

Hybrid Geometry/Property Autoencoders for Multi-Lattice Transitions

Martha Baldwin,¹ Nicholas A. Meisel,² Christopher McComb¹

¹Department of Mechanical Engineering,
Carnegie Mellon University, Pittsburgh, PA 15213

²School of Engineering Design and Innovation,
The Pennsylvania State University, University Park, PA 16802

Abstract

Additive manufacturing has revolutionized structural optimization by enhancing component strength and reducing material requirements. One approach used to achieve these improvements is the application of multi-lattice structures. The performance of these structures heavily relies on the detailed design of mesostructural elements. Many current approaches use data-driven design to generate multi-lattice transition regions, making use of models that jointly address the geometry and properties of the mesostructures. However, it remains unclear whether the integration of mechanical properties into the data set for generating multi-lattice interpolations is beneficial beyond geometry alone. To address this issue, this work implements and evaluates a hybrid geometry/property machine learning model for generating multi-lattice transition regions. We compare the results of this hybrid model to results obtained using a geometry-only model. Our research determined that incorporating physical properties decreased the number of variables to address in the latent space, and therefore improves the ability of generative models for developing transition regions of multi-lattice structures.

1. Introduction

Additive manufacturing has enabled more design freedom, however, these freedoms have forced designers to become more creative due to continuous pushes for optimization and material reduction. One approach to keep pace with these pushes is to utilize lattice structures, which can be used to reduce weight while maintaining performance [1,2]. By simply applying lattice structures to a design, weights have been shown to decrease by as much as 40% while maintaining overall strength [3]. This is one of many benefits that have sparked further interest in lattices, leading to the development of multi-lattice structures which are structures created using multiple types of unit cell topologies [4–7]. Multi-lattice design has become a major area of interest among researchers in additive manufacturing due to its versatility in terms of properties. Of note, multi-lattice structures exhibit better strength and stiffness properties than uniform structures of comparable density [6,8].

However, the success of multi-lattice structures is dependent on their ability to distribute stress evenly, as stress concentrations can cause part failures [6]. Due to the complex nature of this design problem, researchers have explored approaches that use machine learning to design multi-lattice transition regions [9–13]. These models often make use of both geometry and stiffness information in the training data to create generative models that are responsive to both. Stiffness is a physical property that is considered in most of the literature that analyzes the performance of lattices [2,14,15], as it is a primary physical descriptor of the performance of a lattice. While incorporating stiffness properties into these models has proven successful, the relative value of incorporating

physical properties into these generative models has not been evaluated [9–13]. In other words, it is unclear whether the added complexity and computational requirements incurred through the addition of stiffness information provide an associated increase in design performance to justify those costs.

This study aims to determine whether it is necessary to incorporate physical properties into variational autoencoders (VAEs) that use geometrically-defined latent spaces to optimize the physical performance of transition regions of multi-lattice structures. Specifically, we examine *stiffness continuity*, a measure of the change in stiffness across adjacent unit cells, as an indicator of the stiffness throughout a transition region. Our primary research questions address the unknowns regarding the role of physical properties in unit lattice cell latent spaces:

1. How does incorporating physical properties into VAEs affect the relationship between *geometric smoothness*, distance, and transition length in the latent space?
2. How does incorporating physical properties into VAEs affect the relationship between *stiffness continuity*, distance, and transition length in the latent space?

2. Methodology

This section outlines the architectures developed and the methods of testing used to evaluate our VAEs. Additionally, it will briefly outline our use and development of data for this work.

The data used was originally generated in prior work by Wang et al. and consists of 248,396 unique orthotropic microstructures, or unit cells [10,12,16]. In addition to this data, we developed a function for calculating the stiffness tensors of generated geometries based on a MATLAB topology optimization code [17]. This code was validated against those computed by Wang et al in prior work [10,12,16]. A subset of 10,000 random data points was used to develop architecture of the models without excessive training times. The individual pieces of data were binary arrays of size 50×50 that represented an individual unit cell's geometry. Additionally, each unit cell had a corresponding 3×3 stiffness tensor.

2.1 Architectures

This section outlines the VAE architectures used for investigating the research questions. Specifically, this entails two main architectures: (1) a geometry-only VAE based on prior work, and (2) a hybrid representation VAE that encodes both stiffness and geometry information. Both architectures in this work used identical training parameters: training using a batch size of 32 and an Adam optimizer [18], where 85% of the dataset was used for training. Additionally, training was terminated early if the loss failed to improve after 10 full epochs, where the loss term measures the difference between the reconstructed data and the original data.

The geometry VAE simply encodes and decodes the geometry using a standard VAE architecture (see Figure 1). The architecture of the model in Figure 1 was based on our previous works [19,20]. The purpose of this model is to serve as a baseline for comparison against the hybrid architecture.




Figure 1: Geometry VAE Architecture* where the input consists of only the unit cell geometry

The hybrid VAE (see Figure 2) has a similar encoder and decoder framework to the geometry VAE. However the encoded geometry is appended with a flattened stiffness matrix before calculating the mean and log variance. This architecture is intended to encode information relating to the stiffness tensor into the latent space. The decoder of this architecture is identical to the decoder of the geometry decoder architecture to enable a consistent comparison.




Figure 2: Hybrid VAE Architecture* where the input consists of unit cell geometry and the corresponding stiffness tensor

2.2 Evaluating Performance of Architecture

The geometry VAE and hybrid VAE must be evaluated in terms of both geometric smoothness and stiffness continuity. Prior work establishes a procedure for evaluating geometric smoothness and shows that it is related to the distance in latent space and transition length [19,20]. We use the same analysis procedure and geometry smoothness metric here. Specifically, distance is measured with respect to number of standard deviations, and transition length is simply the number of points in a transition region. We also introduce a new metric for evaluating stiffness continuity.

2.2.1 Evaluating Geometric Smoothness

The success of a latent space is dependent on the performance of the interpolations that can be produced from that latent space. When evaluating geometric transitions in a latent space, continuously and smoothly changing geometries are desired. To perform this evaluation, we utilize a smoothness metric that was developed in our previous work to evaluate the geometric smoothness of a 2D interpolation [19,20]. This metric calculates the gradients between multiple layers in an interpolation, essentially measuring the flow between each layer of pixels. The root mean squared error (RMSE) is calculated between the flows and then normalized to produce a

* Image created with <http://alexlenail.me/NN-SVG/AlexNet.html>

value of smoothness over an entire interpolation. This metric will be utilized in this work to evaluate the smoothness of interpolations among the various models. More details on the implementation of this smoothness evaluation are available in prior work by the authors [20].

2.2.2 Evaluating Stiffness Continuity

When evaluating the equivalent of geometric smoothness for stiffness tensors, the changes within a single tensor are not important as each element represents a unique piece of information. Therefore, the tensors are evaluated by comparing values only with the neighboring tensor. To achieve this, we implemented an RMSE framework to calculate a stiffness continuity value. First, the stiffness tensors in the interpolation were normalized with respect to the entirety of the training data. Then, the RMSE must be computed between each pair of values in the stiffness tensor using

$$RMSE_i = \sqrt{\frac{\sum_{j=1}^N (K_{i+1,j} - K_{i,j})^2}{N}} \quad (1)$$

where $RMSE_i$ denotes the root mean squared error between a pair of stiffness tensors at indices i and $i+1$, K is the normalized flattened stiffness tensor, j is the index that identifies the specific term in flattened stiffness tensor, and N is the number of terms in a single stiffness tensor (in this case, 9). Since the stiffness tensors are normalized before computing the RMSE, the maximum possible value is 1. This allows the direct evaluation of the continuity of the stiffness by averaging the values in Equation 1 using

$$C_K = (1 - avg(RMSE_i)) \cdot 100\% \quad (2)$$

where C_K is a value representing the continuity of stiffness over the entire transition region. Higher values indicate a smoother and more continuous transition in stiffness, while lower values indicate abrupt transitions.

3. Results

In this section we discuss results of the models in three main areas: (1) the reconstruction capabilities, (2) the test interpolations outlined in the methods, and (3) an ordinary least squares evaluation based on the interpolation results.

3.1 Model Reconstruction

Initially, we evaluate the performance of a machine learning model by examining the performance during training and testing (here using a mean squared error (MSE) loss and coefficient of determination). This indicates whether or not the model can appropriately reconstruct the desired data. If the model can reconstruct data, then an examination of the learned latent space may be informative. The results from the geometry VAE (see Figure 3) showed that geometry can be effectively reconstructed to 80% accuracy for testing data, and 87% accuracy for training data. This model serves as the baseline and represents the original model developed in previous work [19,20]. The decrease in accuracy from our previous work is likely due to the increase in dimensionality of the data, as well as the randomness of the data. The performance of the validation

data was likely extremely inconsistent due to the lack of similarity between all the data points. This is a matter to consider when evaluating the other model types, as this model serves as the baseline.




Figure 3: Geometry VAE: Plot of Loss and Coefficient of Determination (left) and Visualization of the Latent Space using PaCMAP Dimensionality Reduction (right)

The hybrid VAE was nearly identical to the geometry VAE, with matching performance for accuracy for both training and testing data (see Figure 4). This is a good indication that the model will perform well when reconstructing data, and the results in the evaluation section should be directly comparable.




Figure 4: Hybrid VAE: Plot of Loss and Coefficient of Determination (left) and Visualization of the Latent Space using PaCMAP Dimensionality Reduction (right)

In addition, Pairwise Controlled Manifold Approximation (PaCMAP) was used to visualize the latent space produced by each of the trained models (see Figure 3 and Figure 4) [21]. This approach has been shown in prior work to accurately provide a representation with balanced preservation of local and global features for engineering-relevant data [22]. There are distinct similarities in the learned embeddings, despite the addition of unique performance information. The overall shape is similar, consisting of approximately four lobes. In addition, the unit cells with lower average density are placed at the periphery of the latent space.

3.2 Interpolation Performance

As discussed previously, the evaluation technique for consistent results was traveling through latent space using standard deviations as the distance metric. This metric is discussed in further detail in prior work [19,20]. Each point on a plot represents an interpolation, with the average value of a metric (y-axis) over a distance in the latent space (x-axis). The points are labeled using color to denote the length of the transition region.

The results from Figure 5 (left) are consistent with the results seen in previous work [19,20]. As distance in the latent space increases, the smoothness decreases. Additionally, the higher transition lengths have higher smoothness values relative to one another.




Figure 5: Geometry VAE: Geometry Smoothness versus Number of Standard Deviations in the Latent Space (left) and Stiffness Continuity (C_K) vs Number of Standard Deviations in the Latent Space (right).

The predicted geometries from Figure 5 (left) were used to calculate the corresponding stiffness tensors, which were evaluated using the metrics described in the methods section to produce Figure 5 (right). This figure serves as the baseline of stiffness continuity for our hybrid geometry/property model. From visual observations, the stiffness continuity begins to plateau as the distance in the latent space increases. This indicates that the effect of distance in the latent space is limited to some degree.

The results from the hybrid VAE are consistent with the results from the geometry VAE, where the relationship between smoothness, distance, and transition length are comparable (see Figure 6 (left)). This is a desirable trait, as it indicates that the latent space has similar embeddings to the baseline model. Which means that we can directly compare the performance of the two models in terms of the stiffness continuity. However, the pattern of the results displays a much more aggressive plateau starting at 4 standard deviations between endpoints.




Figure 6: Hybrid VAE: Geometry Smoothness versus Number of Standard Deviations in the Latent Space (left) and Stiffness Continuity (C_K) vs Number of Standard Deviations in the Latent Space (right).

The predicted geometries from Figure 6 (left) were used to calculate the corresponding stiffness tensors, which were evaluated using the metrics described in the methods section to produce Figure 6 (right). The stiffness continuity of the data in Figure 6 (right) appears to plateau much sooner than the points in Figure 5 (right).

Visually, Figure 5 (left) and Figure 6 (left) are consistent with prior work, showing that smoothness has a clear relationship to distance and the length of the transition region. Smoothness appears to be positively correlated with the length of the transition region, and negatively correlated with the distance in the latent space between the endpoints. Figure 5 (right) and Figure 6 (right) show the relationship between stiffness continuity, distance in the latent space, and length of the transition region. These plots demonstrate that the effect of distance and transition length does not have as significant of an impact as they do on geometric smoothness. However, these values may be skewed arbitrarily high, given that symmetric geometries contain zeros in the stiffness tensor [23]. The next section serves to further explore the validity of these visual observations.

3.3 Evaluating Latent Space Relationships

This section will outline the relationships between the primary variables in the latent space. The independent variables are smoothness of geometry and continuity of stiffness properties of the respective geometries. The dependent variables are distance in the latent space and the length of the transition region. Using these variables, an ordinary least squares regression can be executed based on the results in the interpolation performance section.

Table 1 displays the ordinary least squares regression analysis of the results from Figure 5 (left), which evaluates the relationships between geometric smoothness and latent properties for the geometry model. Again, the geometry model serves as the baseline since it has architecture similar to that used in prior work [19,20]. Table 2 displays the ordinary least squares regression analysis of the results from Figure 5 (right), which evaluates the relationships between stiffness continuity and latent properties for the geometry model.

Table 1: Geometry VAE: Ordinary Least Squares Regression Analysis for Geometry Smoothness versus Number of Standard Deviations in the Latent Space

R-Squared = 0.958	Coefficient	Standard Error	p-value
Constant:	95.7815	3.742	< 0.0001
Number of Standard Deviations (Distance):	-8.2559	0.961	< 0.0001
Transition Length:	0.3930	0.346	0.276
Interaction Term:	0.2838	0.089	0.007

Table 2: Geometry VAE: Ordinary Least Squares Regression Analysis for Stiffness Continuity versus Number of Standard Deviations in the Latent Space

R-Squared = 0.767	Coefficient	Standard Error	p-value
Constant:	97.0878	3.482	< 0.0001
Number of Standard Deviations (Distance):	-2.9178	0.894	0.006
Transition Length:	0.1877	0.322	0.570
Interaction Term:	0.1392	0.083	0.115

The results from Table 1, are consistent with prior work, where the transition length alone does not have a significant relationship with geometric smoothness and distance has the most significant impact on geometric smoothness [19,20]. This is evident based on the p-values in Table 1. The high R-squared value is indicative that the variability in the geometric smoothness is almost fully described by the two independent variables, distance and transition length.

Based on the results of Table 1 and Table 2, the relationship between stiffness continuity and the properties of the geometry defined latent space are not as strong as the relationship with geometric smoothness and the properties of the latent space. This is based on the low R-squared value of the stiffness continuity of the geometry VAE in Table 2. An R-squared value of 76% indicates that 24% of the variability of the stiffness continuity is not accounted for by the two independent variables explored. This suggests that there are additional variables that affect the performance of stiffness continuity. Therefore, the true underlying relationship between stiffness continuity, distance, and transition length is unclear.

Table 3 displays the ordinary least squares regression analysis of the results from Figure 6 (left), which evaluates the relationships between geometric smoothness and latent properties for the hybrid model. Table 4 displays the ordinary least squares regression analysis of the results from Figure 6 (right), which evaluates the relationships between stiffness continuity and latent properties for the hybrid model.

Table 3: Hybrid VAE: Ordinary Least Squares Regression Analysis for Geometry Smoothness versus Number of Standard Deviations in the Latent Space

R-Squared = 0.896	Coefficient	Standard Error	p-value
Constant:	92.5313	5.709	< 0.0001

Number of Standard Deviations (Distance):	-7.1776	1.466	< 0.0001
Transition Length:	0.6077	0.529	0.269
Interaction Term:	0.2079	0.136	0.148

Table 4: Hybrid VAE: Ordinary Least Squares Regression Analysis for Stiffness Continuity versus Number of Standard Deviations in the Latent Space

R-Squared = 0.908	Coefficient	Standard Error	p-value
Constant:	101.4797	1.450	< 0.0001
Number of Standard Deviations (Distance):	-2.5201	0.372	< 0.0001
Transition Length:	-0.0682	0.134	0.619
Interaction Term:	0.1162	0.034	0.005

The results from Table 3 illustrate that incorporating physical properties into the model removed the effects of transition length on geometric smoothness. This conclusion is based on the p-values of the transition and interaction terms in Table 3. Alternatively, stiffness continuity in the hybrid model has a relationship with the transition length, based on the p-values from Table 4. However, the R-squared values of nearly 90% suggest that other variables account for approximately 10% of the variability of the geometric smoothness and stiffness continuity.

Overall, the results indicate that incorporating physical properties into the VAE changed the relationships with distance and transition length. The results suggest that incorporating physical properties into the latent space decreased the likelihood that transition length influences geometric smoothness. However, stiffness continuity demonstrated a greater likelihood that the combination of distance and transition length have an effect. As expected, none of the models showed a relationship with transition length alone. The results from the hybrid VAE indicate that its latent space is better suited to create multi-lattice transitions with continuous geometry, given that fewer variables affect its performance. However, none of the results relating the two models were statistically significant, thus further testing is warranted.

4. Conclusion

The development of multi-lattice transition regions is increasingly dominated by generative machine learning models. These models aim to produce multi-lattice transition regions with smooth geometric and physical properties. This work compares two approaches to constructing transition regions: a model that uses only geometric information and another hybrid model that merges geometry and property information. The first model, the geometry VAE, created a geometrically defined latent space only using the unit cell topologies. The second model, the hybrid VAE, utilized a combination of unit cell topologies and unit cell stiffnesses to define the latent space.

The geometry VAE performed much like our previous model, where geometric smoothness is influenced by distance and the combination of distance and transition length. Further testing is

necessary to fully explain the relationships between stiffness continuity and the latent space. The hybrid VAE exhibited very different relationships than the geometry VAE. The results of the ordinary least squares regression indicate that incorporating stiffness into the generative model decreased the effect of transition length with respect to geometric smoothness. This is a desirable trait as it indicates that the geometric smoothness will only be affected by the distance in the latent space. Alternatively, the continuity of stiffness is affected by the combination of transition length and distance when analyzing the hybrid model. However, there appears to be a set of additional variables that account for ~10% of the variability for both stiffness continuity and geometric smoothness.

When analyzing the reconstruction accuracy of the models, they can only perform at 80% accuracy, which is much lower than previous work. The drop in accuracy from our previous work is likely due to the increase in dimensionality of the data, from 28×28 to 50×50 . Although the accuracy is sufficient for the current analysis, further work should seek to train more accurate VAE models. In addition, there are many ways in which property information can be combined with geometric information in the training process of hybrid models. Therefore, future work should also explore the design freedom that is present in the design space of deep learning models.

Acknowledgements

This material is based upon work supported by the National Science Foundation through Grant No. CMMI-1825535. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors. We are also grateful to Yigitcan Comlek from the Integrated Design Automation Laboratory (IDEAL) at Northwestern University for generously sharing their data to support this work.

References

- [1] Hanks, B., Berthel, J., Frecker, M., and Simpson, T. W., 2020, “Mechanical Properties of Additively Manufactured Metal Lattice Structures: Data Review and Design Interface,” *Addit Manuf*, **35**.
- [2] Plocher, J., and Panesar, A., 2020, “Effect of Density and Unit Cell Size Grading on the Stiffness and Energy Absorption of Short Fibre-Reinforced Functionally Graded Lattice Structures,” *Addit Manuf*, **33**, p. 101171.
- [3] Kantareddy, S. N. R., Roh, B. M., Simpson, T. W., Joshi, S., Dickman, C., and Lehtihet, E. A., 2016, “Saving Weight with Metallic Lattice Structures: Design Challenges with Real-World Example,” *Solid Freeform Fabrication Symposium*.
- [4] Sanders, E. D., Pereira, A., and Paulino, G., 2021, “Optimal and Continuous Multilattice Embedding,” *Sci Adv*, **7**(16), p. 4838.
- [5] Sahariah, B. J., Namdeo, A., and Khanikar, P., 2022, “Composite-Inspired Multilattice Metamaterial Structure: An Auxetic Lattice Design with Improved Strength and Energy Absorption,” *Mater Today Commun*, **30**, p. 103159.

- [6] Kang, D., Park, S., Son, Y., Yeon, S., Kim, S. H., and Kim, I., 2019, “Multi-Lattice Inner Structures for High-Strength and Light-Weight in Metal Selective Laser Melting Process,” *Mater Des*, **175**, p. 107786.
- [7] Gok, M. G., 2021, “Creation and Finite-Element Analysis of Multi-Lattice Structure Design in Hip Stem Implant to Reduce the Stress-Shielding Effect,” *Materials: Design and Applications*, **236**, pp. 429–439.
- [8] Wang, C., Zhu, J. H., Zhang, W. H., Li, S. Y., and Kong, J., 2018, “Concurrent Topology Optimization Design of Structures and Non-Uniform Parameterized Lattice Microstructures,” *Structural and Multidisciplinary Optimization*, **58**, pp. 35–50.
- [9] Chan, Y.-C., Ahmed, F., Wang, L., and Chen, W., 2021, “METASET: Exploring Shape and Property Spaces for Data-Driven Metamaterials Design,” *Journal of Mechanical Design*, **143**(3).
- [10] Wang, L., Chan, Y.-C., Ahmed, F., Liu, Z., Zhu, P., and Chen, W., 2020, “Deep Generative Modeling for Mechanistic-Based Learning and Design of Metamaterial Systems,” *Comput Methods Appl Mech Eng*, **372**.
- [11] Wang, J., Chen, W. W., Da, D., Fuge, M., and Rai, R., 2022, “IH-GAN: A Conditional Generative Model for Implicit Surface-Based Inverse Design of Cellular Structures,” *Comput Methods Appl Mech Eng*, **396**, p. 115060.
- [12] Wang, L., Chan, Y.-C., Liu, Z., Zhu, P., and Chen, W., 2020, “Data-Driven Metamaterial Design with Laplace-Beltrami Spectrum as ‘Shape-DNA,’” *Structural and Multidisciplinary Optimization*.
- [13] Wang, L., Tao, S., Zhu, P., and Chen, W., 2021, “Data-Driven Topology Optimization With Multiclass Microstructures Using Latent Variable Gaussian Process.”
- [14] Niknam, H., and Akbarzadeh, A. H., 2020, “Graded Lattice Structures: Simultaneous Enhancement in Stiffness and Energy Absorption.”
- [15] Wang, Y., Zhang, L., Daynes, S., Zhang, H., Feih, S., and Wang, M. Y., 2018, “Design of Graded Lattice Structure with Optimized Mesostructures for Additive Manufacturing,” *Mater Des*, **142**, pp. 114–123.
- [16] Lee, D., Chan, Y.-C., Chen, W. W., Wang, L., van Beek, A., and Chen, W., 2022, “T-METASET: Task-Aware Acquisition of Metamaterial Datasets Through Diversity-Based Active Learning.”
- [17] Xia, L., and Breitkopf, P., 2015, “Design of Materials Using Topology Optimization and Energy-Based Homogenization Approach in Matlab Updated Element Density Variable,” **52**, pp. 1229–1241.
- [18] Kingma, D. P., and Lei Ba, J., 2015, “Adam: A Method for Stochastic Optimization,” *ICLR*.
- [19] Baldwin, M., Meisel, N. A., and McComb, C., 2023, “Smoothing the Rough Edges: Evaluating Automatically Generated Multi-Lattice Transitions,” *3D Print Addit Manuf*.
- [20] Baldwin, M., Meisel, N. A., and McComb, C., 2022, “A Data-Driven Approach for Multi-Lattice Transitions,” *Solid Freeform Fabrication Symposium*.

- [21] Wang, Y., Huang, H., Rudin, C., and Shaposhnik, Y., 2021, “Understanding How Dimension Reduction Tools Work: An Empirical Approach to Deciphering t-SNE, UMAP, TriMap, and PaCMAP for Data Visualization,” *Journal of Machine Learning Research*, pp. 1–73.
- [22] Agrawal, A., and McComb, C., 2022, “Comparing Strategies for Visualizing the High-Dimensional Exploration Behavior of CPS Design Agents,” *IEEE Workshop on Design Automation for CPS and IoT*.
- [23] Zou, W.-N., Tang, C.-X., and Lee, W.-H., 2013, “Identification of Symmetry Type of Linear Elastic Stiffness Tensor in an Arbitrarily Orientated Coordinate System,” *Int J Solids Struct*.