jh240063

Physics Informed Machine Learning for Soft Matter

John Molina (Kyoto University)

Abstract

We have developed (physics-informed) Machine Learning methods to study complex Soft/Active matter systems, which are characterized by hierarchy of length- and time-scales that makes them incredibly challenging to study. We have focused on three basic themes: (A) Polymer Rheology, (B) Stokes Flows, and (C) Optimal Control of Active Systems. (A) We used ML to infer the relationship between the microscopic and macroscopic degrees of freedom of entangled polymer melts, this includes, among others, learning the constitutive relation for the stress, as well as inferring the molecular weight distribution from rheological data. (B) We develeoped a probabilistic inference framework for Stokes flow problems, capabable of inferring the flow solution given noisy and/or partial data. (C) We developed a reinforcement learning protocol to teach swimmers how to navigate complex flows, as well as infer the (hidden) utility of rational agents.

1 Basic information

1.1 Collaborating JHPCN centers

- The University of Tokyo
- Kyoto University

1.2 Theme area

- Large-scale computational science area

1.3 Project members and their roles

- J. J. Molina: ML for Soft Matter.
- T. Iwashita: HPC support.
- H. Shiba: support.
- T. Taniguchi: MSS for polymers.
- R. Yamamoto: Soft Matter theory.
- S. Miyamoto: ML/MSS for polymers.
- T. Sato: MSS for polymers.
- S. K. Schnyder: ML/Optimal control.
- D. Mayank: Polymer simulations.
- M. S. Turner: Soft Matter Theory.
- A. C. Meneses: Active Matter.
- M. Lynch: ML/Optimal control.
- Y. Ueno: ML for polymers.
- G. Iwami: Polymer simulations.
- S. Turley: ML for active Matter.

2 Purpose and Significance of the Research

Soft Matter systems (e.g., polymers, colloids, biofilms) are fundamental to modern industry. Typically, these systems possess a hierarchy of length- and time-scales that provide interesting material properties, while making them incredibly challenging to study. Theoretical approaches are limited to idealized systems, leaving computer simulations as the

preferred approach. However, this still requires vast computational resources. There is now a growing interest to use Machine Learning (ML), either to complement existing approaches, or to replace them altogether. The goal of this project is to develop ML methods to study characteristic Soft Matter systems. We considered three main themes: (A) Polymer rheology, (B) Stokes Flows, and (C) Optimal Control.

(Theme A) Polymer materials are one of the fundamental building blocks of modern industrialized societies. However, there is an urgent need to develop more efficient and sustainable materials. The current state-of-the-art method to study these systems relies on Multi-Scale-Simulations (MSS), which directly couple micro/macro degrees of freedom, but these are too expensive. To address this, we have developed a method to learn the constitutive relation for the stress from microscopic data. We have also developed a method to solve the inverse problem of inferring the microscopic properties given macroscopic data. This year we worked to improve the learning, and focused on solving the inverse problem (e.g., flow optimization, parameter inference).

(Theme B) Stokes flows, relevant to flows at small scales or high-viscosity fluids (e.g., biological flows or high-molecular weight polymer melts), are ubiquitous in life-sciences. While several numerical methods exist for these flow problems, they are not easily applied to experimental setups, where incomplete and noisy data is the norm. Thus, we have developed a physics-informed probabilistic ML framework that allows us to solve arbitrary Stokes flow problems. This year we worked to improve the method, to be able to analyze large scale 3D systems.

(Theme C) Active systems, composed of "agents" capable of responding to their environment to maximize their utility, are widespread in biological and social systems. We have considered two such systems: (1) "smart" microswimmers navigating non-uniform flows and (2) "rational" individuals self-distancing during an epidemic. For (1), we developed a reinforcement learning scheme to teach swimmers how to navigate using local hydrodynamic signals, for (2) we developed a learning method to infer the hidden utility. This year, we focused on (1) cargo-transport and (2) improving the robustness of the model and the learning.

3 Significance as JHPCN Joint Research Project

The goal of this project is to develop ML methods to simulate/analyze complex Soft Matter systems, with a focus on (A) ML for polymer rheology, (B) ML for Stokes flows, (C) ML for optimal control. The complexity of these systems necessarily requires large-scale computing resources. For example, (A) the full MSS for polymer melts flows can require simulating up to $\simeq 10^8$ polymer chains; (B) the Stokes flow inference requires exact Gaussian Process (GP) regression on $\simeq 10^6$ training points; (C) the flow navigation requires training through a hydrodynamic simulation and the utility inference relies on a set of nested neural networks.

4 Outline of Research Achievements until FY2023 (Only for continuous projects)

(Theme A) We developed a Gaussian Process (GP) regression scheme to learn the constitutive relation for the stress of entangled polymer melts (within the Doi-Takimoto Slip-Link model) and used it to predict the flow behavior under varying conditions. The constitutive relation was assumed to be a function of the stress $\boldsymbol{\sigma}$ and the strain rate $\boldsymbol{D} = (\boldsymbol{\kappa} + \boldsymbol{\kappa}^T)/2$ ($\boldsymbol{\kappa} = \boldsymbol{\nabla u}$ the velocity gradient), i.e., $\dot{\boldsymbol{\sigma}} = \dot{\boldsymbol{\sigma}}(\boldsymbol{\sigma}, \boldsymbol{D})$. In particular, the method was extended to allow for multi-deformation mode flows and to satisfy the principle of material objectivity. Comparisons with full MSS showed very good agreement at a fraction of the numerical cost, with the results obtained from the learned constitutive relation showing drastically reduced noise levels.

(Theme B) We developed a physics-informed Stokes flow inference framework capable of inferring the flow solution of arbitrary Stokes problems given noisy and/or partial data. In particular, our GP regression approach is able to exactly encode the underlying physical laws, i.e, the Stokes (force-balance) $\boldsymbol{f} = \boldsymbol{\nabla} p - \mu \nabla^2 \boldsymbol{u} = 0$ and continuity $s = \boldsymbol{\nabla} \cdot \boldsymbol{u} = 0$ equations. We obtain the solution in the form of a (posterior) probability distribution for the velocity and pressure, conditioned on all prior information (i.e., the physical laws, boundary conditions, and any measured data).

(Theme C) We developed a Reinforcement Learning protocol to teach swimmers to navigate non-uniform flow fields using purely local hydrodynamic signals. We trained (1) individual swimmers to navigate zig-zag shear flows, and (2) a collection of swimmers to produce non-trivial collective motion.

5 Details of FY2024 Research Achievements

5.1 ML for Polymer Rheology

We have developed a GP based ML method to solve the (inverse) problem of inferring the molecular weight distribution (MWD) of a linear polydisperse polymer melt from rheological data, as given by stress relaxation curves $G(t)$, which are easy to measure experimentally (in contrast to the MWD themselves). To account for multi-modal solutions, where distinct polymer compositions can give rise to indistinguishable relaxation curves, we estimate the full posterior distribution using Hamiltonian Monte-Carlo simulations. For simplicity, we have considered a heterogeneous mixture of N distinct species of Rouse chains, which differ only in the length of the chains. Training data for $G(t)$ was generated from the known analytical solution given the weight/length (i.e., the number of beads in a chain) distribution f_i ($\sum_{i=1}^{N} f_i = 1$).

Let $\boldsymbol{x} = (t, f_i)$ denote the model input, with $G(\boldsymbol{x})$ the model output. The solution to the inverse problem is given by the posterior probability distribution $p(\boldsymbol{x}_\star | \boldsymbol{G}_\star, \boldsymbol{x}, \boldsymbol{G})$, with $\boldsymbol{x}$ and $\boldsymbol{G}$ the training data sets, $\boldsymbol{x}_\star$ and $\boldsymbol{G}_\star$ the (unknown) test data/predictions. Using Bayes' theorem this distribution

can be shown to be proportional to $\int \mathrm{d}\boldsymbol{\Theta} p(\boldsymbol{G}_\star|\boldsymbol{\Theta}, \boldsymbol{x}_\star, \boldsymbol{x}, \boldsymbol{G}) p(\boldsymbol{G}|\boldsymbol{\Theta}, \boldsymbol{x}) p(\boldsymbol{\Theta}) p(\boldsymbol{x}_\star)$, where $p(\boldsymbol{G}_\star|\boldsymbol{\Theta}, \boldsymbol{x}_\star, \boldsymbol{x}, \boldsymbol{G})$ is the GP-posterior (i.e., the learned mapping between input/output), $p(\boldsymbol{G}|\boldsymbol{\Theta}, \boldsymbol{x})$ is the GP-prior for the stress-relaxation function, and $p(\boldsymbol{\Theta})$ and $p(\boldsymbol{x}_\star)$ are suitable priors for the GP hyperparameters $\boldsymbol{\Theta}$ and the number fraction $\boldsymbol{f}$.

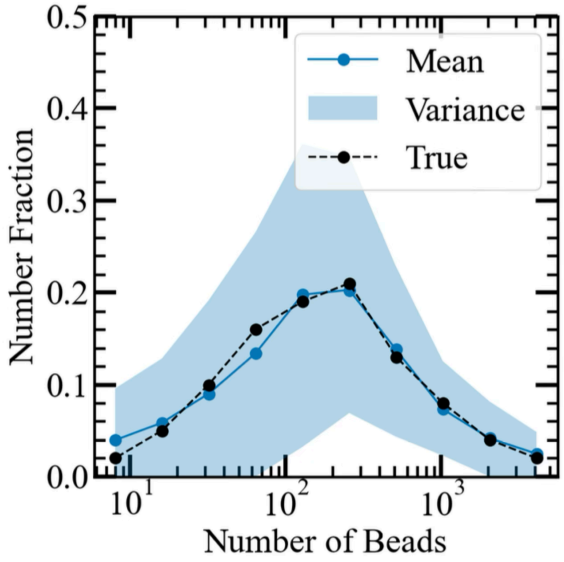

Fig. 1 Number fraction inference for a denary system of heterogeneous Rouse chains. Y. Ueno, Mater Thesis, Kyoto University (2025).

We discuss the results for the denary $N = 10$ case (the largest considered), where the shortest chain had $N = 8$ beads, and the longest chain $N = 4096$. To obtain a uniform sampling of training data points we sampled 500 sets of $\boldsymbol{f}$ using a Centroidal Voronoi Tessellation (CVT) and computed the corresponding $G(t)$ from the analytical solution to the Rouse model. The relation between $\boldsymbol{f}$ and $G(t)$ is learned using a GP regression, the posterior distributions for $\boldsymbol{f}_\star$ given $G_\star$ is sampled using the NUTS Hamiltonian Monte-Carlo method (as implemented in PYMC). The predicted distribution, as a function of the number of beads in a chain is given in Fig. 1. While the variance is relatively high, the mean prediction is in excellent agreement with the true value. To verify the applicability to experimental data, we studied the scalability of the method, by estimating how the cost of the Bayesian inference scales with the number of training data sets, as well as with the dimensions of the problem (i.e., the number of components). Results were obtained by performing each inference ten times for random training and test data sets. The calculation time is seen to scale linearly, both with the number of training data sets, as with the number of dimensions/components d (Fig. 2).

Finally, we have also performed systematic studies to understand the remarkable mechanical properties of natural rubber (NR), as compared to synthetic rubber. This has been attributed to the terminal groups of the polymer chains, and the presence of non-rubber components (NRCs), e.g. proteins, lipids, fatty acids, and hydroxy impurities. Using all-atom molecular dynamics simulations, we investigated interactions between NRCs and terminal groups. From an analysis of the equilibrium properties; e.g., end-to-end distance, radius of gyration, rotational relaxation time, stress-stress autocorrelation, diffusion coefficient, radial distribution function, we found that NRCs preferentially interact with terminal groups, signifi-

cantly slowing chain dynamics and promoting stable NRC – terminal clusters. Fatty acids and hydroxy impurities form star-like physical junctions mediated by hydrogen bonds, restricting polymer mobility. These findings provide molecular-level insight into how NRCs enhance the mechanical strength of NR through intermolecular interactions and junction formation (M. Dixit and T. Taniguchi, ACS Applied Engineering Materials **3**, 337, 2025; M. Dixit and T. Taniguchi, ACS Polymer Au **4**, 273, 2024).

5.2 ML for Stokes Flows

Our proposed probabilistic inference method relies on a physics-informed GP regression that directly incorporates the Stokes and continuity equations into the GP kernel. The method works equally well on 2D or 3D systems. However, exact regression requires us to evaluate large matrix(inverse)-vector products. Using a standard Cholesky-based decomposition limits the calculations to relatively small systems ($\simeq 10^3$ training points), making it unusable for experimental 3D flow data. To overcome this problem, the recommended approach is to use a Black-Box-Matrix-Matrix (BBMM) method (as pioneered by GPyTorch). However, we were unable to achieve the desired speedup/performance (most likely due to the non-sparse nature of the resulting Kernel matrices) and have yet to find a suitable replacement method.

5.3 ML for Optimal Control

We have extended our reinforcement learning method to consider the task of transporting cargo through a non-uniform flow field using only local hydrodynamic signals as a guide. For this, we continue to use Deep Q-Learning (via PyTorch) to learn the optimal action response, coupled to direct numerical simulations (using the KAPSEL colloidal simulator) to resolve the fluid/particle

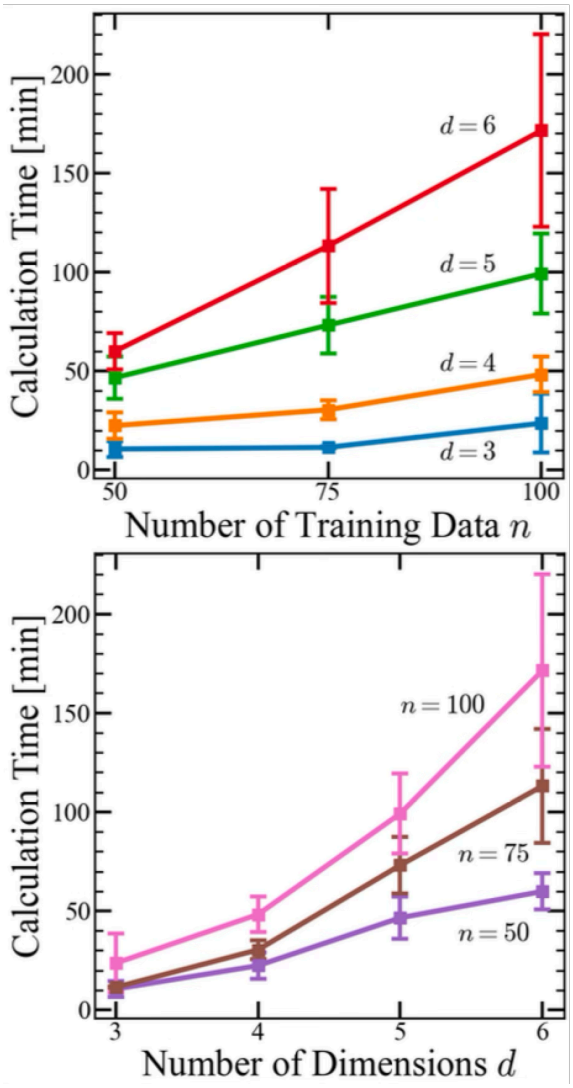

Fig. 2 Calculation time as a function of (top) number of training data sets and (b) number of components or dimensions d; running on a single NVIDIA A100 GPU (Wisteria). Y. Ueno, Master Thesis, Kyoto University (2025).

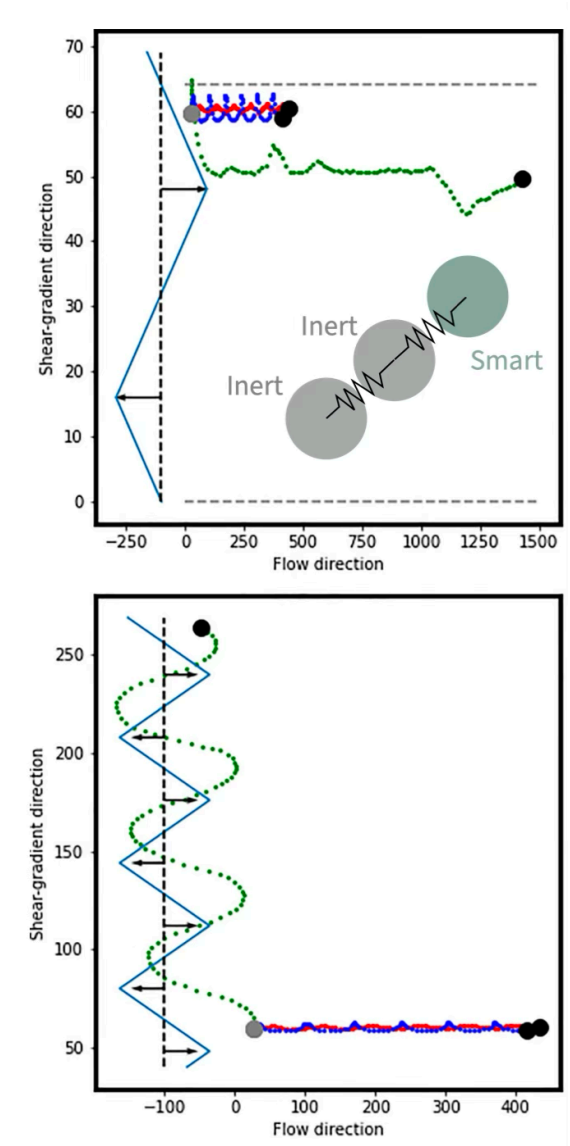

Fig. 3 Simulation trajectory for a "smart" (green) and "naive" (blue) cargo-towing swimmer tasked with moving in the (left) shear-gradient and (right) shear-flow directions; together with an inert chain (red). Adapted from K. Sankaewtong et at., Phys. Rev. Res. **6**, 033305, 2024 (CC BY 4.0).

hydrodynamic coupling. For the active cargo-towing particles we use the canonical squirmer model, while the cargo itself consists of inert tracer/colloidal particles attached to the swimmer by harmonic springs (see Fig.3). The "smart" swimmer particle is endowed with sensors on its surface, which allow it to measure the local hydrodynamic forces. This "smart" swimmer is able to tune its swimming velocity, as well as it's swimming direction, in response to these hydrodynamic signals. We successfully trained such a cargo-towing swimmer to navigate a zig-zag shear flow; it is able to effectively move in both the shear-gradient and shear-flow directions. Performance is dramatically improved compared to "naive" cargo-towing swimmers, as well as inert chains (Fig. 3). Surprisingly, we found that the cargo-towing swimmers outperform single swimmers when trained to swim in the shear-flow direction. We also find that pullers outperform neutral and pusher-type swimmers, regardless of the cargo-load (K. Sankaewtong et al., Phys. Rev. Res. **6**, 033305, 2024).

In addition, we have also considered the optimal control/ game theory problems that explains the behavior of rational/selfish agents socially distancing during an epidemic. To start, we have considered a simple compartmental disease model, e.g., the Susceptible-Infected-Recovered (SIR) model, within a mean-field description in which all individuals are identical. The disease infectivity is used as a proxy for individual behavior, which serves as the control variable that individuals can tune to maximize (minimize) their total utility (loss). Given the form of the utility function, it is straightforward to derive the appropriate Euler-Lagrange equa-

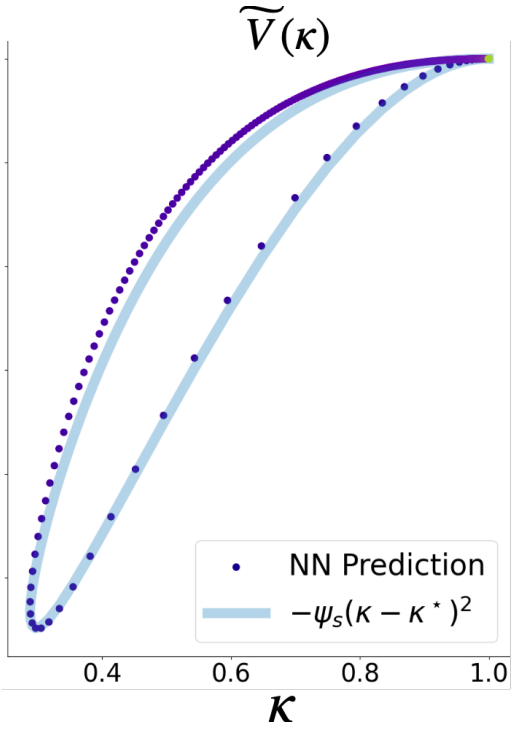

Fig. 4 Effective utility $\widetilde{V}$ as a function of the behavior κ, for individuals socially distancing during an epidemic. The learned NN prediction is compared to the true function that determines the behavior.

tions for this constrained optimization problem. However, solving the inverse problem is incredibly challenging. We have identified an analytic solutions to the direct problem (Schnyder et al., PNAS **122**, e2409362122, 2025), which has allowed us to gain a much deeper understanding of the model. We have also investigated the role of empathy in the disease outcome, and found that even a small degree of empathy can have large effects on the population and disease dynamics. Mainly, we have developed a physics-informed Neural-Network (PINN) that incorporates both the physical constraints (i.e., the Lagrange/Hamiltonian equations of motion), as well as the game-theoretic constraints, into the structure of the networks. Using this approach, we have successfully inferred hidden utilities (Fig. 4) from behavior data. Unfortunately, this inference requires information on the Lagrange multipliers used to constrain the disease dynamics. We have investigated how to accelerate the training, as well as how to completely eliminate the unwanted (experimentally unknown) degrees of freedom from the learning.

6 Self-review of Current Progress and Future Prospects

For theme (A), we focused on understanding the coupling between micro/macro properties of polymer systems. This included tackling the inverse problem of inferring microscopic data given macroscopic information (e.g., molecular weight distributions from stress relaxation curves), as well as improving our knowledge of the microscopic origins of the enhanced properties of natural rubber. Our efforts to improve the learning of the constitutive relations are still ongoing, in particular, we must still improve the reliability of the ML predictions, as well as the performance of our MSS code (e.g., by efficiently parallelizing the calculation of polymer chain entanglements on GPUs). Likewise, for theme (B) we are still in the process of improving the efficiency and scalability of the inference beyond $10^3 - 10^4$ training points. For theme (C), we extended our Reinforcement Learning method to consider "smart" cargo-towing swimmers. We also improved upon our optimal control/game-

theory model of social distancing during an epidemic, though our solution of the inverse problem still requires some knowledge of the (hidden) Lagrange multipliers.

Our future prospects are clear, as we have shown that the methods we developed work as intended. From this point, we will work on optimizing them. For theme (A) this involves improving the predictive capabilities of the learned constitutive relations, as well as developing a fully-differentiable flow-solver. For theme (B) we need to find a more efficient method to perform the matrix(inverse)-vector products for our custom physics-informed Gaussian Process kernels. For theme (C), we will focus on the (inverse) optimal control problem, to develop physics-informed neural networks that can be trained solely on observable behavioral data.