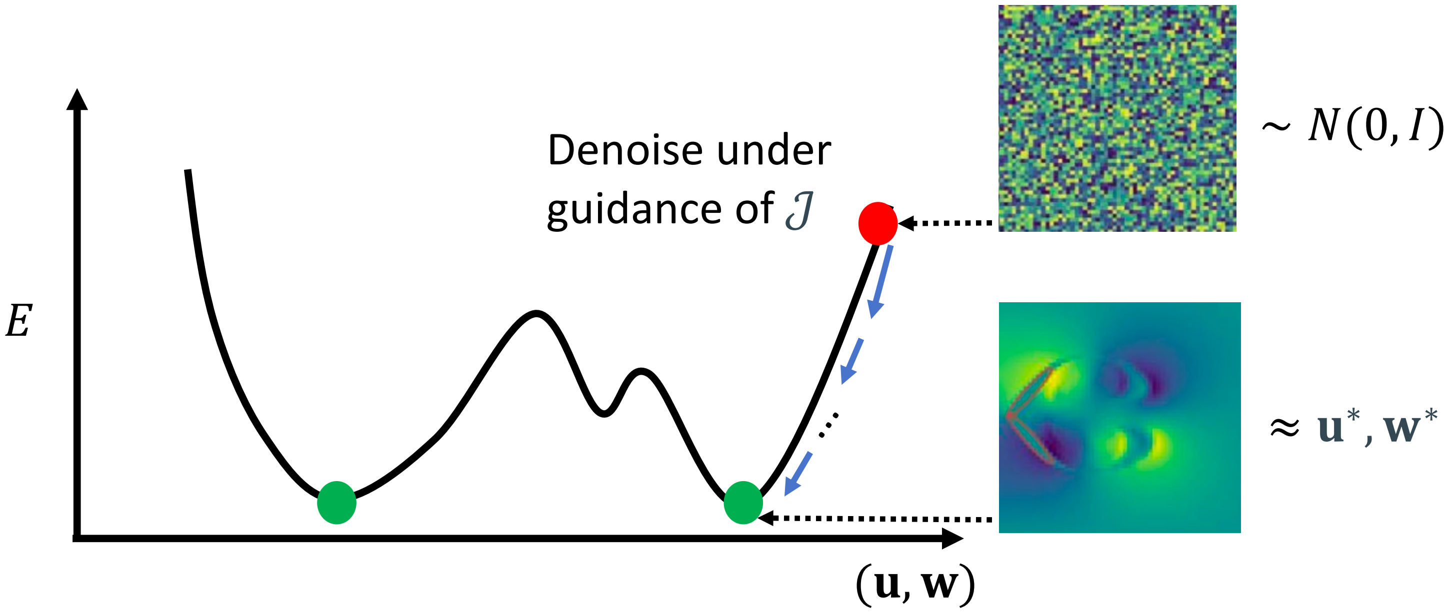

Approach - EBM Perspective

- Reformulate the physical system control task as:

$$\mathbf{u}^*, \mathbf{w}^* = \text{argmin}_{\mathbf{u},\mathbf{w}}[E_\theta(\mathbf{u},\mathbf{w}) + \lambda \cdot \mathcal{J}(\mathbf{u},\mathbf{w})]$$

E_θ: energy-based model (EBM); serves the purpose of a surrogate model in approximating PDE constraints

E_θ is learned by diffusion models $\boldsymbol{\epsilon}_\theta$: $\boldsymbol{\nabla_\theta E_\theta \approx \epsilon_\theta}$

- Training:

$$\text{Loss } \mathcal{L} = \mathbb{E}_{k\sim U(0,K),\mathbf{z}\sim p(\mathbf{z}),\epsilon\sim\mathcal{N}(\mathbf{0},\mathbf{I})}[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_k}\mathbf{z} + \sqrt{1-\bar{\alpha}_k}\boldsymbol{\epsilon}, k)\|_2^2], \text{ where } \mathbf{z} = [\mathbf{u},\mathbf{w}]$$

- Inference (sampling)

$$\mathbf{z}_K \sim \mathcal{N}(\mathbf{0},\mathbf{I}),$$

$$\mathbf{z}_{k-1} = \mathbf{z}_k - \eta\big(\boldsymbol{\epsilon}_\theta(\mathbf{z}_k, k) + \lambda\nabla_{\mathbf{z}}\mathcal{J}(\hat{\mathbf{z}}_k)\big) + \boldsymbol{\xi}_1, \boldsymbol{\xi}_1 \sim \mathcal{N}(\mathbf{0},\mathbf{I}),$$

where $\hat{\mathbf{z}}_k$ is a noise-free estimation of $\mathbf{z}_0$, $\mathbf{z}_k = [\mathbf{u}_k, \mathbf{w}_k]$

Approach - Prior Reweighting

- Motivation: how to obtain control sequences superior to those in training dataset?
- Reweighted joint distribution ($0 < \gamma \leq 1$)

$$p_\gamma(\mathbf{u}, \mathbf{w}) := p^\gamma(\mathbf{w}) p(\mathbf{u} \mid \mathbf{w})/Z = p^{\gamma-1}(\mathbf{w}) p(\mathbf{u}, \mathbf{w})/Z$$ (Z is the normalization constant)

- Reweighted energy based model form by taking logarithm:

$$E^{(\gamma)}(\mathbf{u}, \mathbf{w}) = (\gamma - 1) E_\phi(\mathbf{w}) + E_\theta(\mathbf{u}, \mathbf{w}) - \log Z$$

- Similarly, learn $\nabla E_\phi(\mathbf{w})$ by diffusion model ϵ_ϕ
- Inference (sampling):

$$\mathbf{z}_{k-1} = \mathbf{z}_k - \eta\left(\epsilon_\theta(\mathbf{z}_k, k) + \lambda \nabla_{\mathbf{z}} \mathcal{J}(\hat{\mathbf{z}}_k)\right) + \boldsymbol{\xi}_1,$$

$$\mathbf{w}_{k-1} = \mathbf{w}_{k-1} - \eta(\gamma - 1)\epsilon_\phi(\mathbf{w}_k, k) + \boldsymbol{\xi}_2,$$

where $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{z}_k = [\mathbf{u}_k, \mathbf{w}_k]$

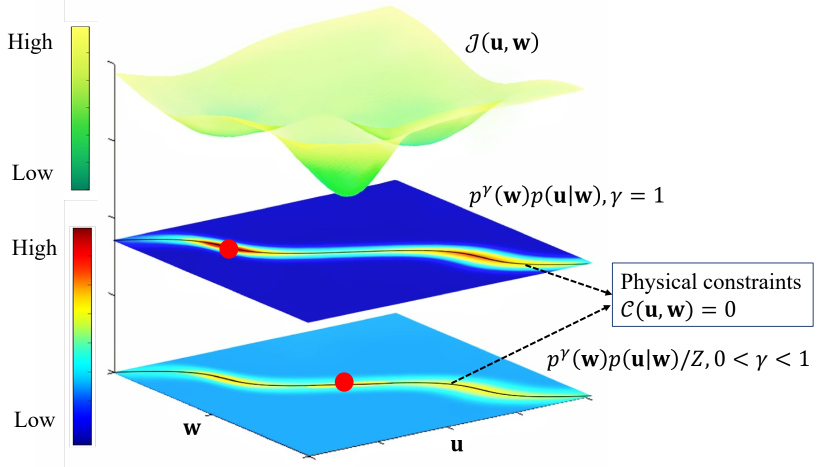

Results

- Our method tested in 3 different control tasks across 1D Burgers' equation and 2D Navier-Stokes equation.
  - 1D Burgers' equation state control
  - 2D jellyfish control
  - 2D smoke movement control

- DiffPhyCon demonstrates superior control performance
  - Better control metrics compared widely used RL methods.
  - A fast-close-slow-open pattern unveiled in 2D jellyfish movement, aligning with established findings in fluid dynamics

Results - 1D Burgers' Equation

$$\begin{cases} \dfrac{\partial \mathbf{u}}{\partial t} = -\mathbf{u} \cdot \dfrac{\partial \mathbf{u}}{\partial \boldsymbol{x}} + \nu \dfrac{\partial^2 \mathbf{u}}{\partial \boldsymbol{x}^2} + \mathbf{w}(t, \boldsymbol{x}), \text{in } [0,T] \times \Omega \\ \mathrm{u}(t, \boldsymbol{x}) = \mathbf{0}, \quad \text{on } [0,T] \times \partial\Omega \\ \mathrm{u}(0, \boldsymbol{x}) = \mathbf{u_0}(\boldsymbol{x}), \quad \text{in } \{0\} \times \Omega \end{cases}$$

Control objective: ($u_d(x)$ is target state)

$$J_{actual} = \int |u(T,x) - u_d(x)|^2 \, \mathrm{d}x$$

Energy cost: $\int |w(t,x)|^2 \mathrm{d}t\mathrm{d}x$

	PO-FC	FO-PC	PO-PC
PID (surrogate-solver)	-	0.09115	0.09631
SL	0.09752	0.00078	0.02328
SAC (surrogate-solver)	0.01577	0.03426	0.02149
SAC (offline)	0.03201	0.04333	0.03328
BC	0.02836	0.00856	0.00952
BPPO	0.02771	0.00852	0.00891
DiffPhyCon-lite (ours)	0.01139	**0.00037**	**0.00494**
DiffPhyCon (ours)	**0.01103**	**0.00037**	**0.00494**

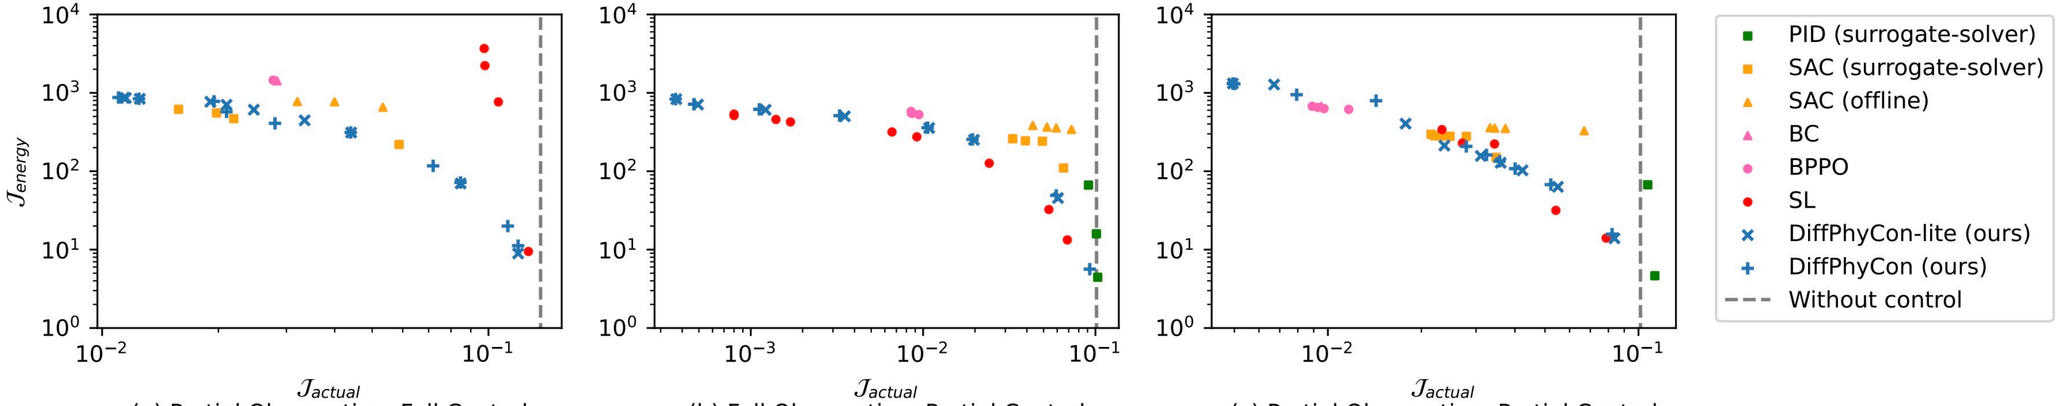

(a) Partial Observation, Full Control (b) Full Observation, Partial Control (c) Partial Observation, Partial Control

DiffPhyCon achieves the best Performance

Results - 2D Jellyfish Movement Control

The implicit physical dynamic is Navier-Stokes Equation:

$$\begin{cases} \dfrac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} - \nu \nabla^2 \mathbf{v} + \nabla p = 0 \\ \nabla \cdot v = 0 \\ \mathbf{v}(0, \boldsymbol{x}) = \mathrm{v}_{\mathbf{0}}(\boldsymbol{x}) \end{cases}$$

Control objective: maximize average moving speed $\bar{v}$ of the jellyfish, under energy cost constraints $R(\mathbf{w})$ of opening angles w of it wings:

$$\mathcal{J} = -\bar{v} + \zeta \cdot R(\mathbf{w})$$

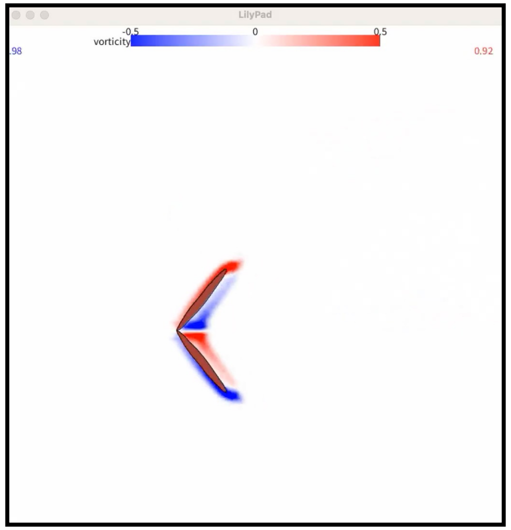

Results - 2D Jellyfish Movement Control

The implicit physical dynamic is Navier-Stokes Equation:

$$\begin{cases} \frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} - \nu \nabla^2 \mathbf{v} + \nabla p = 0 \\ \nabla \cdot v = 0 \\ \mathbf{v}(0, \boldsymbol{x}) = \mathrm{v}_{\mathbf{0}}(\boldsymbol{x}) \end{cases}$$

Control objective: maximize average moving speed $\bar{v}$ of the jellyfish, under energy cost constraints $R(\mathbf{w})$ of opening angles w of it wings:

$$\mathcal{J} = -\bar{v} + \zeta \cdot R(\mathbf{w})$$

	Full observation			Partial observation		
	$\bar{v}\uparrow$	$R(\mathbf{w})\downarrow$	$\mathcal{J}\downarrow$	$\bar{v}\uparrow$	$R(\mathbf{w})\downarrow$	$\mathcal{J}\downarrow$
MPC	25.72	0.0112	109.17	-150.51	0.1791	329.59
SL	-76.94	0.1286	205.57	-102.98	0.1188	221.79
SAC (surrogate-solver)	-166.96	0.0069	18.14	-153.09	0.0057	158.82
SAC (offline)	-158.66	0.0069	165.58	-206.21	0.0058	211.96
BC	30.48	0.0629	32.44	20.08	0.0556	35.48
BPPO	107.67	0.0867	-20.93	54.83	0.0518	-3.02
DiffPhyCon-lite (ours)	95.04	0.0746	-20.47	2.92	0.0779	74.97
DiffPhyCon (ours)	**279.87**	0.2058	**-74.11**	**150.21**	0.1269	**-23.32**

DiffPhyCon achieves the highest moving speed and lowest control objective.

Results - 2D Jellyfish Movement Control

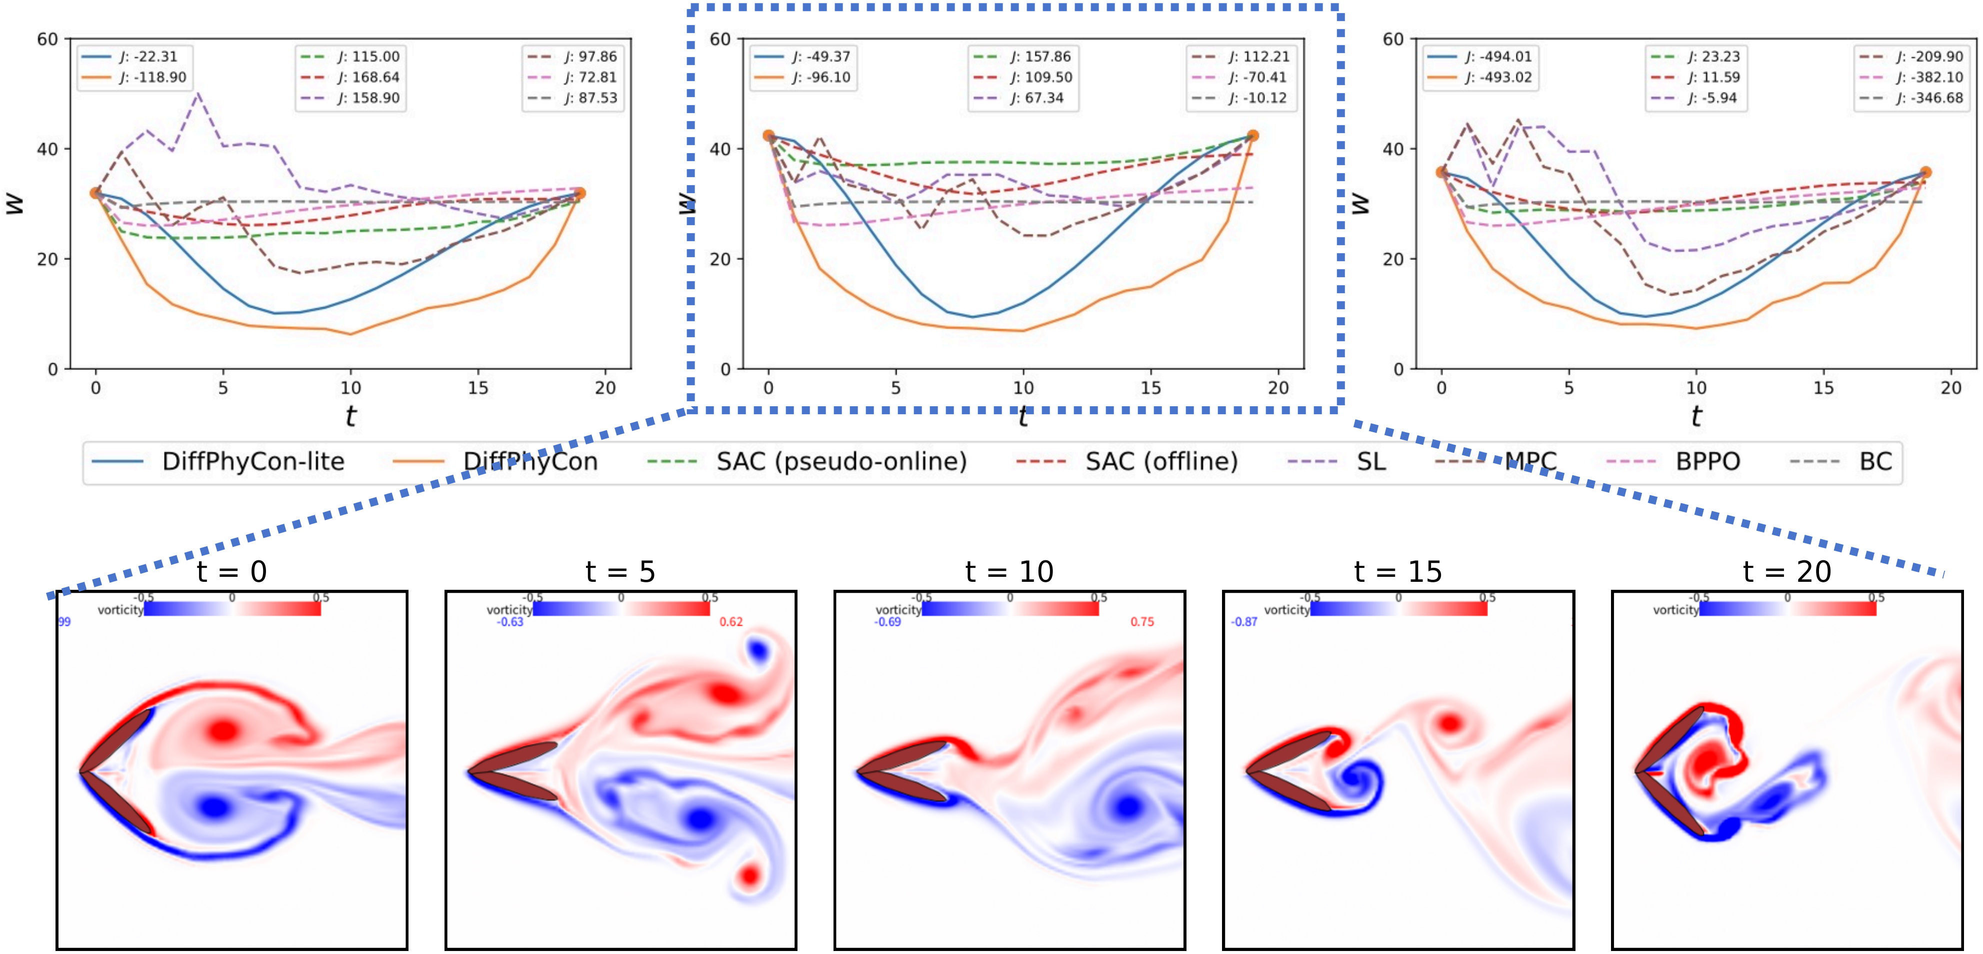

Our method presents a desired fast-close-slow-open pattern.

Results - 2D Jellyfish Movement Control

Propulsive performance and vortex dynamics of jellyfish-like propulsion with burst-and-coast strategy

Cite as: Phys. Fluids **35**, 091904 (2023); doi: 10.1063/5.0160878
Submitted: 6 June 2023 · Accepted: 12 August 2023 ·
Published Online: 6 September 2023

View Online Export Citation CrossMark

Linlin Kang (康林林),[1] An-Kang Gao (高安康),[2,a)] Fei Han (韩非),[1] Weicheng Cui (崔维成),[1,b)]
and Xi-Yun Lu (陆夕云)[2]

AFFILIATIONS

[1]Key Laboratory of Coastal Environment and Resources of Zhejiang Province, School of Engineering, Westlake University, Hangzhou, Zhejiang 310030, China
[2]Department of Modern Mechanics, University of Science and Technology of China, Hefei 230026, Anhui, China

[a)]Author to whom correspondence should be addressed: ankanggao@ustc.edu.cn
[b)]Electronic mail: cuiweicheng@westlake.edu.cn

ABSTRACT

The propulsive performance and vortex dynamics of a two-dimensional model for the jellyfish-like propulsion with burst-and-coast strategy are investigated using a penalty-immersed boundary method. The simplified model comprises a pair of pitching flexible plates with their leading edges connected. The effects of two key parameters are considered, i.e., the duty cycle (*DC*, the ratio of the closing phase to the whole period) and the bending stiffness (*K*). Three different wake patterns, i.e., periodic symmetric, periodic asymmetric, and chaotic wakes, are identified in the *DC*–*K* plane. Numerical results indicate that a significant fast-close-slow-open motion is more likely to achieve higher speed, efficiency, and stability than a slow-close-fast-open motion, and proper higher bending stiffness is conducive to improving efficiency. A force decomposition based on the weighted integral of the second invariant of the velocity gradient tensor is performed to gain physics insight into the self-propulsive mechanism. It is found that the repulsive force induced by the strain-rate field between the body and the previous vortex pair is the main driving force of the jellyfish-like motion and that capturing the previous vortex pair during the closing phase can significantly enhance the strain rate as well as the thrust. This clarifies why the jellyfish can achieve thrust by pushing back vortex pairs. This study provides inspiration for the design and control of flexible jet propulsion devices.

"Numerical results indicate that a significant fast-close-slow-open motion is more likely to achieve higher speed, efficiency, and stability"

--Kang et al, *Physics of Fluids*, 2023

Control results of DiffPhyCon are aligning with established findings in fluid dynamics

Results - 2D Smoke Control

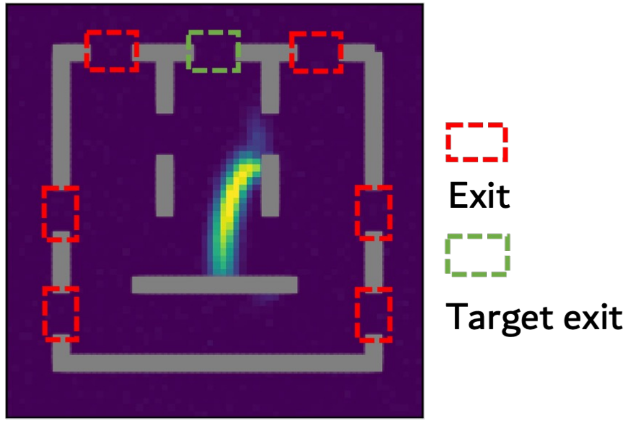

(a) Locations of exits and obstacles

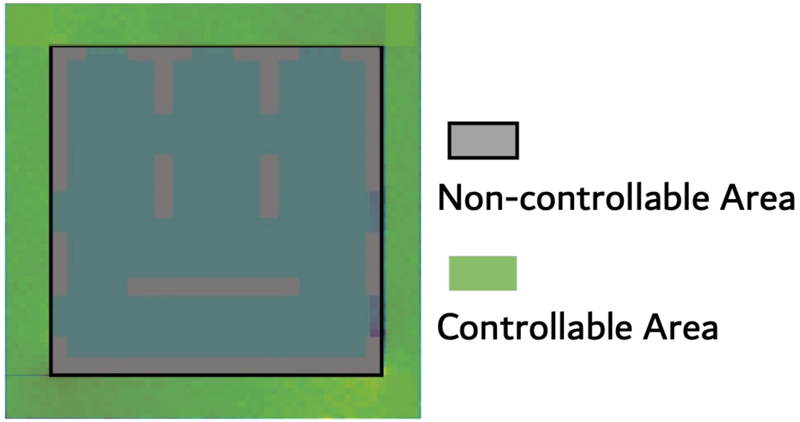

(b) Locations of controllable area

Method	$\mathcal{J}\downarrow$
BC	0.3085
BPPO	0.3066
SAC (surrogate-solver)	0.3212
SAC(offline)	0.6503
DiffPhyCon-lite (ours)	<u>0.2324</u>
DiffPhyCon (ours)	**0.2254**

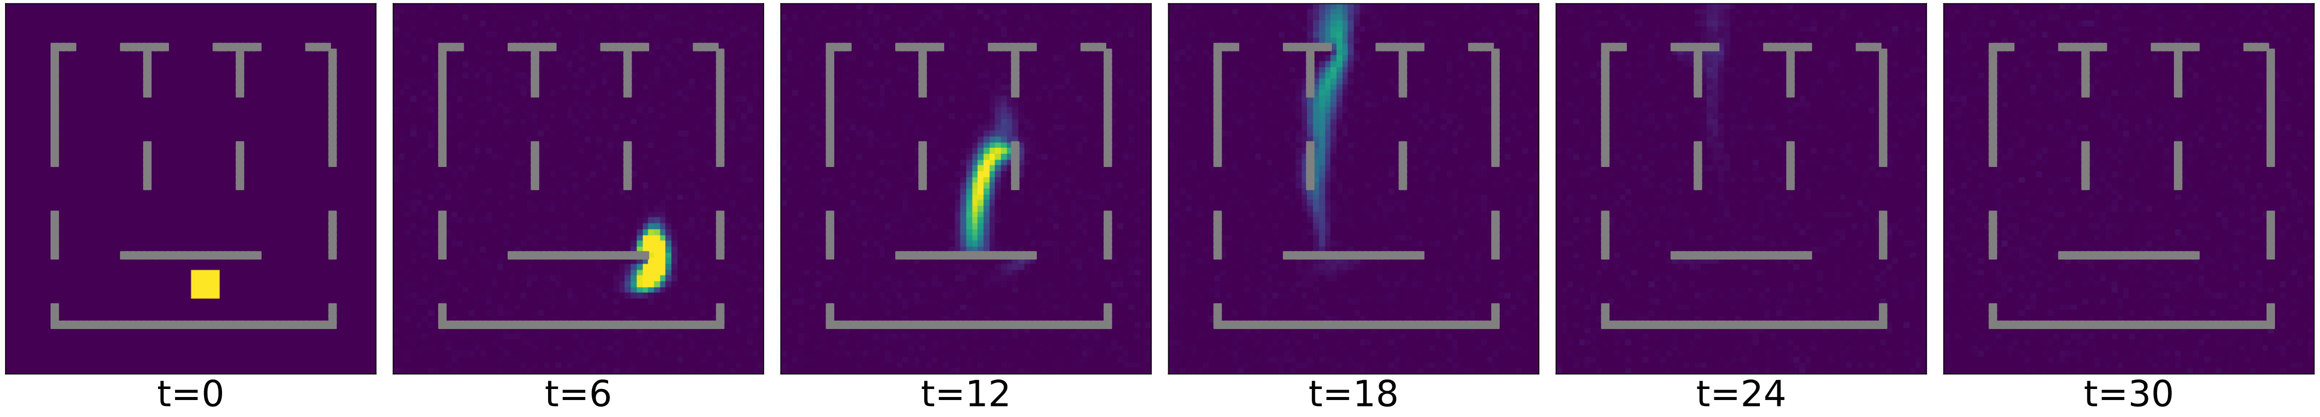

Limitation and Future Work

- Efficiency
  - The inference currently involves hundreds of denoising steps. How to accelerate inference process by using e.g. , distillation or DDIM sampling methods?
- Online training
  - The training is currently conducted in an offline fashion, lacking interaction with a ground-truth solver. Incorporating solvers into the training framework could adapt to dynamicl environment and discover novel strategies and solutions
- Closed-loop inference
  - Inference presently operates in an open-loop manner. Integrating feedback from environments would empower the algorithm to adjust subsequent control decisions based on the evolving state of the environment

Our group: AI for Scientific Simulation & Discovery Lab @ Westlake University


Long Wei

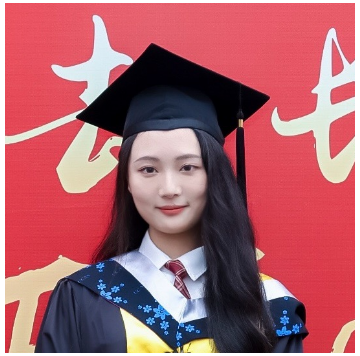
Peiyan Hu


Ruiqi Feng

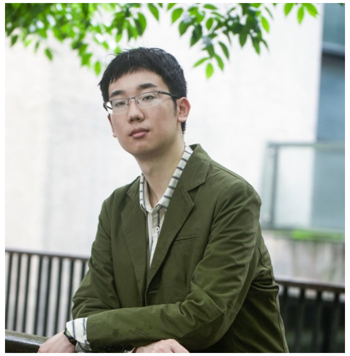
Haodong Feng

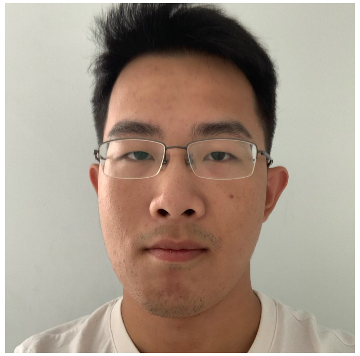
Yixuan Du


Tao Zhang

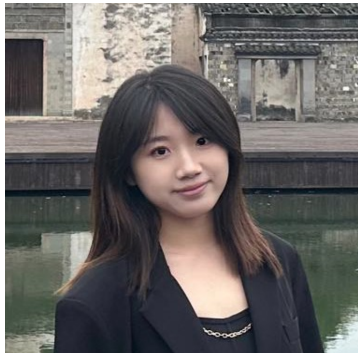
Rui Wang


Yue Wang


Zhi-Ming Ma


Tailin Wu

Paper:

Group Website:

Thank you!

Paper:

Group Website:

If you have any questions, please feel free to contact us at:

weilong@westlake.edu.cn

hupeiyan18@mails.ucas.ac.cn

fengruiqi@westlake.edu.cn

wutailin@westlake.edu.cn