CONTROLLING DYNAMIC SPATIAL LIGHT MODULATORS USING EQUIVARIANT NEURAL NETWORKS

Sumukh Vasisht Shankar
Khoury College of Computer Sciences
Northeastern University
Boston, MA 02115, USA
vasishtshankar.s@northeastern.edu

Darrel D'Souza
Mechanical and Aerospace Engineering
Rutgers University
Piscataway, NJ 08854, USA
ddd102@scarletmail.rutgers.edu

Jonathan P. Singer
Mechanical and Aerospace Engineering
Rutgers University
Piscataway, NJ 08854, USA
jonathan.singer@rutgers.edu

Robin Walters
Khoury College of Computer Sciences
Northeastern University
Boston, MA 02115, USA
r.walters@northeastern.edu

ABSTRACT

Spatial Light Modulators (SLMs) are devices that can modulate the amplitude or the phase of a beam of light. These devices are used in applications such as beam front aberration and microscopic manipulation with optical tweezers. Here, we study the problem of learning to modulate light in a new type of temperature-controlled SLM. These SLMs are panels that use a thin viscous film in which shallow wave patterns can be induced by varying the temperature of the panel. This method can be used for modulating light such as high-power lasers. The problem here is to learn which input temperature signal is necessary in order to induce a given pattern in the reflected light. We propose a deep E(2)-equivariant model to learn this relationship. We generate a synthetic dataset consisting of temperature signals and corresponding light patterns by simulating the thin film lubrication equation that governs the phenomenon of thermocapillary dewetting. We use this dataset to train our networks. We demonstrate the advantage of using equivariant neural networks over convolutional neural networks in order to learn the mapping.

1 INTRODUCTION

Deep learning models have proven helpful in analyzing and understanding complex material structures and properties. The most difficult such tasks involve predicting the changes in material properties over time. However, recent work has shown deep learning can be effective in modeling many different dynamical systems and can provide a method for solving difficult differential equations (Wang et al. (2020a), Wang et al. (2020b)). Data-driven methods have an advantage over traditional computational methods in that they are several orders of magnitude faster, can operate in partially observed systems, and are robust to noise. Perhaps most importantly, a differentiable model of system dynamics can be optimized for control or design applications.

In this paper, we propose to use deep learning to predict the temperature signal required to induce a given pattern of reflected light in a thermocapillary dewetting-based dynamic spatial light modulator (Kovacevich et al. (2021)). Solving the thin film lubrication equation that governs the phenomenon of thermocapillary dewetting is computationally expensive. We demonstrate that deep models can predict the signal quickly and accurately.

For this particular problem, traditional neural networks need a large dataset to learn the relationship between the height profile and the temperature profile required to induce it. However, in experimental applications like this, real data is scarce. By incorporating the symmetries of the lubrication equation into the neural network, we build equivariant networks which are able to learn with fewer samples. Another advantage is that such networks have better generalization ability and physical accuracy than unconstrained models Wang et al. (2020b). Specifically, we incorporate rotational and translational symmetry into the model using $E(2)$-equivariant networks Weiler & Cesa (2019).

Our contributions include:

- We study the problem of controlling the input temperature profile required to induce a given height profile in a thermocapillary dewetting-based dynamic spatial light modulator.
- We generate a synthetic dataset by simulating the thin film lubrication equation for two-dimensional inputs.
- We use equivariant neural networks to enforce rotational symmetry from the physical system and demonstrate improved accuracy.

2 BACKGROUND

2.1 SLMS AND THERMOCAPILLARY DEWETTING

Spatial light modulators. Dynamic SLMs are devices that can precisely modulate a beam of light by tuning either the phase or the intensity of an array of pixels in parallel. SLMs are used in optics to control the behavior of light and are used for a number of applications like projectors, laser beam shaping, beam front aberration, etc. They can also be used to manipulate the properties of microscopic particles and materials by applying different light patterns to them.

Conventional dynamic SLMs are normally incompatible with high-powered sources like lasers. To address this issue, Kovacevich et al. (2021) present the usage of thermocapillary dewetting to dynamically control the thickness of a thin viscous, reflowable film. This film, when used with an SLM, results in the incident beam of light being phase-shifted and focused based on the height map (pattern) of the film.

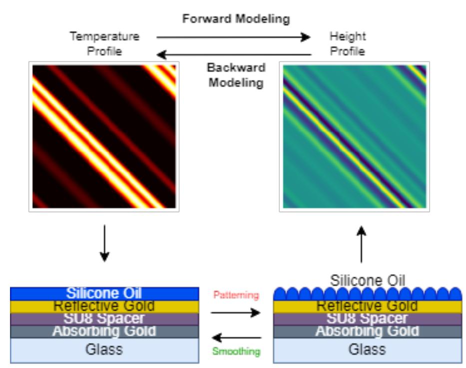

Figure 1: Schematic of multilayer SLM showing forward and backward modeling tasks.

Thermocapillary dewetting is a phenomenon where a temperature gradient acting on a thin liquid film creates a thermocapillary force that leads to the deformation of the liquid film (Bénard, 1900; C.E., 1855; Darhuber & Troian, 2005; Bénard, Henri, 1901). The driving force of thermocapillary dewetting is the thermocapillary shear $\tau = \beta \boldsymbol{\nabla} T$. The evolution of the height profile of the film is described by the thin film lubrication equation,

$$\frac{\mathrm{d}h}{\mathrm{d}t} = -\boldsymbol{\nabla}.\left(\frac{h^2\beta\boldsymbol{\nabla}T}{2\mu} + \frac{h^3}{3\mu}\left(\gamma\boldsymbol{\nabla}^2 h + \frac{\mathrm{d}V}{\mathrm{d}h}\right)\right) \tag{1}$$

where μ is the fluid viscosity, h is the film thickness, and V encapsulates surface interactions (Becker et al. (2003)). When the thin film is heated, it becomes thin in the heated areas and becomes thick around it, which leads to the formation of trench and ridge structures. Kovacevich et al. (2021) explains the real-world experimental setup along with images.

2.2 EQUIVARIANT NETWORKS

We incorporate equivariance into two architectures - ResNet (He et al. (2016)) and U-Net (Ronneberger et al. (2015)). The architectures of the equivariant models are adapted from Wang et al. (2020b). The models from Wang et al. (2020b) were built for fluid flow vector fields, and were applied autoregressively for dynamic prediction whereas here we consider scalar fields. We incorporate rotational and translational symmetry into our models by enforcing equivariance.

A function f is said to be equivariant if when the input of f is transformed by a symmetry group element g, the output of the function is also transformed by the same symmetry group element g:

$$f(g.x) = g.f(x). \tag{2}$$

In this paper, both the forward modeling task (predicting the height profile for a given temperature profile) and the backward modeling task (predicting the temperature profile for a given height profile) are rotationally and translation equivariant due to invariances in the underlying equations Jafari & Tanhaeivash (2021).

Rotational equivariance. To incorporate rotational symmetry from the underlying physical system into our networks, we use the results from Cohen & Welling (2016a) and Cohen & Welling (2016b) that describe how the convolution operation can be constrained by addition weight sharing to enforce rotational equivariance on each layer of CNN. The `RotEq` models are built using the E(2)-CNN architecture by Weiler & Cesa (2019). Figure 2 shows an illustration of an output being transformed by the symmetry that acts upon an input. For this paper, we use the discrete cyclic rotation group $G = C_4$.

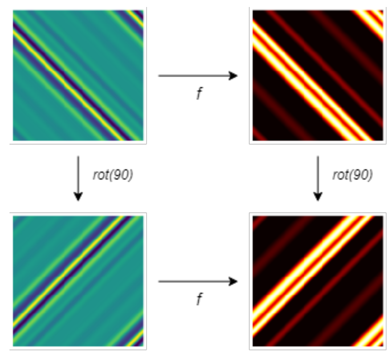

Figure 2: Rotational equivariance of the backward task mapping temperature to height with respect to $g = \mathrm{rot}(\pi/2)$

3 EXPERIMENTS

We compare rotationally equivariant versions of ResNet and U-net called `RotEq-ResNet` and `RotEq-Unet` (Wang et al. (2020b)) with convolutional neural network architectures (`CNN`, `U-Net`, and `ResNet`). We test our models on the synthetic dataset generated by simulating Equation 1.

Evaluation Metrics. For accuracy, we use Root Mean Square Error (RMSE) between the predictions and the ground truth over all pixels.

Experimental Setup. Unless otherwise mentioned, the models are trained for 60 epochs. We use a training set of size 300, validation set of size 70 and test set of size 100.

Model Architecture. We implement a `CNN` as our baseline model. We use 4 hidden layers with a kernel size of 2 and stride of 1. `ResNet` He et al. (2016) and `U-Net` (Ronneberger et al. (2015)) are effective deep convolutional architectures and are well-suited for our tasks. Hence, we implement these two architectures with rotational symmetry, which we name `RotEq-ResNet` and `RotEq-Unet`. We use a kernel size of 3. All models predict the input temperature signal required to induce a given pattern of light in the SLM as the backward task (backward modeling) and predict the height profile for a given temperature profile as the forward task (forward modeling) as shown in Figure 1.

3.1 DATA SIMULATION AND DESCRIPTION

The experiments necessary for real-world data are expensive. For our proof of concept study, we simulate the thin film lubrication equation (TFLE) shown in equation 1, which is the governing equation for thermocapillary dewetting, using finite-difference methods to generate a synthetic dataset.

We simulate the TFLE for an input 200×200 array representing the *temperature profile*. The output of the simulation is a 200×200 array representing the *height profile*. The temperature profile consists of peaks of temperatures in random patterns across the thin film after a sufficient amount of time such that the profile has stabilized. This results in the ridge and trench pattern in the resulting height profile. Temperatures are in Celcius and heights are in nanometers. We resize it to 192×192 for our experiments.

3.2 EXPERIMENTS ON SIMULATED DATA

3.2.1 FORWARD MODELING

While our primary goal is the backward task of finding the temperature field to produce a desired height profile, we also demonstrate our model can learn the forward task to show the model can capture system dynamics. All the models are trained with temperature profiles as the input and height profiles as the output.

		Forward Modeling		Backward Modeling	
Config	Parameters	Train RMSE	Test RMSE	Train RMSE	Test RMSE
CNN	3.5M	0.29 ± 0.02	0.23 ± 0.03	0.35 ± 0.02	0.32 ± 0.03
U-Net	3.6M	0.28 ± 0.02	0.17 ± 0.03	0.27 ± 0.03	0.31 ± 0.03
ResNet	3.8M	0.23 ± 0.03	0.12 ± 0.02	0.26 ± 0.03	0.18 ± 0.04
RotEq U-Net	3.6M	0.24 ± 0.02	0.15 ± 0.03	0.18 ± 0.01	0.16 ± 0.02
RotEq ResNet	3.9M	$\mathbf{0.19 \pm 0.02}$	$\mathbf{0.11 \pm 0.02}$	$\mathbf{0.16 \pm 0.02}$	$\mathbf{0.11 \pm 0.03}$

Table 1: RMSE of `CNN`, `ResNet`, `U-Net` and the `RotEq-ResNet`, `RotEq-UNet` trained and tested on the simulated data for forward and backward modeling.

Prediction Performance. Table 1 shows the RMSE of `CNN`, `ResNet`, `U-Net` and the `RotEq-ResNet`, `RotEq-UNet` on the train and test sets. We report mean errors over three random runs. Both equivariant models perform better than the non-equivariant baseline on RMSE. `RotEq-ResNet` achieves the lowest RMSE. The baseline models have a higher error in capturing the peaks in the height profile. `ResNet` and `U-Net` have similar RMSE while `RotEq-ResNet` performs better than `RotEq-UNet` and generalizes better. Refer to Appendix A for the ground truth and the predicted height profiles from forward modeling as well as the ground truth and prediction by the models for a random row of the height profile.

3.2.2 BACKWARD MODELING

Our primary goal in this paper is the backward task of finding the temperature profile required to produce a desired height profile. All the models are trained with height profiles as the input and temperature profiles as the output.

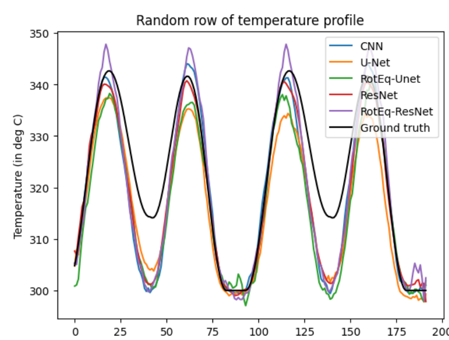

Figure 3: Random row of temperature profile (backward modeling)

Prediction Performance. Table 1 shows the RMSE of `CNN`, `ResNet`, `U-Net` and the `RotEq-ResNet`, `RotEq-UNet` on the train and test sets. We report mean errors over three random runs. Both equivariant models perform better than the non-equivariant baseline on RMSE. `RotEq-ResNet` achieves the lowest RMSE. Figure 3 shows the ground truth and prediction by the models for a random row of the temperature profile. We can see that the baseline models do not capture the peaks well compared to the equivariant models. `ResNet` and `U-Net` have similar RMSE while `RotEq-ResNet` performs better than `RotEq-UNet`. Figure 4 shows the ground truth and the predicted temperature profiles by all the models for a given height profile. All the models either have a similar or better RMSE for the backward modeling compared to the forward modeling.

4 CONCLUSION AND FUTURE WORK

We develop methods to predict the temperature profile required to induce a given height profile in a thermocapillary dewetting-based dynamic spatial light modulator. We incorporate rotational symmetry by using equivariant neural networks. We show that the equivariant neural networks out-

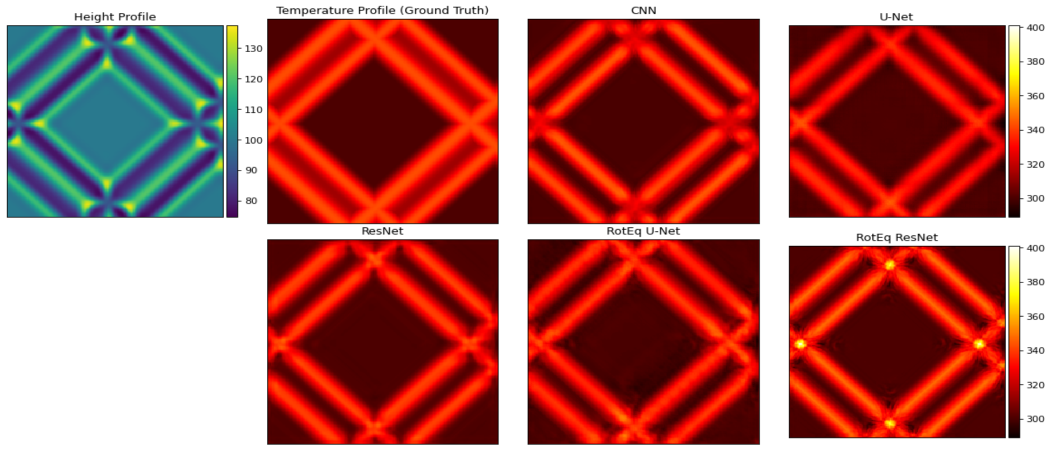

Figure 4: **Height profile**, ground truth and predicted **temperature profiles** by `CNN`, `U-Net`, `ResNet`, `RotEq-Unet`, `RotEq-ResNet`

perform the traditional architectures experimentally in terms of accuracy because they incorporate physical symmetries.

Future work includes carrying out simulations to generate a more diversified synthetic dataset with different patterns for the temperature and height profiles. We will be incorporating other symmetries into the equivariant models. Lastly, we plan to test our models on real-world data.

5 Acknowledgement

We thank Jonathan P. Singer and Darrel D'Souza for helpful discussions and for providing us with the initial MATLAB code for data simulations. We also thank the reviewers for taking the necessary time and effort to review the manuscript. We sincerely appreciate all your valuable comments and suggestions, which helped us in improving the quality of the manuscript. This work is supported by the Khoury Research Apprenticeship Program and the Geometric Learning Lab of Northeastern University.

References

Jürgen Becker, Günther Grün, Ralf Seemann, Hubert Mantz, Karin Jacobs, Klaus R. Mecke, and Ralf Blossey. Complex dewetting scenarios captured by thin-film models. *Nature Materials*, 2 (1):59–63, Jan 2003. ISSN 1476-4660. doi: 10.1038/nmat788. URL `https://doi.org/10.1038/nmat788`.

Henri Bénard. Les tourbillons cellulaires dans une nappe liquide. *Revue Gen. Sci. Pure Appl.*, 11: 1261–1271, 1900.

Bénard, Henri. Les tourbillons cellulaires dans une nappe liquide. - méthodes optiques d'observation et d'enregistrement. *J. Phys. Theor. Appl.*, 10(1):254–266, 1901. doi: 10.1051/jphystap:0190100100025400. URL `https://doi.org/10.1051/jphystap:0190100100025400`.

James Thomson A.M. C.E. Xlii. on certain curious motions observable at the surfaces of wine and other alcoholic liquors. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 10(67):330–333, 1855. doi: 10.1080/14786445508641982. URL `https://doi.org/10.1080/14786445508641982`.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016a.

Taco S Cohen and Max Welling. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016b.

Anton A. Darhuber and Sandra M. Troian. Principles of microfluidic actuation by modulation of surface stresses. *Annual Review of Fluid Mechanics*, 37(1):425–455, 2005. doi: 10.1146/annurev.fluid.36.050802.122052. URL `https://doi.org/10.1146/annurev.fluid.36.050802.122052`.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

M Jafari and A Tanhaeivash. Symmetry group analysis and similarity reductions of the thin film equation. *Journal of Linear and Topological Algebra*, 10(04):287–293, 2021.

Dylan A Kovacevich, Tianxing Ma, Arielle R Gamboa, Michael P Nitzsche, Valeria Saro-Cortes, Emily Davis, and Jonathan P Singer. Thermocapillary dewetting-based dynamic spatial light modulator. *Optics Letters*, 46(15):3721–3724, 2021.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pp. 234–241. Springer, 2015.

Rui Wang, Karthik Kashinath, Mustafa Mustafa, Adrian Albert, and Rose Yu. Towards physics-informed deep learning for turbulent flow prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1457–1466, 2020a.

Rui Wang, Robin Walters, and Rose Yu. Incorporating symmetry into deep dynamics models for improved generalization. *arXiv preprint arXiv:2002.03061*, 2020b.

Maurice Weiler and Gabriele Cesa. General e (2)-equivariant steerable cnns. *Advances in Neural Information Processing Systems*, 32, 2019.

A PREDICTIONS FROM FORWARD MODELING

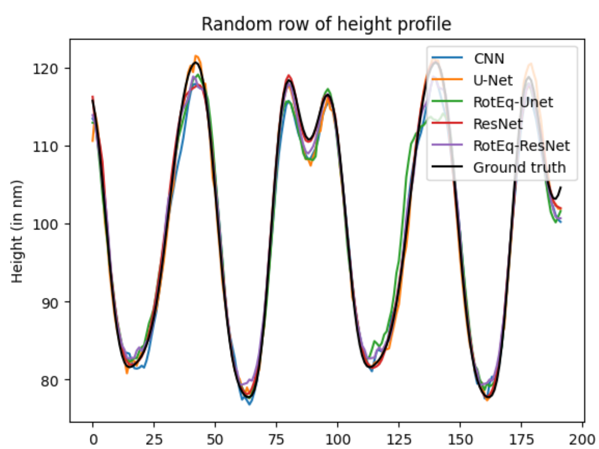

Figure 5: Random row of height profile (forward modeling)

Figure 6 shows the ground truth and predicted height profiles by all the models for a given temperature profile.

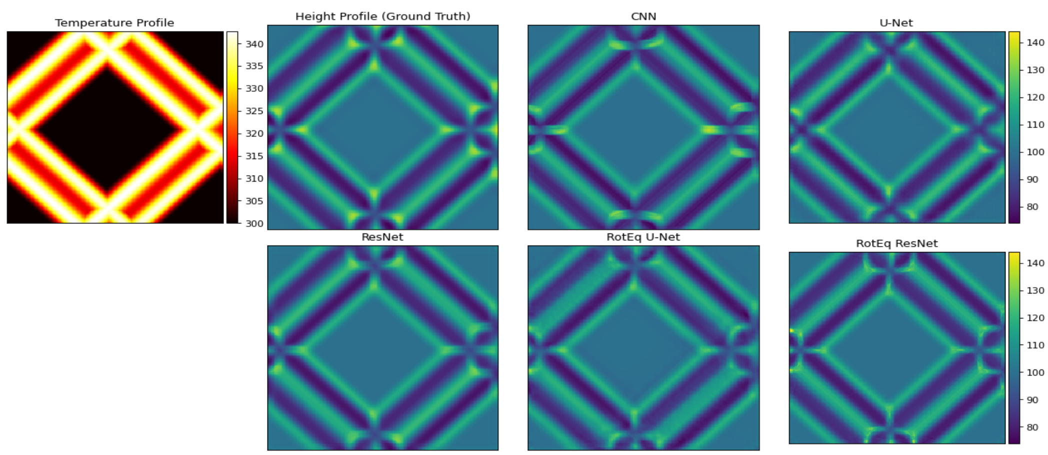

Figure 6: **Temperature profile**, ground truth and predicted **height profiles** by `CNN`, `U-Net`, `ResNet`, `RotEq-Unet`, `RotEq-ResNet`