2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Learning Deposition Policies for Fused Multi-Material 3D Printing

Kang Liao*,1,2, Thibault Tricard*,1, Michal Piovarči[3], Hans-Peter Seidel[1], and Vahid Babaei[1]
[1]Max Planck Institute for Informatics [2]Beijing Jiaotong University [3]Institute of Science and Technology Austria
Email: {kangliao, ttricard, hpseidel, vbabaei}@mpi-inf.mpg.de

***Abstract*— 3D printing based on continuous deposition of materials, such as filament-based 3D printing, has seen widespread adoption thanks to its versatility in working with a wide range of materials. An important shortcoming of this type of technology is its limited multi-material capabilities. While there are simple hardware designs that enable multi-material printing in principle, the required software is heavily underdeveloped. A typical hardware design *fuses* together individual materials fed into a single chamber from multiple inlets before they are deposited. This design, however, introduces a time delay between the intended material mixture and its actual deposition. In this work, inspired by diverse path planning research in robotics, we show that this mechanical challenge can be addressed via improved printer control. We propose to formulate the search for optimal multi-material printing policies in a reinforcement learning setup. We put forward a simple numerical deposition model that takes into account the non-linear material mixing and delayed material deposition. To validate our system we focus on color fabrication, a problem known for its strict requirements for varying material mixtures at a high spatial frequency. We demonstrate that our learned control policy outperforms state-of-the-art hand-crafted algorithms.**

I. Introduction

An increasingly large number of additive manufacturing (AM), also known as 3D printing, technologies rely on continuous deposition where the deposited material follows a continuous path. Such systems are highly preferred over the alternative *raster* deposition, such as inkjet printing, thanks to their capability of printing with a wide range of materials, from electronics to biomaterials [1], [2]. A popular continuous-deposition printing hardware, considered as a representative of this technology in our work, is Fused Deposition Modeling (FDM). FDM works by extruding a plastic filament heated to its melting temperature which solidifies quickly after deposition. While simple in concept, FDM fabrication has seen widespread adoption by a variety of users, including hobbyists, artists, educators, and researchers. Even though FDM printers can process a wide range of filaments, most commercial devices offer at most two separate nozzles: one for the model material and one for the support material. The constraints of a two-nozzle system, *e.g.*, registration of the nozzles, separated paths of the nozzles, and a limited number of materials hinder achieving functionally graded materials that are critical in many applications, such as metamaterial design [3], color fabrication [4], or self-assembling [5].

The quest for multi-material, continuous deposition has motivated various hardware improvements [6], [7], [8]. The most successful to date builds on a *fusion* strategy, much related to the initial FDM concept (but not necessary limited to FDM). In this design, multiple materials are fed into an expanded heated chamber simultaneously as shown in Figure 1 (c). As they melt, they fuse into each other forming a composite material. The proportion of different filaments inside the composite is governed by the *feed rate* of the primary filaments. Such a system can generalize to an array of nozzles and even be adapted to varying viscosities [9]. A disadvantage of this design is that the expanded chamber leads to a considerable delay when switching from a composite with a desirable mixture ratio to a different one. This manifests as a region where the materials smoothly vary between the two target mixtures, limiting the achievable spatial resolution.

To tackle this problem we take inspiration from the interconnection between robotics and additive manufacturing. Robotic systems are extensively used for material deposition [10], [11] where their large range of motion allows for building support free structures [12] or even deposit material in air [13]. The key in deploying deposition systems with such a high degree of freedom in movement lies in careful robotic control [14], [15], [16]. Similarly, in this work, we seek to improve the multi-material continuous deposition by finding optimal printing policies. To this end, we propose a numerical approximation of the FDM printing process based on empirical observations. We integrate our numerical model into a reinforcement learning framework. To accommodate high-frequency changes we adapt the learning process for fused deposition. These improvements enable us to discover optimal printing strategies by experimenting in a simulated environment. To validate our approach, we present a case-study in color fabrication, notorious for its requirements for reproducing high-frequency changes present in the input content (*i.e.*, images). We demonstrate that our control policies are capable of reproducing such a level of detail. Moreover, we show that our control strategies significantly outperform the current state-of-the-art in color fused fabrication. We believe that our framework is the first step toward practical, high-quality multi-material deposition and can be expanded to generate policies for physical hardware. Our code is available at https://github.com/ThibaultTricard/material-fusion.

*The first two authors contributed equally. This work is graciously supported by FWF Lise Meitner (Grant M 3319). Kang Liao sincerely thank Emiliano Luci, Chunyu Lin, and Yao Zhao for their huge support.

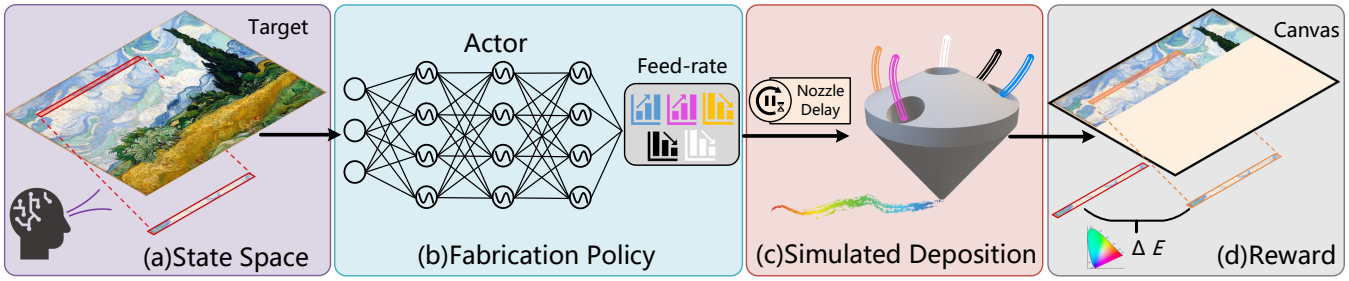

Fig. 1. Overview of our method. We propose to formulate the search for optimal multi-material printing policies in a reinforcement learning setup. First, the agent observes the environment through a 1D sliding window, constructing the state space from the target image (a). Then, our fabrication policy (b) takes the state as input and predicts the feed-rate in regards to different filament materials (in this work we consider the CMYKW setup with five materials). Besides, a numerical deposition model (c) is designed to take into account the non-linear material mixing and delayed material change. Finally, the agent gains the reward in CIELab color space by simulating the resulting pattern on the canvas (d).

II. Related Work

Creating objects made of multiple materials at a high spatial resolution, leading to functionally graded materials, is one of the key competitive advantages of additive manufacturing (AM) over traditional manufacturing methods. Multi-material printing, especially at a high resolution, is typically associated with a single AM technology: inkjet printing. Given the inherent limitation of inkjet 3D printing, we are witnessing hardware improvements that bring multi-material printing to continuous-deposition systems, such as FDM printers [8], [6], [7], [17], [18], [19]. The most common approach is to utilize multiple-in-one-out mixing nozzle as illustrated in Figure 1 (c). Such a nozzle fuses the *base* materials at the output providing excellent mixture quality. Such a high mechanical performance motivated researchers to bring this kind of nozzle to other systems such as Direct-Ink Writing and to generalize the system to multi-nozzle output [9]. A natural drawback of mixing materials within the nozzle is the increased delay required for material swaps. In the case of FDM printers discharging the in-nozzle material can take up to 500 milliseconds. Such a delay causes visible artifacts that affect the functional behavior of printed objects.

In recent years we have seen significant improvements in machine learning based control strategies [20], [21]. Most notably Reinforcement Learning (RL) has been shown to be a powerful tool for discovering control policies for robotic arms [22], [23], hands [24], [25], [26], walker robots [27], [28], [29], and their simulated counterparts [30], [31]. Inspired by this success in robotics and in a transition from hand-designed control algorithms, researchers have started to explore RL for controlling additive manufacturing hardware. The method has been successfully applied to stabilization of filament feed rate [14], tool-path optimization [32], thermal control [15], or control of arc welding [33]. Most notably, Piovarci et al. [16] presented a first closed-loop control policy, discovered entirely in simulation, for printing uniformly covered layers in the presence of considerable material spreading. In this work we assume the material spreading is negligible, plausible for FDM filaments. Having a completely different goal in mind, we take inspiration from this work to control multi-material fusion-based printing. We seek an optimal control strategy that would enable reproducing features beyond the originally assumed capabilities of the physical hardware.

A notable application for multi-material 3D printing is color reproduction [34], [35], [36], [4]. To create a wide range of colors the printers rely on a set of material primaries. The typical primaries are Cyan (C), Magenta (M), Yellow (Y), Black (K), and White (W). Typically, it is assumed that the printer can switch between material instantaneously [37] (the case for the dominant inkjet 3D printing). The instantaneous materials swaps are specially important for images with high spatial frequency material contents. Because of these challenges, we opted for color fabrication as a stress test of the capabilities of our control system for multi-material fusion 3D printers.

III. Numerical Model of Fusion Printer

We propose an empirical model of the deposition process. During the fabrication, a set of N materials is pushed towards the nozzle. By varying the feed rates of individual materials one can achieve various material mixtures. Typically, to maintain a consistent deposition the sum of the feed rates is constrained to a constant value. As the materials are pushed towards the nozzle they melt and fuse together. We mimic this behavior in our numerical model:

$$m = \sum_{i}^{N} w_i \mathscr{P}_i, \tag{1}$$

where m is a new mixed material, w_i are the mixing ratios subject to $\sum_i^N w_i = 1$, and $\mathscr{P}_i$ are material parameters. In practice the mixing weights w_i are scaled to a feed rate appropriate for the physical machine.

The fused material mixture accumulates at the nozzle. The accumulated material needs to be discharged before a change in material mixture can be visible. We model this hysteresis using an exponential decay of material present at the nozzle:

$$M_t = (1-\xi)^t m_{-n} \sum_{i=-n+1}^{t-n} m_i \cdot \xi(1-\xi)^{t-n-i}, \tag{2}$$

where M_t is the material mixture at location t, ξ denotes the changing stride regarding to the current ratios, and n represent the delay between nozzle input and output change, with all m_i where $i < 0$ set to a default value. Intuitively,

our model emulates the mixing of material inside the nozzle by relying on an exponential averaging of the previous input feeding rates to which we add a delay. We empirically set the ξ (mixing ratio) and n (time delay) to 0.6 and 5, respectively to match the observed color transition on a typical nozzle.

Having the material ratios at each location t, we now seek to compute the color of the fused composite. The color of a mixture is derived from its reflected spectrum $R(\lambda)$ at each visible wavelength λ. Similar to many other material properties, the color of a mixture has a non-linear relationship with the color of the base components that make up the mixture. The gold-standard color prediction model for mixing paints or pigments is the Kubelka-Munk (K-M) model [38]. Rooted in material properties, the model represents each primary material as wavelength-dependent absorption $K(\lambda)$ and scattering $S(\lambda)$ coefficients. To predict a new material we first mix these coefficients linearly:

$$K_{mix}(\lambda) = \sum_{i}^{N} M_{t_i} K_i(\lambda), \qquad S_{mix}(\lambda) = \sum_{i}^{N} M_{t_i} S_i(\lambda), \qquad (3)$$

where M_{t_i} are the mixing ratios of our materials. Next, to compute the visible spectra, Kubelka and Munk derived the following analytical relationship:

$$R_{mix}(\lambda) = 1 + \alpha(\lambda) - \sqrt{(\alpha(\lambda))^2 + 2\alpha(\lambda)}, \qquad (4)$$

where $R_{mix}(\lambda)$ is the visible spectra, and $\alpha(\lambda) = \frac{K_{mix}(\lambda)}{S_{mix}(\lambda)}$. Note that we here don't take into account all wavelengths and use *broadband* RGB channels, and thus obtain 3D RGB colors R^{rgb}_{mix} in Equation 4. In practice we rely on the implementation in [39] which augments the K-M 'latent' space with additional imaginary bases to cover the whole RGB color space.

IV. Learning Framework

Our main goal is to learn a deposition policy that reproduces a continuous (but not necessarily smooth) multi-material density field. We focus on color reproduction of a 2D image using a simulated multi-filament FDM printer (Section III). The input to the policy is a target image. The output is a set of feed rate instructions along a predefined printing path to replicate the input as closely as possible. The main challenge stems from the hardware limitations that prevent the deposition of materials with precise mixing ratios at a high spatial resolution.

To find a suitable printing policy we formulate the search in a reinforcement learning setup. More precisely, we cast the policy search into a Markov Decision Process $(\mathscr{S}, \mathscr{A}, \mathscr{P}, \mathscr{R})$, where $\mathscr{S}$ is a set of valid inputs to the policy, $\mathscr{A}$ is the set of actions available to the policy, $\mathscr{P}$ is a transition function that models the printing process, and $\mathscr{R}$ is the reward function that evaluates the quality of the deposition. In this section, we describe how to apply this framework for colored FDM deposition. The overview of our method is shown in Figure 1.

A. State Space

The state space defines the input to the control policy. A natural selection would be to use the complete input image. However, the space of all possible input images is too large to be explored efficiently. To facilitate the exploration we can leverage the mixing properties of the physical nozzles. Mixing different material ratios takes a predefined amount of time. This naturally discounts the influence of the current printer state on future states as the material in the chamber will be gradually dispensed. Therefore, we suggest to model the state space as a 1D sliding window centered at the current nozzle location as shown in Figure 1 (a). We allow the control policy to observe the RGB color of L pixels in the future to adjust deposition. We empirically set L to 15 for a sufficient receptive field while eliminating other irrelevant context.

B. Action Space

A typical 3D printer operates by receiving time-based instructions on nozzle position (printing path) and process parameters (feed rate in our case). We fix the path and focus on learning to output appropriate process parameters. While any space filling curve could be used in our setting, we choose to work with a rectilinear (zigzag) path common for FDM printers. Thus, we tune the feed rate of each input filament to control the resulting color mixture. In this work we consider a *full-color* printer with CMYKW filaments.

C. Transition Function

The transition function in our problem represents the printing process. We model it using our customized simulator from Section III. The simulator takes the current state of the system (*i.e.*, the printing bed, and the leftover material in the nozzle) and the action (*i.e.*, material feed rates) as input and outputs the system stateafter simulating the deposition.

D. Reward Function

Our reward function measures the difference between the target color and the deposited color at each location. Although an Euclidean distance in RGB could be used, we opt for the CIELab color space [40] because of its superior perceptual uniformity. The CIELab colors can be derived directly from RGB colors using a set of analytical transforms [41]. The reward function is therefore the most well-known *color difference formula* (ΔE):

$$\mathscr{R} := \Delta E = -\|R^{Lab}_{mix} - R^{Lab}_{target}\|_2. \qquad (5)$$

V. Multi-material Fusion Fabrication Policies

To discover a printing policy we adopt the actor-critic framework [42]. We model both the actor and the critic using convolutional neural networks (CNNs). Our networks have 3 layers with 1×3 convolutional filters, which progressively extract the high-level features while keeping the spatial alignment among different channels. The layers have 32, 64, and 128 filter kernels respectively. The final CNN output is

Fig. 2. We show the reproduced results by our method and the targets (in gamut-mapped format). The zoom-in views of different samples are exhibited upon the results.

flatten and passed through a fully connected layer of 512 neurons.

A key feature of colored fabrication is the need for high-frequency changes. Unfortunately, the state-of-the-art in CNN relies on rectified linear units (ReLU) activation. ReLUs are incapable of representing high-frequency information due to their incapability to model the derivatives of input signals [43]. To ameliorate this problem, we propose to replace the activation function with a periodic function such as the sinusoid. A sinusoid activation function can preserve arbitrary order of signals derivatives, which allows it to capture high-frequency details. A disadvantage of sinusoidal activation functions is their notoriously hard initialization that significantly affects the learned output [43]. To address this issue, we first map the activation range into a monotone interval by tuning the amplitude of the sinusoid. Then, the LayerNorm [44] is exploited before each activation. This normalization ensures stable regularization with low dependencies on mini-batch data.

At each iteration, we optimize the printing policy by interacting with the environment. We start by printing a batch of exemplars. Next, we optimize the policy parameters by maximizing the following function:

$$\arg\max_{\theta} \mathbb{E}_e \left[\frac{\pi_{\theta_e}(a_e|s_e)}{\pi_{\theta_{e-1}}(a_e|s_e)} \hat{A}_e \right], \tag{6}$$

where e is the current iteration, θ are the learned parameters of CNN, which encodes the policy π that generates an action a_t in terms of the observed states s_t, $\hat{A}_t$ is the estimator of the advantage function generated by our critic network, and the expectation $\mathbb{E}_t$ is an average over our printing batch.

To maximize Equation 6, we use the proximal policy optimization (PPO) algorithm [45]. We set the learning rate and entropy coefficient to 0.001 and 0.0001, both of which are annealed to 0 during training. We use a discount factor of 0.99 and a batch size of 1.

VI. Evaluation

In this section, we conduct extensive experiments to demonstrate the performance of our printing policy.

A. Dataset and Implementation Details

We use a random set of 20 images as our dataset. Each image has approximately the resolution of 1000×1000 pixels. Given the color of the assumed primary materials (CMYKW), one cannot usually cover the gamut (*i.e.*, range) of all (s)RGB colors. Thus, all images in this work are first *gamut mapped* into the space of colors achievable by our primaries (that is the colors our simulated printer can create). We follow standard gamut mapping algorithms for this step [41]. During the training, we randomly crop patches with the size of 75×75 from the training dataset, which are sequentially fed into the agent in the form of 1D sliding window. For keeping the input size of networks, we pad the 1D window with black value when it passes the image boundaries. During the validation, our agent sequentially takes the whole target image as input, and reproduces the result with the full resolution.

B. Reproduction Performance

As illustrated in Figure 2, our results show high-fidelity to the target images. The colors reproduced by our printing policy are bright, vivid, and manifest high-contrast. Even more importantly, our printing policy excels at reproducing the image structure. By carefully controlling the fusion ratio at each printing location, our policy is capable of minimizing the inherent horizontal artifacts and reproduce features with extremely high spatial details (Figure 2 insets). In certain images there is an observable loss of details in our reproduction. We attribute this to the challenging task. Note that our policy has to learn two different tasks: 1) color reproduction, *i.e.*, for a given input RGB color what are the mixing ratio of base materials 2) structure reproduction,

Fig. 3. Samples of our training dataset. We emphasize there is a significant domain gap between the training dataset and test dataset. Our validation results demonstrate the proposed method can bridge this gap by learning an effective multi-material fusion policy.

by taking into account the inherent delay of the simulated printer. It is worth noting that all the results derive from the test dataset, which means our agent is capable of generalizing to unseen data distribution. The training samples are shown in Figure 3, where a considerable domain gap with the test data can be noticed.

C. Comparison with Baseline Deposition

We design a baseline multi-material fusion to show the benefit of our optimized policy. For the color reproduction task (see above), we use our own policy to determine the mixing ratios of primary materials. The baseline, however, makes a naive assumption concerning the structure reproduction (task 2 above): it treats the material changes as if they can occur instantaneously while in reality the computed mixing ratios are fed directly into our simulator. We seek to show that such naive strategy does not lead to optimal deposition. We can see the result of such an approach in Figure 4 (middle). The naive baseline is successful in reproducing the colors of the image. However, the appearance is significantly degraded due to the artifacts in image structure. We can observe severe line artifacts that are caused by the delay between material changes. In contrast, our fusion control policy performs significantly better on the same input (Figure 4 right). In our result, the line artifacts are minimized and the detailed texture from the input image is well preserved. Additionally, we compare our method and the naive baseline on the whole test dataset. The quantitative evaluation is listed in Table I, it shows our superiority especially in structural measurement.

TABLE I

QUANTITATIVE EVALUATIONS ON COMPARISON METHODS.

Metrics	PSNR↑	SSIM↑	PSNR_{std} ↓	SSIM_{std} ↓
Naive Baseline	18.38	0.4081	1.94	0.1117
Ours	22.21	0.7003	1.64	0.0783

D. Comparison With State of the Art

We compare our method with the state-of-the-art method for color fusion in FDM [37]. This method leverages both

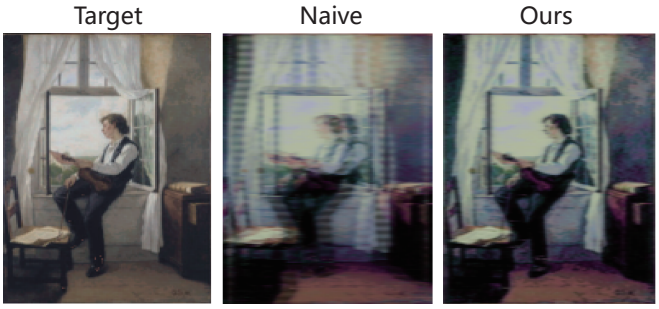

Fig. 4. Comparison between a naive solution to the color mixing problem and our method. Left: The target image. Middle: Result with a naive approach where the feed rate are adjusted only considering the current pixel. Right: Our approach.

fusion *and* superposition of materials to obtain the final color. First, each layer is decomposed into a few sub-layers. Subsequently, a fixed fusion ratio for each sub-layer is computed. Assuming N filaments, to compute the fusion ratio for each sub-layer and the height of the sub-layer at each printing location, the smallest simplex in $\mathbb{R}^{N-1}$ that contains all colors of the layer should be found. As noted in the original work [37], this combinatorial problem becomes prohibitively complex with increasing the number of base filaments. Therefore, the authors limit their color output to CMY.

We can see a result of this limitation in Figure 5 top right. The lack of black and white primaries results in loss of contrast. Moreover, as the mixing ratio and the output color are not linearly linked, the choice of using a linear interpolation-based approach [37] to compute the mixing ratio creates errors in the final result. We can observe the effect of this assumption as undesired color shifts in the final output. In contrast, our method can leverage the full color gamut of fusing five materials, (Figure 5 bottom left). We can also observe that the superposition of single-colored sub-layers used in this method [37] reduces the spatial artifacts (Figure 5 bottom right). However, the sub-layer strategy creates two major drawbacks:

1) To switch from a sub-layer to the next, the material present in the nozzle need to be purged. This require building a separate structure close to the print [46];
2) The printing time multiplies by the number of sub-layers.

Our method does not suffer from these drawbacks as it deposits the materials without any pauses.

E. Ablation Study on Activation Function

To capture the high-frequency details of the input images, we propose to replace the activation function of the prevalent ReLU-based networks with a periodic activation function. Then, we further tame it with a handy mapping-to-normalization strategy to relieve the burdens of initialization. In this section we compare our activation function with existing alternatives. For comparisons we use the ReLU activation function due to its widespread use and a sinusoid activation function initialized as proposed in [43]. The results of our ablation study are depicted in Figure 6. Our proposed

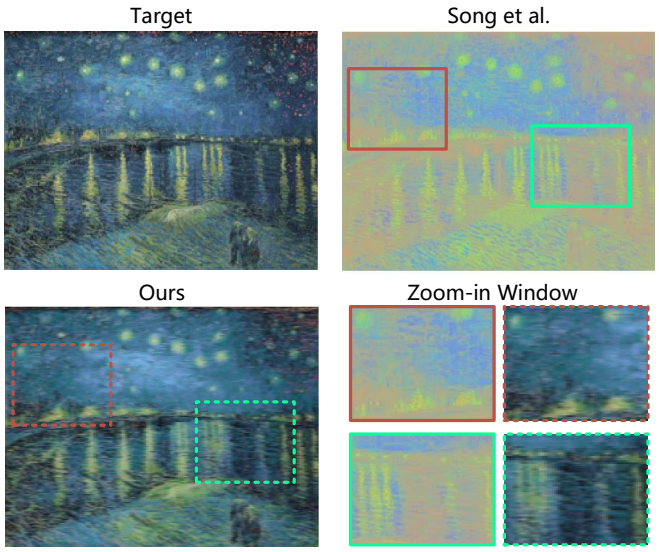

Fig. 5. Comparison with the state-of-the-art method Song et al. [37]. We show the gamut mapped target (we chose an image that was reachable by the comparison method for fair evaluation), the reproduction result of Song et al. [37] and ours, and the zoom-in details.

activation function achieved the best performance in the fusion printing task, showing faster convergence and more stable training during RL exploration. We can also observe a significant improvement in image texture. Thanks to the ability to process high-frequencies our activation function enables learning of control strategies that successfully eliminate the line artifacts, (Figure 6 middle). Moreover, in regions with lower frequency texture our activation function achieves more realistic results, resulting in a better color match, (Figure 6 right).

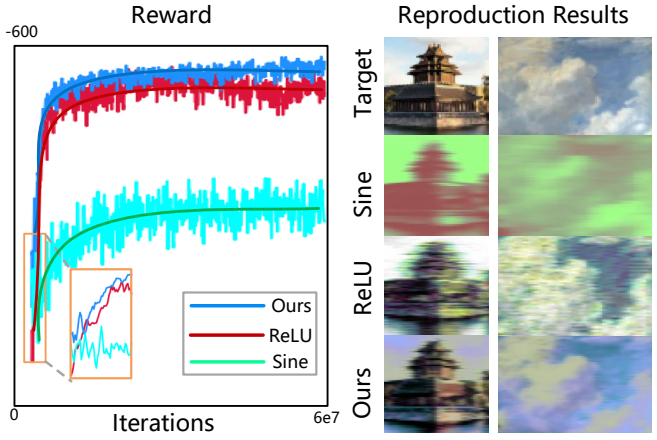

Fig. 6. Ablation study on activation function of our control policy. Left: the reward curves of three comparison methods; Middle and Right: the reproduction results in terms of the high-frequency and low-frequency targets.

VII. Limitations and Future Work

Our control policies are trained in a simulated environment. An interesting direction for future work is to deploy the policies on fabrication hardware. The largest sim-to-real gap is introduced at our simulation step. For efficiency we assume that each material gbvt stays solid immediately after deposition. This is a fair assumption for many fabrication processes such as FDM. However, some interesting materials, *e.g.*, silicons or hydrogels, do not posses this property. To control the deposition of such materials we would need to model their rheological properties similar to [16].

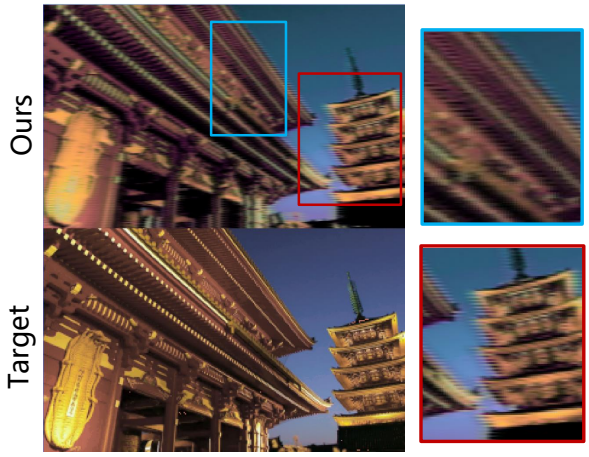

Fig. 7. A failure case on the transition between the low-frequency and high-frequency details.

Our method can handle regions of both high- and low-frequency spatial details. However, at borders between a high- and low-frequency regions we can observe transition artifacts (Figure 7). These artifacts are caused by a difference in control strategies required to optimally deposit each region. The shift in the behavior requires some adaptation time that results in visible artifacts. A potential future work lies in investigating why there is a need for different strategies to handle different regions of the image. The insights gained from our agent could inform future design of multi-material objects.

VIII. Conclusion

In this work we presented a numerical model for multi-material fusion 3D printing. We utilize our model in a reinforcement learning framework to discover optimal policies for multi-material deposition. To handle material deposition at high spatial resolution we propose a custom activation function for our neural network. We evaluate our learned printing policy on a case-study of color fabrication. We demonstrate that our reproduced color samples achieve excellent contrast and chromaticity. Even more importantly, we show that by carefully controlling the deposition it is possible to minimize the artifacts inherent in fusing multiple materials. Our results compare favorable against both commercial baseline controllers and state-of-the-art hand-design algorithms. We believe that our findings illustrate that by using optimal control strategies we can break the hardware limits. We believe that the future of fabrication lies in the joint design of hardware and software.

References

[1] F. Tricot, C. Venet, D. Beneventi, D. Curtil, D. Chaussy, T. Vuong, J. Broquin, and N. Reverdy-Bruas, "Fabrication of 3d conductive

circuits: print quality evaluation of a direct ink writing process," *RSC advances*, vol. 8, no. 46, pp. 26 036–26 046, 2018.

[2] Y. Cheng, K. H. Chan, X.-Q. Wang, T. Ding, T. Li, X. Lu, and G. W. Ho, "Direct-ink-write 3d printing of hydrogels into biomimetic soft robots," *ACS nano*, vol. 13, no. 11, pp. 13 176–13 184, 2019.

[3] B. Zhu, M. Skouras, D. Chen, and W. Matusik, "Two-scale topology optimization with microstructures," *ACM Trans. Graph.*, vol. 36, no. 4, jul 2017. [Online]. Available: https://doi.org/10.1145/3072959.3095815

[4] V. Babaei, K. Vidimče, M. Foshey, A. Kaspar, P. Didyk, and W. Matusik, "Color contoning for 3d printing," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–15, 2017.

[5] T. Chen and K. Shea, "Computational design of multi-stable, reconfigurable surfaces," *Materials & Design*, vol. 205, p. 109688, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0264127521002409

[6] J. Prusa, "Prusa i3 mk2 multi-material upgrade," http://www.prusaprinters.org/original-prusa-i3-mk2-multi-material-upgrade-release/, 2016.

[7] E. Sammut, "The prometheus system," https://www.kickstarter.com/projects/811909269/the-prometheus-system-intuitive-multi-filament-3d, 2016.

[8] Mosaic, "Palette," https://www.kickstarter.com/projects/mosaic3d/the-palette-3d-printing-evolved, 2015.

[9] M. A. Skylar-Scott, J. Mueller, C. W. Visser, and J. A. Lewis, "Voxelated soft matter via multimaterial multinozzle 3d printing," *Nature*, vol. 575, no. 7782, pp. 330–335, Nov 2019. [Online]. Available: https://doi.org/10.1038/s41586-019-1736-8

[10] J. Z. Gul, M. Sajid, M. M. Rehman, G. U. Siddiqui, I. Shah, K.-H. Kim, J.-W. Lee, and K. H. Choi, "3d printing for soft robotics–a review," *Science and technology of advanced materials*, vol. 19, no. 1, pp. 243–262, 2018.

[11] P. M. Bhatt, A. M. Kabir, R. K. Malhan, B. Shah, A. V. Shembekar, Y. J. Yoon, and S. K. Gupta, "A robotic cell for multi-resolution additive manufacturing," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 2800–2807.

[12] C. Wu, C. Dai, G. Fang, Y.-J. Liu, and C. C. Wang, "Robofdm: A robotic system for support-free fabrication using fdm," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1175–1180.

[13] G. Hunt, F. Mitzalis, T. Alhinai, P. A. Hooper, and M. Kovac, "3d printing with flying robots," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 4493–4499.

[14] S. Kim, D. D. Kim, and B. Anthony, "Dynamic control of a fiber manufacturing process using deep reinforcement learning," 2019. [Online]. Available: https://arxiv.org/abs/1911.10286

[15] F. Ogoke and A. B. Farimani, "Thermal control of laser powder bed fusion using deep reinforcement learning," *Additive Manufacturing*, vol. 46, p. 102033, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214860421001986

[16] M. Piovarci, M. Foshey, J. Xu, T. Erps, V. Babaei, P. Didyk, S. Rusinkiewicz, W. Matusik, and B. Bickel, "Closed-loop control of direct ink writing via reinforcement learning," *SIGGRAPH*, 2022.

[17] ORD-Solutions, "Rova4d full color blender 3d printer," http://www.ordsolutions.com/rova4d-full-color-blender-3d-printer-pre-order/, 2016.

[18] RepRap.me, "Diamond nozzle," http://reprap.org/wiki/Diamond_Hotend, 2015.

[19] F. Chan, "Nix color printer," https://www.kickstarter.com/projects/517014753/the-world-first-high-resolution-full-color-3d-prin?lang=fr, 2017.

[20] N. Siegel, J. T. Springenberg, F. Berkenkamp, A. Abdolmaleki, M. Neunert, T. Lampe, R. Hafner, N. Heess, and M. Riedmiller, "Keep doing what worked: Behavior modelling priors for offline reinforcement learning," in *International Conference on Learning Representations*, 2020. [Online]. Available: https://openreview.net/forum?id=rke7geHtwH

[21] S. Reddy, A. D. Dragan, and S. Levine, "SQIL: imitation learning via reinforcement learning with sparse rewards," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [Online]. Available: https://openreview.net/forum?id=S1xKd24twB

[22] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3389–3396.

[23] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "Tossingbot: Learning to throw arbitrary objects with residual physics," 2019.

[24] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.

[25] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.

[26] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.

[27] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, "Sim-to-real: Learning agile locomotion for quadruped robots," *arXiv preprint arXiv:1804.10332*, 2018.

[28] T. Haarnoja, S. Ha, A. Zhou, J. Tan, G. Tucker, and S. Levine, "Learning to walk via deep reinforcement learning," *arXiv preprint arXiv:1812.11103*, 2018.

[29] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, p. eabc5986, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.abc5986

[30] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Overcoming exploration in reinforcement learning with demonstrations," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 6292–6299.

[31] M. Zamora, R. Poranne, and S. Coros, "Learning solution manifolds for control problems via energy minimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7912–7919, 2022.

[32] S. D. Patrick, A. Nycz, M. W. Noakes, and K. T. Gaul, "Reinforcement learning for generating toolpaths in additive manufacturing," 8 2018. [Online]. Available: https://www.osti.gov/biblio/1474597

[33] A. G. Dharmawan, Y. Xiong, S. Foong, and G. Song Soh, "A model-based reinforcement learning and correction framework for process control of robotic wire arc additive manufacturing," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4030–4036.

[34] L. Shi, V. Babaei, C. Kim, M. Foshey, Y. Hu, P. Sitthi-Amorn, S. Rusinkiewicz, and W. Matusik, "Deep multispectral painting reproduction via multi-layer, custom-ink printing," *ACM Trans. Graph.*, vol. 37, no. 6, dec 2018. [Online]. Available: https://doi.org/10.1145/3272127.3275057

[35] D. Sumin, T. Rittig, V. Babaei, T. Nindel, A. Wilkie, P. Didyk, B. Bickel, J. Křivánek, K. Myszkowski, and T. Weyrich, "Geometry-aware scattering compensation for 3D printing," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 38, Jul. 2019.

[36] T. K. Nindel, T. Iser, T. Rittig, A. Wilkie, and J. Křivánek, "A gradient-based framework for 3d print appearance optimization," *ACM Trans. Graph.*, vol. 40, no. 4, Jul. 2021. [Online]. Available: https://doi.org/10.1145/3450626.3459844

[37] H. Song, J. Martínez, P. Bedell, N. Vennin, and S. Lefebvre, "Colored fused filament fabrication," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 5, pp. 1–11, 2019.

[38] P. Kubelka, "Ein beitrag zur optik der farbanstriche (contribution to the optic of paint)," *Zeitschrift fur technische Physik*, vol. 12, pp. 593–601, 1931.

[39] Šárka Sochorová and O. Jamriška, "Practical pigment mixing for digital painting," *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2021)*, vol. 40, no. 6, p. 234, 2021.

[40] G. Wyszecki and W. S. Stiles, *Color science*. Wiley New York, 1982, vol. 8.

[41] J. Morovic and M. R. Luo, "The fundamentals of gamut mapping: A survey," *Journal of Imaging Science and Technology*, vol. 45, no. 3, pp. 283–290, 2001.

[42] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International conference on machine learning*. PMLR, 2016, pp. 1928–1937.

[43] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions,"

Advances in Neural Information Processing Systems, vol. 33, pp. 7462–7473, 2020.

[44] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.

[45] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[46] S. Hornus, S. Lefebvre, J. Dumas, and F. Claux, "Tight printable enclosures for additive manufacturing," Inria, Research Report RR-8712, Apr. 2015. [Online]. Available: https://hal.inria.fr/hal-01141706