
Invertible Temper Modeling using Normalizing Flows and the Effects of Structure Preserving Loss

Sylvia Howland¹ Keerti Sahithi Kappagantula¹ Henry Kvinge^{1,2} Tegan Emerson^{1,3,4}

Abstract

Advanced manufacturing research and development is typically small-scale, owing to costly experiments associated with these novel processes. Deep learning techniques could help accelerate this development cycle but frequently struggle in small-data regimes like the advanced manufacturing space. While prior work has applied deep learning to modeling visually plausible advanced manufacturing microstructures, little work has been done on data-driven modeling of how microstructures are affected by heat treatment, or assessing the degree to which synthetic microstructures are able to support existing workflows. We propose to address this gap by using invertible neural networks (normalizing flows) to model the effects of heat treatment, e.g., tempering. The model is developed using scanning electron microscope imagery from samples produced using shear-assisted processing and extrusion (ShAPE) manufacturing. This approach not only produces visually and topologically plausible samples, but also captures information related to a sample’s material properties or experimental process parameters. We also demonstrate that topological data analysis, used in prior work to characterize microstructures, can also be used to stabilize model training, preserve structure, and improve downstream results. We assess directions for future work and identify our approach as an important step towards end-to-end deep learning system for accelerating advanced manufacturing research and development.

¹Pacific Northwest National Laboratory, Washington, USA
²Department of Mathematics, University of Washington, Washington, USA
³Department of Mathematics, Colorado State University, Colorado, USA
⁴Department of Mathematical Sciences, University of Texas El Paso, Texas, USA. Correspondence to: Tegan Emerson <tegan.emerson@pnl.gov>.

Proceedings of the Geometry-grounded Representation Learning and Generative Modeling at 41st International Conference on Machine Learning, Vienna, Austria. PMLR Vol Number, 2024. Copyright 2024 by the author(s).

1. Introduction and Motivation

Effective manufacturing processes underpin enormous swathes of the contemporary economy and everyday life: device miniaturization, scaling sustainable energy infrastructure, and the efficacy or longevity of endless varieties of construction materials and other goods are just a few foundational ways in which improved materials and manufacture impact the world at large. Cutting-edge advanced manufacturing processes, by nature of their novelty and presently small scale, do not have access to the same kinds of established, reliable simulations used to probe hypotheses available to more conventional processes. First-principle simulation capabilities are under development for these advanced processes, but are time- and resource-intensive approaches. In the meantime practitioners frequently rely on small-scale datasets and their learning when undertaking research and development. Introducing data driven tools that help guide and accelerate the R&D cycle could go a long way in mitigating research risk and cost.

Shear-Assisted Processing and Extrusion (ShAPE) (Whalen et al., 2021) is an emerging advanced manufacturing technology that has been shown to produce parts with state-of-the-art material performance in forms ranging from rods, tubes, and wires (Li et al., 2021; Reza-E-Rabby et al., 2022; Kalsar et al., 2022; Li et al., 2022), from various forms of metallic feedstock. Notably, ShAPE’s enhanced performance carries over to bulk-scale components, giving the technology a pathway to industrial-scale adoption and presenting a need to accelerate research and development towards this end.

In the absence of complete first-principle, physics-based models advanced manufacturing researchers heavily rely on a scaffolding of process-structure-property relationships when developing new materials or exploring advanced manufacturing pathways such as ShAPE. Process corresponds to the various manufacturing settings or choices that can be parameterized. Structure corresponds to actual material microstructures that can be inspected by a subject matter expert and if frequently captures using sophisticated imaging technologies. For ShAPE processing of aluminum alloy 7075 (considered in this work), microstructures are observed through the use of Scanning Electron Microscopy (SEM). Finally, the properties can be measured using a variety of

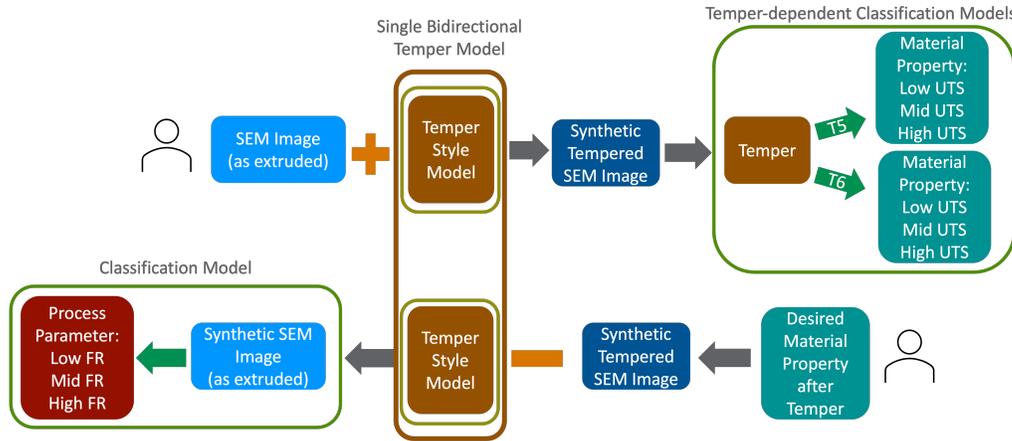


Figure 1. Illustration of an AI-enabled pipeline to accelerate research and development in advanced manufacturing; specifically for Shear Assisted Processing and Extrusion. “FR” stands for “feed rate” and represents the relevant processing parameter considered. “UTS” is ultimate tensile strength and is the material property considered in this initial demonstration.

destructive and nondestructive means based on the objective of manufacturing. Some technologies rely on a separate or secondary heat treating process called tempering that can strongly influence the resulting material properties as well as the visual appearance of the associated microstructures.

During the development of novel manufacturing techniques, subject matter experts develop an intuition for the relationships between process parameters and resulting material properties primarily through analysis of microstructure. Ideally, a materials science would like to understand this relationships in two directions. The “forward” direction is the ability to predict resulting material properties for a set of selected processing parameters. Alternatively, the “backward” direction indicates the ability to identify the set of processing parameters that would give rise to a desired material property.

Recently, machine learning, and more specifically deep learning, has emerged as a potential accelerant of developing process-structure-property relationships and is a framework we believe could be useful in answering the above questions. Deep learning techniques have found application in tasks ranging from SEM image classification and segmentation (Tsutsui et al., 2020; Azimi et al., 2018; Müller et al., 2020; Durmaz et al., 2021) to generating synthetic microstructures (Iyer et al., 2019; Howland et al., 2023). While these generative approaches aim to solve similar problems to our “forward” and “backward” questions, none to our knowledge address them while incorporating advanced manufacturing and post-processing procedures into their modeling framework.

There are numerous challenges associated with trying to directly model the relationships between process parameters and material properties. Among the challenges are that the

backward direction is a many-to-one map (e.g., multiple processing parameter combinations can result in the same resulting material property) and significant data limitations due to the time and cost associated with experimentation. In Figure 1 we provide an illustration for how a subject matter expert could interact with a machine learning assisted workflow for ShAPE. The upper and lower portions of the diagram correspond to the forward and backward directions, respectively, that were previously described.

Unlike prior work, we treat tempering as a secondary process and recognize its importance in both the forward and backward paths. This work demonstrates a first foray into developing a single bi-directional temper model through the use of normalizing flows. Building on the prior work in (Emerson et al., 2022), we further seek to leverage techniques in topological data analysis that were shown to preserve relevant structure and useful in capturing temper dependent features. In this work, we seek to answer the following research questions. First, can we train a deep neural network to successfully perform both visually and topologically plausible ShAPE microstructure generation for different temper conditions? Second, can this temper transfer model preserve microstructure label information or equivalently, can we infer the original process parameter or material property labels post-transfer?

2. Background and Related Work

2.1. Shear-Assisted Processing and Extrusion (ShAPE) Manufacturing

Shear-Assisted Processing and Extrusion (Whalen et al., 2021) is an advanced manufacturing process wherein a billet is extruded from a container equipped with a coaxial man-

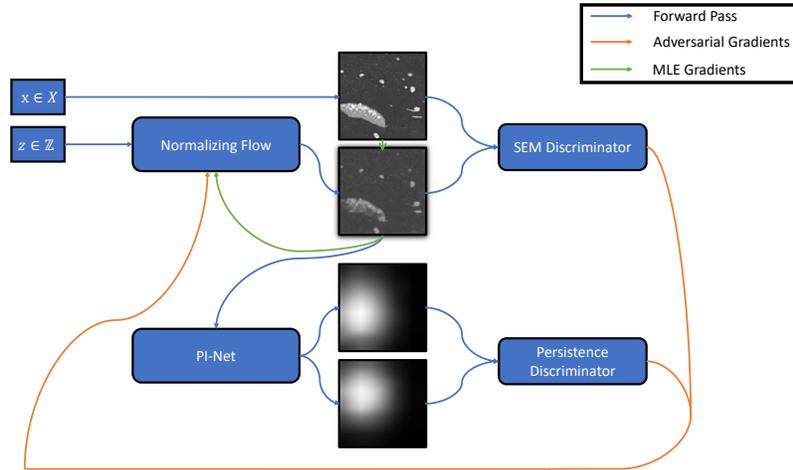


Figure 2. A high-level overview of our data flow and model training for a given normalizing flow. Experimental samples are drawn from our ShAPE SEM dataset ($x \in \mathbb{X}$) to act as ground truth samples for adversarial discriminator networks. We sample standard Gaussian noise as our models’ base distribution ($z \in \mathbb{Z}$), which is transformed into an SEM microstructure sample by the normalizing flow. Synthetic samples are used to train the flow using maximum likelihood estimation. We use a frozen persistence homology approximator, PI-Net, to generate topological persistence imagery (PIs) for experimental and synthetic samples. Separate SEM and PI discriminators provide adversarial feedback to the normalizing flow to encourage plausible generations.

drel when impinged upon by a spinning die. Friction heats the billet, working in tandem with the attendant shearing forces to plasticize it. ShAPE experiments are parameterized by various input streams or process parameters, such as temperature, torque, power, and the most important to our work being the the traverse rate or “feed rate”.

This study uses process and property data derived from ShAPE-manufactured AA7075 tubes previously explored by (Howland et al., 2023) and (Emerson et al., 2022) but originally developed by (Whalen et al., 2021). In the original study by Whalen et al., they manufactured various tubes using aluminum alloy (AA) 7075 feedstock via ShAPE, of which some were subjected to different heat treatments. The two heat treatments applied are identified as T5 and T6 heat treatments which correspond to subjecting the ShAPE samples to different temperature profiles over time. Post heat treatment, the tubes’ material property performance was evaluated to determine their percent-elongation, yield strength, and most importantly for our purposes, ultimate tensile strength (UTS) in the tensile loading condition. The SEM images of the heat treated AA7075 ShAPE tubes with different heat treatments collected by the authors as a part of the second study form the basis of the data analyzed herein. In previous work, as in this study as well, entire images were not used; rather image ‘chips’ with specific dimensions were used to train the models.

2.2. Normalizing Flows

Normalizing flows (Dinh et al., 2014; Rezende & Mohamed, 2015; Kobyzev et al., 2021) are a family of generative neural

networks which are able to perform efficient, exact sampling and density evaluation. Normalizing flows, like generative adversarial networks (GANs) (Goodfellow et al., 2014), typically operate by transforming random samples from a simple base distribution – such as a standard normal – into a more complex target distribution of interest via a deep neural generator network. Unlike GANs, however, normalizing flow networks are built using components which are differentiable *as well as* invertible and possessing easily-computed Jacobian determinants. Their “forward” direction, from base to target distribution, allows for efficient sampling while the “backward” direction, due to the flow’s exact invertibility and accessible Jacobian determinants, allows for precise and tractable likelihood evaluation. Flows tend to have lower sample quality than GANs, but train more stably and consistently. For our purposes, where invertibility is desirable and structural plausibility is more important than impeccable visual quality, normalizing flows are a natural fit.

The normalizing flow models used in this study are adapted from Alignflow (Grover et al., 2020), which incorporates a GAN-esque adversarial loss and multi-network configuration to extend flows to unsupervised domain adaptation and image-to-image translation problems. Alignflow uses a shared base distribution (latent space) between per-domain flows. Due to the invertibility of these sub-flows, they can be composed in arbitrary order to consistently translate a sample from any one domain into all other domains. These attributes are a natural fit for our temper transfer problem. Alignflow has a limitation shared with other normalizing

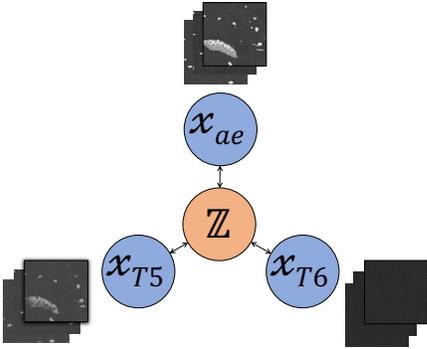


Figure 3. An overview of our temper transfer problem setup. We consider three domains: one for the untempered condition (as extruded, or *ae*) as well as one each for the T5 and T6 temper conditions. We train one invertible deep neural network per domain, sharing a latent base distribution \mathbb{Z} . This allows us to sample synthetic microstructures in any of our domains, as well as (de)temper synthetic *and* experimental microstructures.

flows: its source and target distributions must have the same dimensionality to maintain invertibility. Since Alignflow is built for image-to-image problems, the high dimensionality of each data sample can quickly pose a prohibitive cost. Fortunately, our image data (chips from the SEM images) are on the scale of 128x128 pixels – small enough to sidestep this limitation even using three normalizing flows instead of the original authors’ two. An illustration of this set up is provided in Figure 3.

2.3. Topological Analysis of Material Microstructures

Topological data analysis possesses myriad tools useful for studying the shape of data, including SEM microstructures corresponding to specific ShAPE process parameters or with specific material properties. Persistent homology (Edelsbrunner et al., 2000) is one such family of tools. While there are many different techniques for capturing the notion of topology in data, for analysis of the grayscale SEM images the authors utilize sublevel set persistent homology. Sublevel set filtration on grayscale images captures topological structure using pixel intensity relationships. Precipitates manifest as higher contrast, bright regions in the SEM imagery creating a natural connection to the filtration mechanism used in sublevel set filtration. By treating the image as a 2-dimensional function, we are able to capture both 0- and 1-dimensional topological features based on a nested sequence of cubical complexes. The topological features are first captured in a persistence diagram and then converted to persistence images.

Persistence images have been shown to be useful as both SEM microstructure feature descriptors (Emerson et al., 2022) and as a useful diagnostic tool for synthetic SEM

evaluation (Howland et al., 2023), at least when applied to images that can be described in terms of their precipitate, pore, or second phase distribution. Both of these methods examined ShAPE AA7075 tubes as well, though ours is the first work in this area to move beyond persistence images as descriptive or analytical tools and incorporate them as signal into the model training process itself. For a more thorough treatment of persistence images, we refer readers to their introduction in (Adams et al., 2017). Persistence images capturing the topological features derived from the sublevel set filtration provide a global summary of the precipitate distribution and intensities. The proposed approach provides explainability without the need for image segmentation as a preprocessing step for subsequent statistical summaries.

All persistence homology calculations and associated persistence images are computed using the Ripser (Bauer, 2021) package with 10x10 pixel persistence images.

3. Experimental Design

In our experimental setup we hope to evaluate two related questions. The first focuses on the visual plausibility of synthesizing the effects of heat treatment with and without structure preserving loss which is grounded in a more qualitative discussion. Secondly, we seek to quantify the benefit of our novel combination of techniques for topologically penalized, joint latent-space learning. We describe the dataset used for analysis, our label preserving experimental design, our normalizing flow model, and finally how we implement our topological regularization.

3.1. SEM Microstructure Dataset

We train and evaluate all of our models using the ShAPE AA7075 process-structure-property dataset used by (Howland et al., 2023). The dataset consists of roughly 437,000, partially overlapping, 128x128 pixel image chips derived from 32 distinct experimental samples. Each experimental run results in a single SEM image corresponding to a single ShAPE AA7075 sample produced by recorded process parameters. We associate each chip with its experiment’s corresponding feed rate (process parameter). For the samples that undergo tempering to T5 or T6 the resulting ultimate tensile strength (material property) is also recorded. Inspired by ML literature on utilizing continuous labels in low-data regimes (Truong et al., 2021; Ahmad et al., 2018; Workman et al., 2016; Ammar Abbas & Zisserman, 2019), these labels are evenly discretized into “low”, “mid”, and “high” bins to create a surrogate classification problem in place of regression given the sparsity of samples.

3.2. Label Preservation Experimental Setup

To assess the utility of normalizing flows for modeling heat treatment on this data set, we measure the performance of models trained on features derived from real data representations to generalize to features extracted from SEM images resulting from synthetic (de)tempering.

More concretely for the quantitative analysis, for the forward model we train a simple support vector machine classifier on real SEM images (or persistence images derived from real SEM images) to assign a material property class label per temper type. This results in two classifiers for the forward direction; one for T5 and one for T6. The 3-class classification accuracy for predicting material property from real T5 and T6 imagery is shown in Figure 5. For the backward direction, we train a classifier on features from real, As Extruded, SEM images to predict process parameter class label also discretized into a surrogate 3-class problem. The baseline results for this direction are also shown in Figure 5.

After training these 3 classification models, we can assess their performance on synthetic tempering of As Extruded to T5 or T6 or (de)tempering real T5 or T6 imaged to simulated As Extruded images. We perform 3 different runs for each classifier trained on real data and report the 95% confidence intervals with errorbars in Figure 5.

3.3. Alignflow

Our normalizing flows are adapted from Alignflow (Grover et al., 2020) and its official Pytorch (Paszke et al., 2019) implementation. Our models, both the generative flow components and their associated discriminators, are trained using the default hyperparameters from the official implementation except as follows. Since our dataset has three domains (As Extruded samples as well as T5 and T6 treatment conditions), we modify the repository’s “flow2flow” model to create a “multiflow2flow” which generalizes the architecture to handle an arbitrary number of image domains. Due to the increased memory costs associated with adding a third domain, we found it necessary to halve the batch size from 16 to 8. Finally, we set the number of input channels to 1 since we are working with grayscale imagery and double the number of mid-network flow channels from 32 to 64.

3.4. Topological Regularization Implementation

A well-known bottleneck for incorporating topological features into machine learning is how topological feature extraction methods scale. This has been a significant barrier to incorporating persistence homology into training iterations. In order to incorporate topological features into our model training we supplement Alignflow’s GAN-esque visual adversarial loss with a similar, *topological* adversarial loss following the same structure. The associated discrimina-

tors, instead of making authenticity classifications based on experimental or synthetic SEM images, do so using 10x10 persistence images derived from the SEM chips. This approach allows our models to receive more direct feedback on their topological fidelity – a desirable property to optimize, given the ability of persistence images to describe ShAPE SEM microstructures (Emerson et al., 2022).

A well-known bottleneck for incorporating topological features into machine learning is how topological feature extraction methods scale. This has been a significant barrier to incorporating persistence homology into training iterations. For example, it would fail to scale if we sought to calculate the sublevel set persistence homology on each synthetic SEM image at every training iteration. To resolve this issue we adapt PI-Net (Som et al., 2020), a convolutional network (O’Shea & Nash, 2015) trained to efficiently approximate persistence imagery given a source SEM image.

PI-Net is small due to the small dimensionality and visual simplicity of both our SEM chips and persistence imagery. We train our PI-Net to approximate persistence images of the same SEM microstructures used in our Alignflow training set. Ground truth persistence images were computed using the Ripser and Persim Python libraries (Bauer, 2021; Saul & Tralie, 2019). The network is trained to convergence with a batch size of 128, taking roughly an hour on a single Nvidia Tesla P100 GPU. Our PI-Net consists of five convolutional layers with 256, 256, 128, and 128 channels respectively with a batch normalization layer and ReLU nonlinearity between each layer. PI-Net then applies global average pooling and instance normalization before a final convolutional layer, which produces a single-channel persistence image. Each convolutional layer uses a 3x3 kernel and a stride length of one.

By leveraging PI-Net to approximate the topological summaries of the synthetic images during training, we can successfully mitigate the computational bottleneck. We observed negligible L1 discrepancies between experimental persistence images and PI-Net’s predictions on both our training set and on held-out experimental chips. Finally, we observed that PI-Net could approximate the persistence images of our entire training set within several minutes – roughly two orders of magnitude faster than the twelve to fourteen hours needed to compute the original ground truths. We use these same, frozen PI-Net weights in all of our experiments.

3.5. Model Training and Selection

Figure 2 depicts our overall architecture and training loop. We train normalizing flows with and without PI-Net loss. To determine the effects of topological regularization on model stability and consistency, we train three normalizing flows for each variant for a total of six models. Each flow is

trained for 30 epochs over our SEM image dataset. Training takes roughly 12 hours on a single Nvidia A100 GPU. We observe that our flows tend to slowly degrade in image quality towards the end of their training time. We suspect this may simply be due to the difficulty of balancing the challenging tempering and detempering directions of our task, but have as yet been unable to confirm this and leave a more in-depth exploration for future work. Since our flows are trained in an unsupervised fashion we needed to find a metric to use for automatic model selection, avoiding arbitrarily picking evaluation checkpoints whose outputs simply “looked right” after limited manual inspection. Since degradation of our flows manifested in increasing noise and artifacting, we tested using the synthetic chips’ average signal-to-noise ratio and found it to be a good proxy for sample quality.

4. Results

In this work we aim to demonstrate that normalizing flows can be used to plausibly model, and remove, the effects of tempering heat treatments. As a second line of inquiry, we explore our ability to correctly identify labels associated with the forward (e.g., predicting resulting material properties) and backward (e.g., predicting the process parameter that produced the associated real image). The qualitative results of this first line of inquiry are presented in Section 4.1. Quantitative results are provided in Section 4.2.

4.1. Synthetic SEM Analysis

We first perform a qualitative visual comparison of SEM chips to complement our persistence image analysis. Figure 4 shows examples of temper transfer in all directions and with and without topological regularization. All models perform transfer on the same source SEMs for the sake of comparison. In the As Extruded to T5 direction, we see that topological regularization has only modest impact: absent any regularization we consistently observe faint, “ghostly” gray-scale features unlike those seen in our experimental data. Introducing topological regularization leads to shrinkage or vanishing of small gray-scale features, with larger ones undergoing only minor changes in shape, which may be expected outcomes of T5 treatment applied to an As Extruded sample. In the reverse T5 to As Extruded direction we see substantial artifacting without topological regularization.

Temper transfer to and from T6 has more mixed success. Tempering from As Extruded to T6 without any topological regularization produces sparser precipitate patterns than when transferring to T5 but still maintains faint, unrealistic, “phantom” gray-scale features we might interpret as precipitates with similar structure to those in the As Extruded sample. In experimental T6 data, on the other hand,

we know that large precipitates are quite rare (although not non-existent as evident from the experimental SEM images). Introducing topological regularization appears to resolve this problem (as shown in Figure 4). However, in the example imagery shown in Figure 4, the removal of large precipitates comes at the expense of producing *any* significant precipitates, which is also inconsistent with experimental T6 imagery.

Achieving a synthetic image by detempering in the T6 to As Extruded direction appears exceptionally difficult. Without topological regularization we see the production of tiny, scattered precipitates far more in line with other T6 samples than with the As Extruded target condition. Topological regularization gets closest to the mark by producing less evenly distributed, larger precipitates, though still not as large and irregularly sized as those seen in experimental data. We suspect that because T6 images have very little visible structure as a source domain in the SEM images at the magnifications where data was collected in the second study, models are especially prone to degenerated or mode collapsed behavior. Ensuring that adversarial loss plays a strong role during training, at least in transfer directions like this with low signal-to-noise, is likely crucial. This combination of persistence image and SEM chip results affirm our main hypotheses: augmenting our normalizing flows with topological regularization seems to help produce more stable generated samples, and particularly in the T6 directions to induce more topologically realistic samples. However, topological regularization can lead to *more* artifacting in transfer directions where the source domain has low signal overall. Investigating techniques to adaptively select or learn the proper weighting of topological regularization for different input domains would be a productive direction for future work.

4.2. Label Preservation Analysis

In addition to the visual and topological properties of our synthetic chips, we would also like to evaluate whether they encode features that represent the process parameters and material properties associated with those microstructures. Figure 5 shows the accuracy of SVM classifiers fit on a training set of experimental persistence images, then evaluated on both the held-out experimental samples and all synthetic images from a given model. As Extruded classifiers are fit using binned feed rate (i.e., process parameter) labels, whereas T5 and T6 classifiers are fit using binned UTS (i.e., material property) labels. We compare classification performance using both SEM chips and persistence images derived from the SEM imagery and find in nearly all cases the persistence images achieve comparable or superior performance.

We see that parameter and property classification is chal-

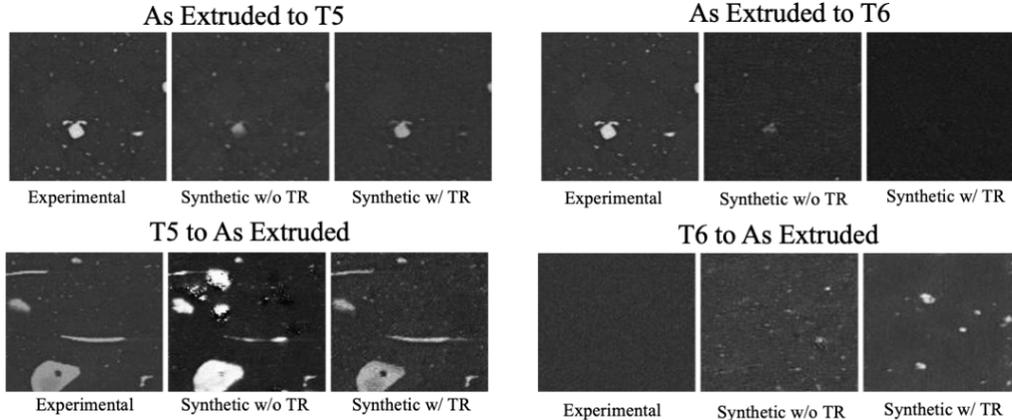


Figure 4. Synthetic (de)tempered SEM imagery generated by our normalizing flows with various levels of topological regularization. All synthetic images within a given direction are transferred from the same source. The labels “w/ TR” and “w/o TR” indicate with or without topological regularization, respectively.

lenging even using experimental imagery, with classifiers topping out at around 70% accuracy. Our synthetic models perform modestly but noticeably better in the “forward” tempering task (e.g., As Extruded to T5 and T6 directions) than the “backward” directions detempering to As Extruded. However, a number of results call into question just how useful this approach of classifying directly from microstructures or persistence images is. For example, topological regularization performs competitively with other synthetic models in the T6 to As Extruded direction despite producing heavily artifacted, implausible precipitates. In addition, the artifacting and unrealistic precipitate coloration in the baseline T5 to As Extruded direction does not stop those samples from being classified better than PI-Net samples.

These classification results for generated images largely mirror our other findings. Accuracy is highest in the As Extruded to T5 direction, which also has the visually cleanest synthetic images. Interestingly, it also achieves the highest accuracy on As Extruded to T6 samples by a wide margin. We suspect this may be due to the partial mode collapse in the regularized models; since those models tend to produce T6 images with either abnormally low or high gray-scale precipitate-like features, links between T6 microstructure and material property would necessarily be lost. The T5 and T6 to As Extruded directions proved exceptionally difficult to classify, with nearly all classifiers performing around or below chance. Given the visual coherence and reasonable persistence agreement for most models’ samples in these directions, we suspect that feed rate alone (or at least coarsely-binned feed rate) may be intrinsically difficult to recover from tempered samples.

A particularly significant result emerges when comparing our results to synthetic images generated by the conditional

GANs from (Howland et al., 2023), which achieve only chance-level performance on our label preservation task despite being visually and topologically plausible. Not only is the performance improved under this work’s demonstrated workflow, but this improvement comes despite the former approach directly conditioning on the target attributes during training. Incorporating conditioning information into our normalizing flows and verifying whether this label preservation result holds, particularly when performing sampling rather than temper transfer, would be a compelling direction for future work.

5. Conclusions and Future Work

This work introduced the use of normalizing flows for invertible temper transfer of ShAPE AA7075 microstructures, improving upon prior work in multiple ways. While previous work could generate heat-tempered microstructures, either by using per-temper generators or by incorporating it as conditioning information to the network, our approach can consistently transfer back and forth between them via a shared latent space. We observe that microstructures (de)tempered by our models preserve information about their experimental process parameters and material properties. This is particularly notable because our models are not exposed to any such information during training, and existing models which *do* condition on this information fail to produce samples that reflect it. We also find that optimizing our models to generate visually as well as topologically realistic samples helps to increase model consistency across multiple runs and improve the overall plausibility of synthetic imagery.

There are a number of promising avenues for future work. While our models already preserve some process param-

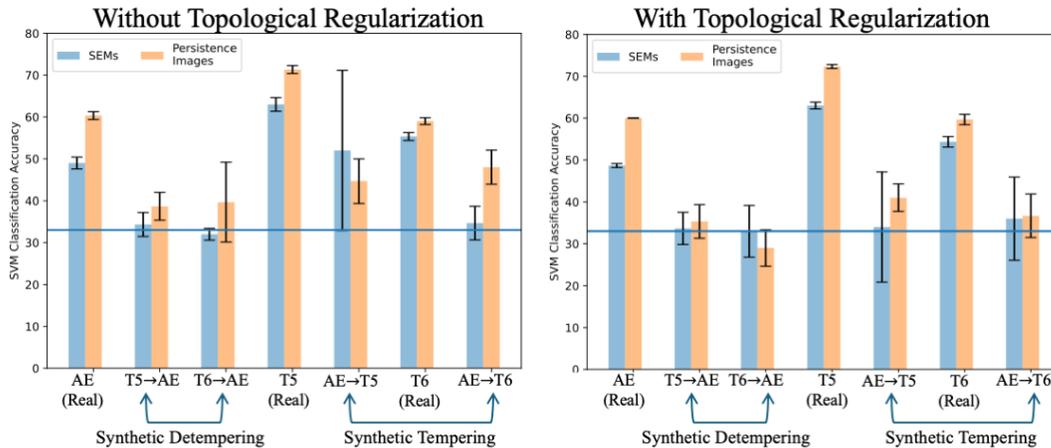


Figure 5. Label preservation accuracies for experimental and synthetic imagery without topological regularization either omitted (left) or with topological regularization (right). The notation AE in the figure indicates “As Extruded.” Accuracies and standard deviations are derived from three runs per model configuration. The blue, horizontal lines represent an at-chance performance of 33% accuracy.

ter and material property information when they perform domain transfer, a logical next step would be to see how sample quality and post-transfer information is preserved if we train our flows with explicit conditioning on these experimental attributes. An ideal solution would allow our flows to condition on a possibly imputed combination of process parameters *and* material properties; this would allow practitioners to explore potential microstructures based on desired properties or proposed experiments as desired. Incorporating less coarsely-binned conditioning labels, and incorporating multiple conditioning labels at once, remain open questions that could help produce more powerful generative models if solved but which are challenging due to the data constraints. If these problems could be adequately addressed, the field would be much closer to having a deep learning system applicable to design of experiments in advanced manufacturing. This effort is a successful first step in developing the framework for forward and invertible models that aid in accelerating advanced manufacturing research. While there are crucial additional steps essential for developing models that can be deployed in the research regime, efforts like this demonstrate the viability of data driven approaches to assist the R&D cycle.

Acknowledgments

This work was performed using the resources available at the Pacific Northwest National Laboratory (PNNL) and funded by the Mathematics for Artificial Reasoning in Science (MARS) Initiative as a Laboratory Directed Research and Development Project. PNNL is a multi-program national laboratory operated by Battelle Memorial Institute for the U.S. Department of Energy under contract DE-AC05-76RL01830.

References

- Adams, H., Emerson, T., Kirby, M., Neville, R., Peterson, C., Shipman, P., Chepushtanova, S., Hanson, E., Motta, F., and Ziegelmeier, L. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18, 2017.
- Ahmad, A., Khan, S. S., and Kumar, A. Learning regression problems by using classifiers. *Journal of Intelligent & Fuzzy Systems*, 35(1):945–955, 2018.
- Ammar Abbas, S. and Zisserman, A. A geometric approach to obtain a bird’s eye view from an image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pp. 0–0, 2019.
- Azimi, S. M., Britz, D., Engstler, M., Fritz, M., and Mücklich, F. Advanced steel microstructural classification by deep learning methods. *Scientific reports*, 8(1): 1–14, 2018.
- Bauer, U. Ripser: efficient computation of Vietoris-Rips persistence barcodes. *J. Appl. Comput. Topol.*, 5 (3):391–423, 2021. ISSN 2367-1726. doi: 10.1007/s41468-021-00071-5. URL <https://doi.org/10.1007/s41468-021-00071-5>.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation, 2014. URL <https://arxiv.org/abs/1410.8516>.
- Durmaz, A. R., Müller, M., Lei, B., Thomas, A., Britz, D., Holm, E. A., Eberl, C., Mücklich, F., and Gumbusch, P. A deep learning approach for complex microstructure inference. *Nature communications*, 12(1): 6272, November 2021. ISSN 2041-1723. doi: 10.1038/

- s41467-021-26565-5. URL <https://europepmc.org/articles/PMC8560760>.
- Edelsbrunner, H., Letscher, D., and Zomorodian, A. Topological persistence and simplification. In *Proceedings 41st annual symposium on foundations of computer science*, pp. 454–463. IEEE, 2000.
- Emerson, T., Kassab, L., Howland, S., Kvinge, H., and Kappagantula, K. S. Toptemp: Parsing precipitate structure from temper topology. In *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, 2022.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Grover, A., Chute, C., Shu, R., Cao, Z., and Ermon, S. Align-flow: Cycle consistent learning from multiple domains via normalizing flows. 34:4028–4035, Apr. 2020. doi: 10.1609/aaai.v34i04.5820. URL <https://ojs.aaai.org/index.php/AAAI/article/view/5820>.
- Howland, S., Kassab, L., Kappagantula, K., Kvinge, H., and Emerson, T. Parameters, properties, and process: Conditional neural generation of realistic sem imagery toward ml-assisted advanced manufacturing. *Integrating Materials and Manufacturing Innovation*, pp. 1–10, 2023.
- Iyer, A., Dey, B., Dasgupta, A., Chen, W., and Chakraborty, A. A conditional generative model for predicting material microstructures from processing methods, 2019.
- Kalsar, R., Ma, X., Darsell, J., Zhang, D., Kappagantula, K., Herling, D. R., and Joshi, V. V. Microstructure evolution, enhanced aging kinetics, and mechanical properties of aa7075 alloy after friction extrusion. *Materials Science and Engineering: A*, 833:142575, 2022.
- Kobyzev, I., Prince, S. J., and Brubaker, M. A. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, 2021. doi: 10.1109/TPAMI.2020.2992934.
- Li, X., Zhou, C., Overman, N., Ma, X., Canfield, N., Kappagantula, K., Schroth, J., and Grant, G. Copper carbon composite wire with a uniform carbon dispersion made by friction extrusion. *Journal of Manufacturing Processes*, 65:397–406, 2021.
- Li, X., Wang, T., Ma, X., Overman, N., Whalen, S., Herling, D., and Kappagantula, K. Manufacture aluminum alloy tube from powder with a single-step extrusion via shape. *Journal of Manufacturing Processes*, 80:108–115, 2022.
- Müller, M., Britz, D., Ulrich, L., Staudt, T., and Mücklich, F. Classification of bainitic structures using textural parameters and machine learning techniques. *Metals*, 10(5), 2020. ISSN 2075-4701. doi: 10.3390/met10050630. URL <https://www.mdpi.com/2075-4701/10/5/630>.
- O’Shea, K. and Nash, R. An introduction to convolutional neural networks. *arXiv preprint arXiv:1511.08458*, 2015.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance.pdf>.
- Reza-E-Rabby, M., Wang, T., Canfield, N., Roosendaal, T., Taysom, B. S., Graff, D., Herling, D., and Whalen, S. Effect of various post-extrusion tempering on performance of aa2024 tubes fabricated by shear assisted processing and extrusion. *CIRP Journal of Manufacturing Science and Technology*, 37:454–463, 2022.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1530–1538, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/rezende15.html>.
- Saul, N. and Tralie, C. Scikit-tda: Topological data analysis for python, 2019. URL <https://doi.org/10.5281/zenodo.2533369>.
- Som, A., Choi, H., Ramamurthy, K. N., Buman, M. P., and Turaga, P. Pi-net: A deep learning approach to extract topological persistence images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- Truong, L., Choin, W., Wight, C., Coda, E., Emerson, T., Kappagantula, K., and Kvinge, H. Differential property prediction: A machine learning approach to experimental design in advanced manufacturing. In *AAAI 2022 Workshop on AI for Design and Manufacturing (ADAM)*, 2021.
- Tsutsui, K., Terasaki, H., Uto, K., Maemura, T., Hiramatsu, S., Hayashi, K., Moriguchi, K., and Morito, S. A methodology of steel microstructure

recognition using sem images by machine learning based on textural analysis. *Materials Today Communications*, 25:101514, 2020. ISSN 2352-4928. doi: <https://doi.org/10.1016/j.mtcomm.2020.101514>. URL <https://www.sciencedirect.com/science/article/pii/S2352492820325253>.

Whalen, S. A., Kappagantula, K. S., Reza-E-Rabby, M., Li, X., Overman, N. R., Olszta, M. J., Wang, T., Herling, D. R., Suffield, S. R., Roosendaal, T. J., Taysom, B. S., Escobar Atehortua, J. D., Silverstein, J. A., Canfield, N. L., and Graff, D. D. Shear assisted processing and extrusion (shape) of aluminum alloy 7075, 2024, and al-12.4tm. 12 2021. doi: 10.2172/1843596. URL <https://www.osti.gov/biblio/1843596>.

Workman, S., Zhai, M., and Jacobs, N. Horizon lines in the wild. *arXiv preprint arXiv:1604.02129*, 2016.