

# GANTL: Towards Practical and Real-Time Topology Optimization with Conditional GANs and Transfer Learning

**Mohammad Mahdi Behzadi**  
Computational Design Laboratory  
Department of Mechanical Engineering  
University of Connecticut  
Email: mohammad.behzadi@uconn.edu

**Horea T. Ilieș\***  
Computational Design Laboratory  
Department of Mechanical Engineering  
University of Connecticut  
Email: horea.ilies@uconn.edu

*Many machine learning methods have been recently developed to circumvent the high computational cost of the gradient-based topology optimization. These methods typically require extensive and costly datasets for training, have a difficult time generalizing to unseen boundary and loading conditions and to new domains, and do not take into consideration topological constraints of the predictions, which produces predictions with inconsistent topologies.*

*We present a deep learning method based on generative adversarial networks for generative design exploration. The proposed method combines the generative power of conditional GANs with the knowledge transfer capabilities of transfer learning methods to predict optimal topologies for unseen boundary conditions. We also show that the knowledge transfer capabilities embedded in the design of the proposed algorithm significantly reduces the size of the training dataset compared to the traditional deep learning neural or adversarial networks. Moreover, we formulate a topological loss function based on the bottleneck distance obtained from the persistent diagram of the structures and demonstrate a significant improvement in the topological connectivity of the predicted structures. We use numerous examples to explore the efficiency and accuracy of the proposed approach for both seen and unseen boundary conditions in 2D.*

## 1 Introduction

Topology optimization (TO) is an iterative procedure that outputs the optimum material distribution of a structure within a specified domain subject to prescribed boundary conditions and external loads [1, 2]. Typically, gradient-

based techniques combined with model parameter continuation are used to promote convergence to optimal solutions. The process often requires hundreds of iterations, each involving complete finite element solutions, which are time consuming [3]. For example, the SIMP solver presented in [4] needs 1.85s per iteration and hundreds of iterations to solve a simple MMB beam with resolution (300 x 100).

Over the past few years, various machine learning approaches have been proposed to improve the computational cost of gradient-based topology optimization, but the existing methods have significant limitations, which limit their utility in practice. Without a doubt, a practical machine-learning based topology optimization tool that could be used to explore the associated design space needs to: (1) be trainable with a relatively small training data set; (2) have a strong generalization ability to boundary and loading conditions as well as design domains that are not part of the training data set; and (3) enforce the topological connectivity of the predicted structure.

For now, the proposed machine learning methods that make pixel-wise predictions of the optimum structures need extensive training datasets, and especially those with generative characteristics, and the generation of these large datasets is very expensive from a computational standpoint. Moreover, only a few of the proposed algorithms have been shown to handle boundary conditions that are not part of the training data, but in all these cases, the unseen boundary conditions are very close to those used in the training data. At the same time, and without exception, the structural predictions made by all these methods have not only poor but also inconsistent topologies compared to the ground truth. The latter limitation, which increases in severity with the increase in resolu-

---

\*Corresponding author.

tion, is due to the fact that none of these methods takes into consideration the topology of the prediction.

In our recent paper [5], we introduced a transfer learning approach based on Convolutional Neural Networks (CNN) that was capable of handling high-resolution 2D & 3D design domains of various shapes and topologies and supporting real-time design space explorations. We showed that the knowledge transfer capabilities of that network significantly reduced the size of the training dataset compared to the traditional deep learning neural networks. Consequently, we showed that the proposed architecture was capable of handling boundary conditions that were unseen to the source network by only fine-tuning the corresponding target network with a much smaller training data set. While the method demonstrated the knowledge transfer capabilities of transfer learning applied to topology optimization, it requires fine-tuning of the target network for unseen domains and boundary conditions.

Generative Adversarial Networks (GAN) are one of the most promising developments in machine learning because they have the ability to learn low and high-level patterns in datasets and use them to generate new data that bears a (statistical) resemblance to the original dataset. As detailed in section 2, GANs have been used to generate realistic but artificially generated images, and these same generative properties prompted GANs to be used in making predictions of optimal topologies. However, standard GANs need very large training data sets to be able to generate the new images, and their training can be notoriously difficult due to mode collapses and diminished gradients.

In this paper, we combine the generative power of conditional GANs with the knowledge transfer capabilities of transfer learning methods to predict optimal topologies for unseen boundary conditions. We show that the knowledge transfer capabilities embedded in the design of the proposed algorithm significantly reduces the size of the training dataset compared to the traditional deep learning neural or adversarial networks. In fact, we show that using the same amount of data as in [5], the proposed GANTL can generate the structures with new boundary conditions that are not included in the target model training data. Moreover, we formulate a topological loss function based on the bottleneck distance obtained from the persistent diagram of the structures and demonstrate a significant improvement in the topological connectivity of the predicted structures. We use numerous examples to explore the efficiency and accuracy of the proposed approach for both seen and unseen boundary conditions. To the best of our knowledge, this is the first time that the topological loss function is combined with GAN to tackle the connectivity issue of the predicted optimum structures.

## 2 Background

The significant computational cost of the traditional gradient-based TO prompted the investigation of various alternatives to the highly iterative process. For instance, parallelized algorithms run on high-performance computing platforms with a multigrid preconditioner was discussed in [6] to

optimize ‘ultra large’ scale shell models comprised of over 11M elements in 17 hours using 800 compute cores. The parallelized level set method [7] solves a  $96 \times 48 \times 24$  cantilever beam in 45 minutes on a desktop computer with 4 Intel i7-6700 CPU cores and 16 GB memory. Moreover, the work described in [8] employs the Portable and Extendable Toolkit for Scientific Computing (PETSc) to produce the optimal topology starting from a cantilever beam with  $480 \times 240 \times 240$  elements in 14 hours using 144 CPU cores.

The high computational cost of the gradient-based TO algorithms prompted the development of a number of machine learning methods that can make predictions almost instantaneously once trained. Recently proposed CNN-based algorithms can predict the optimum structure for low resolution 2D domains ( $50 \times 50$ , and  $40 \times 40$ , respectively) [9, 10] and low resolution 3D domain ( $24 \times 12 \times 12$ ) [11].

Several recent papers in the realm of topology optimization were inspired by the ability of GANs to synthesize images of human faces that don’t belong to a real person [12], which learn the patterns and the distribution of the input data - see also section 3.2. Despite this, many TO approaches that use GANs do not explore this important aspect. For example, conditional Wasserstein generative adversarial networks (CWGAN), which minimize an approximation of the Earth Mover’s distance [13], were used in [14] for 2D TO. That method was trained on 3,024 cantilever beams ( $120 \times 120$ ) with design variables that include the volume fraction, penalty factor, and radius of the filter. They concatenate their design variables with the noise vector and feed them to their CWGAN to generate the cantilever beam that corresponds to the design variables. A two-stage hierarchical prediction-refinement GAN-based framework [15] is used to predict the low resolution near-optimal structure and its corresponding refined structure in high resolution. Their training data set contains 9,900 pairs of low-resolution ( $40 \times 40$ ) and high-resolution ( $160 \times 160$ ) data, and the method achieves a 9.1 % MSE for high resolution predictions. Similarly, [16] describes a two-stage refinement cGAN to predict the high resolution optimum structure, which is trained with 64,000 low ( $32 \times 32$ ) resolution and their high ( $128 \times 128$ ) resolution counterparts.

However, all these algorithms require large datasets for training, lack generalization abilities, and do not have a mechanism to enforce the topology of their predictions.

Two recent papers do provide some evidence of generalizability. Specifically, the method described in [17] uses 64,000 cantilever beam to train a deep CNN to predict the optimum structure for simple low-resolution ( $40 \times 80$ ) 2D domains. Their network uses the displacement and the strain field as inputs, and the authors showed that their method could generate a solution for the simply supported and the two-span continuous boundary conditions that are not included in the training data set. Moreover, a conditional GAN, named TopologyGAN, is described in [18], which used 49,078 2D beams ( $64 \times 128$ ) with 42 different displacement boundary conditions to train for their model. Their results display a nearly three times reduction in the MSE on test problems involving four new boundary conditions that

are not part of the training data set comparing to the baseline cGAN. Although both methods have shown some success in making predictions for domains with unseen boundary conditions, they do require large training datasets, and they do not consider the topology of the predictions that they make, often producing severely branched (tree-like) predictions.

To date, the only known method to produce well-connected predictions involves variants of the gradient-based TO methods, which trade computational efficiency for an improved connectedness. For instance, [19], and [20] proposed an iterative topology optimization method using neural networks for single material (TOuNN) and multi-material structures (MM-TOuNN) for direct execution of TO. In this work, the NN weights and bias are used to parameterize the density function, thus disconnecting the density from the finite element mesh and obtain a crisp and differentiable material interface and boundary. The NN is used to predict the density values at each iteration, and part of the associated sensitivity analysis was computed via backpropagation. Importantly, the iterations performed by this approach are still driven by FEA. Clearly, this variation of the gradient-based TO is flexible in terms of the domains and the boundary conditions it can tackle, since it continues to be driven by FEA, but remains computationally expensive and seems to lose some of SIMP's ability to reproduce the solutions generated by SIMP. For example, some examples from [19] require twice as much time per iteration compared to SIMP [4].

In this paper, we take a different approach and integrate notions from persistence homology [21, 22] to consider the connectedness of the predicted structure. As argued above, such topological properties represent one of the key criteria for making these machine learning algorithms practical. However, generating predictions with topological guarantees is certainly not specific to the field of topology optimization. In fact, the image processing community has been interested in techniques that capture the topological properties of image predictions. For example, [23] uses feature maps from pre-trained very deep convolutional networks, dubbed as VGG, obtained for the ground truth and the predicted image to construct a topology-aware loss function. They showed that by minimizing this loss function, one can significantly improve the topological features of the predicted image. Another way to improve the topological feature of the predicted images is using the persistent homology to capture the topological information of the images in order to construct the topological loss function. The work described in [24] used the modified Wasserstein distance between the persistent diagram of the ground truth and the prediction to construct the loss function. This method can achieve much better performance in terms of the Betti number<sup>1</sup> error between the prediction and ground truth than any other prior method.

## 2.1 Contributions and Outline

The data-driven topology optimization method proposed in this paper combines the generative power of conditional

GANs with the knowledge transfer capabilities of transfer learning methods. We show that this combination results in a practical TO tool based for design space exploration that: (1) requires a much smaller training data set than any other data-driven TO method; (2) has appealing generalization abilities for unseen boundary conditions; (3) displays a significant improvement in the topological connectivity of the predicted structures due to the new loss function that uses key concepts from persistence homology.

The rest of the paper is organized as follows. Section 3 contains the preliminaries, including formulation, GANTL architecture, data generation, and metrics used for model evaluation. We demonstrate the generality and flexibility of our approach by exploring various 2D examples with different resolutions, boundary conditions, and design spaces, including boundary conditions that are *unseen* the network. These examples show that the proposed method supports real-time generative design space explorations for both *seen* and *unseen* boundary condition and the topological improvements of our predictions due to the novel topology-aware loss function. Finally, Section 5 summarizes the key advantages and limitations of the proposed method and its potential applications.

## 3 Problem Formulation

### 3.1 Minimum Compliance TO problem

Topology optimization seeks the material distribution  $\rho(\mathbf{x})$  that minimizes an objective function  $f(\Omega, \rho)$  that is subjected to various constraints  $g_i < 0$ . A common objective function in structural TO is the global structural compliance [26, 27, 28], but local objective functions, such as those depending on the von Mises stresses, have been proposed [29, 30]. The minimization is usually carried out by solving a finite element analysis. The well known SIMP approach [31] associates each element with a density variable  $\rho_e \in [0, 1]$ , where 0 corresponds to an empty element, and 1 to an element completely filled with material. The corresponding optimization problem can be formulated as

$$\min : f(\Omega, \rho) = \mathbf{U}^t \mathbf{K} \mathbf{U} = \sum E_e(\rho_e) \mathbf{u}_e^t \mathbf{k}_e^0 \mathbf{u}_e \quad (1)$$

$$s.t : \mathbf{K} \mathbf{U} = \mathbf{F} \quad (2)$$

$$\sum \rho_e v_e \leq V_{max} \quad (3)$$

$$0 \leq \rho \leq 1 \quad (4)$$

where  $\Omega$  is the design domain;  $\rho$  is the density vector;  $\mathbf{U}$  and  $\mathbf{F}$  are the global displacement and force vectors, respectively;  $\mathbf{K}$  is the global stiffness matrix;  $\mathbf{u}_e$  is the element displacement vector;  $\mathbf{k}_e^0$  is the element stiffness matrix for an element with unit Young's modulus;  $v_e$  and  $\rho_e$  are the element volume and density of element  $e$ , respectively; and  $V_{max}$  is the volume upper bound.  $E_e(\rho_e)$  is the element's Young's modulus determined by the element density  $\rho_e$ :

$$E_e(\rho_e) = E_{min} + \rho_e^p (E_0 - E_{min}), \quad \rho_e \in [0, 1] \quad (5)$$

<sup>1</sup>A reasonably accessible description of homology groups and Betti numbers can be found in [25].

In equation (5),  $E_0$  represents the stiffness of the material,  $E_{min}$  is a very small number, and  $p$  is a penalty factor [2,32].

### 3.2 Generative Adversarial Networks (GANs)

Generative Adversarial Networks are one of the most exciting and most promising innovations in machine learning that learn the patterns and statistical distributions of the input data in order to synthesize new samples [33] that do not exist in the original data. The synthesizing ability of GANs has been used in semantic image editing [34], data augmentation [35], and style transfer [36]. Furthermore, one can use the patterns learned by GANs for other data-intensive tasks, such as classification [37].

GANs consist of a generator and a discriminator, which are both neural networks that work together. During training, the generator learns to generate plausible data, and the discriminator learns to distinguish the generator's fake data from real data and penalizes the generator for producing implausible results. As training progresses, the generator becomes better at generating plausible data, and the discriminator has a harder time distinguishing between real and fake data. Using GANs in practice must overcome several challenges, including those induced by vanishing gradients, mode collapse, and convergence failures. A good and recent review on algorithms, theory, and applications of GANs can be found in [38].

### Transferring GANs

Deep Neural Networks achieve promising results on image classification problems as long as large datasets are available for training. For example, Xception [39] is trained on 350M images. However, many practical use cases do not have the required abundance of data needed for training of these networks. This is where transfer learning can come in.

As the name implies, transfer learning reuses the knowledge learned by a source network to improve the performance of a target network trained on a different but related task. Transfer learning can significantly increase the network performance for situations when the data is scarce [40].

Recent results of applying transfer learning in conjunction with generative adversarial networks [41, 42, 43] show that by transferring pre-trained knowledge, one can effectively reduce the amount of training data in the context of image generation. In addition, this knowledge increases the GAN convergence rate, and improves the convergence failures due to the well-known mode collapse [33].

### 3.3 Proposed GAN Architecture

As mentioned above, GANs consist of a generator and a discriminator, and in our case both contain a source as well as a target neural networks. The overall architecture of the source generator is inspired by [44] and separates the subnets handling the feature extraction from the synthesis. The synthesis subnet is responsible for generating the optimum structure from a noise vector using the feature maps obtained by the feature extraction subnet. The latter is inspired by

SE-Res Net introduced in [45], but we remove the addition layer [46] from the SE-Res block and use a concatenation layer instead to increase the information flow between the layers and reuse the extracted features [47].

The architecture of the source generator is shown in Figure 1, and the corresponding input data is described in the next section. The synthesis subnet uses a series of synthesis blocks described in Fig.2 that helps to combine the feature maps of boundary conditions with the previous layer and determine the amount of the feature information that should be retained for the following blocks. The SE-Concat block contains a squeeze and excitation (SE) block [45], which improves channel interdependencies and determines the important features of the input, as shown in Figure. 2.

The source discriminator is built using convolutional layers and a dense layer to decide whether an optimum structure output by the generator is fake, as shown in Fig.3. The discriminator's inputs are a predicted structure as well as its corresponding boundary information, and the output is 0 for the fake structures and 1 for the real structures.

The target generator is built by adding a convolutional transpose layer and four convolutional layers on top of the source generator. Before using the source generator on the target GAN, we remove the last two-layers of the source generator. Furthermore, we add a downsampling function at the front of the target generator to reduce the high resolution input dimension to the low-resolution input required by the pre-trained generator. The architecture of the target discriminator is very similar with that of the source discriminator, with the exception of an additional convolutional layer and of a convolutional transpose layer tasked with matching the output resolution. The structure of the target generator and discriminator are shown in Fig.4.

### 3.4 Dataset

We used a SIMP-based topology optimization code [4] to generate our training and test cases. The code is modified to automate the data generation for different domains and boundary conditions. The inputs are the voxelized domain geometry, volume fraction, filter radius, load, and displacement boundary conditions. The volume fraction and filter radius are prescribed to 0.5 and 1.5, respectively.

The magnitudes of the force components were sampled using uniform random sampling within the range  $[-100, 100]$ N. The location of the external load is selected according to a uniform random sampling within prescribed ranges along the coordinate axes. For example, the force components  $F_x$  and  $F_y$  for a 2D beam domain are applied within the range  $[\frac{b_x}{2}, b_x]$  and  $[0, b_y]$ , respectively, where  $b_x$  and  $b_y$  are the beam dimensions in the  $x$  and  $y$  directions. We also used a discrete random sampling to select one of the displacement boundary conditions shown in fig. 5 (a) and (b).

Our inputs were captured in five channels for the 2D cases represented as matrices:

1. First channel: Initial density value for each voxel.

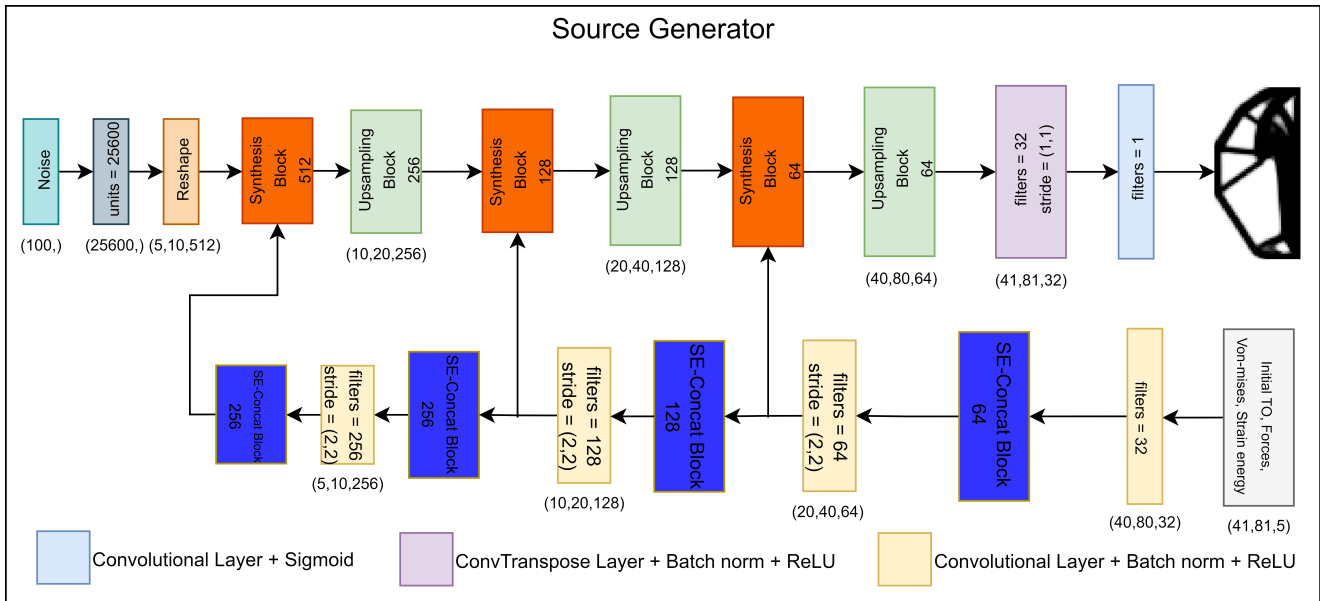


Fig. 1: The architecture of the source generator. The numbers below the blocks show the size of the out put of blocks.

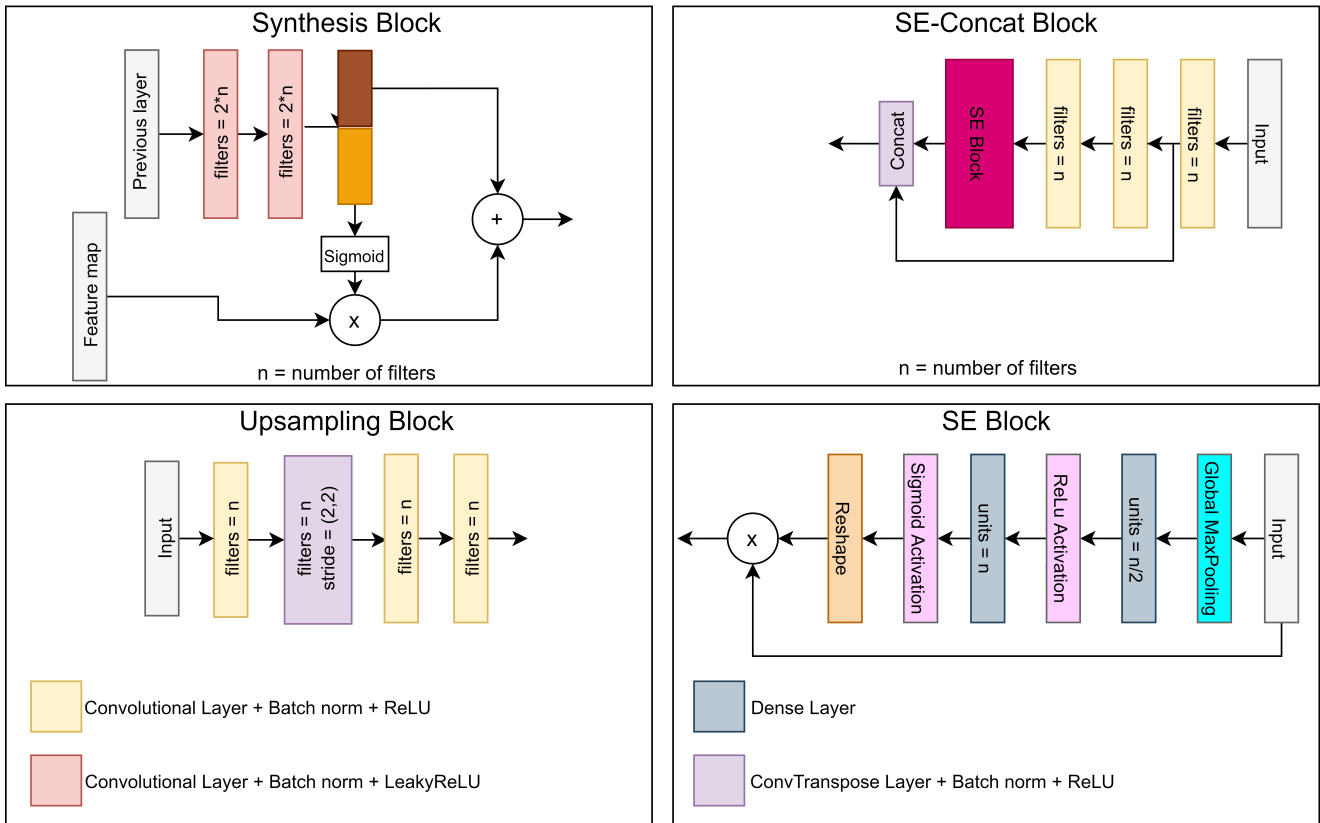


Fig. 2: The architecture of the synthesis block, upsampling block, SE-Res block and SE block.

2. Second channel: Force component in the  $x$  direction at each voxel,
3. Third channel: Force component in the  $y$  direction at each voxel,
4. Fourth channel: The von Mises stress of the initial domain defined as:

$$\sigma_v = \sqrt{\sigma_x^2 + \sigma_y^2 - \sigma_x \sigma_y + 3\sigma_{xy}^2}$$

5. Fifth channel: The strain energy density function of the initial domain defined as:

$$W = \frac{\sigma_x \epsilon_x + \sigma_y \epsilon_y + 2\sigma_{xy} \epsilon_{xy}}{2}$$

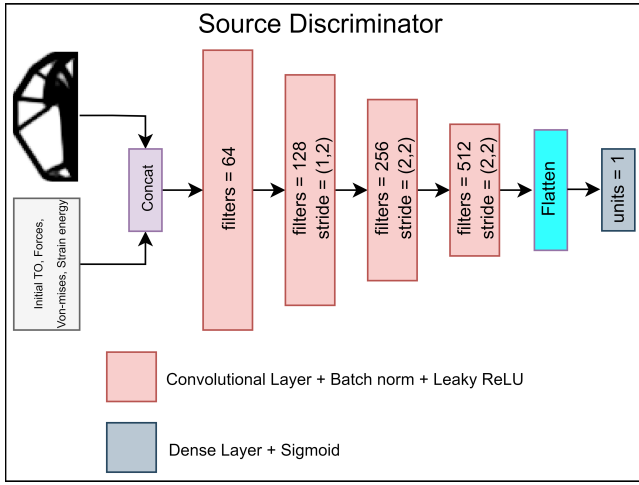


Fig. 3: The structure of the source discriminator.

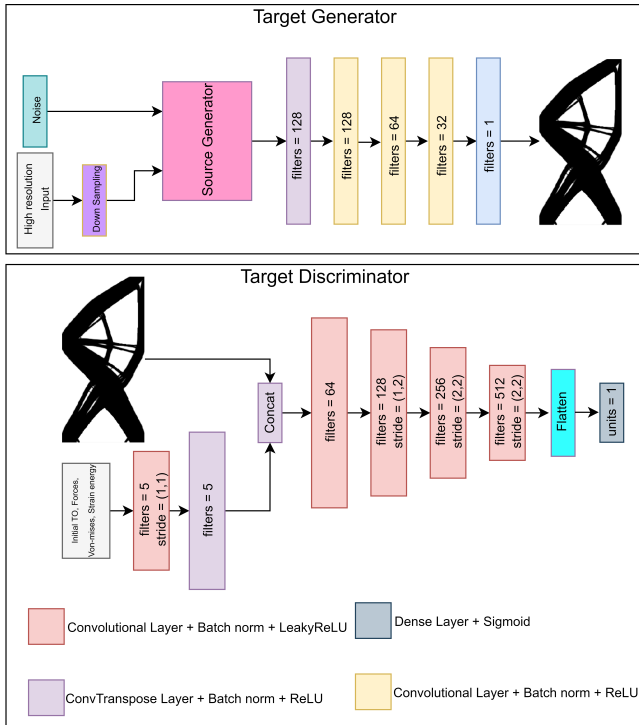


Fig. 4: The architecture of the target generator and discriminator.

where  $\sigma_i$  and  $\epsilon_i$  are stress and strain in direction  $i$ , respectively. The stress and strain corresponding to the initial domain were computed with SolidsPy library [48] in Python.

Based on the information given above, a low resolution dataset (11,000 cases) and a much smaller high resolution dataset (1500 data) were generated for the source and target GAN.

### 3.5 Training Process and Performance Metrics

The generator and discriminator of the GAN should be trained simultaneously, and we followed the following training process: (1) trained the discriminator for one batch, (2) trained the generator for one batch, (3) repeat step (1) and (2) to finish the batches and go to the next epoch. In the first step, the algorithm generates the labels for the real structures and the fake structures, and then it trains the discriminator on the real structures with label one and the fake structures generated by the generator with label 0. In the next step, the generator is trained on the ground truth data to generate accurate pictures. As the generator getting better at generating the pictures, the discriminator is getting worse in distinguishing between the real and the fake data, and when the accuracy of the discriminator for classifying the data is less than 50%, the training process finishes.

Our GAN is trained on the domain with boundary conditions shown in fig. 5(a) and (b). The loss function for both generator and the discriminator is binary cross-entropy, and Adam is used as the optimizer. We employed the same three criteria that we used in [5] to evaluate the quality of the GAN predictions:

1. *Compliance Error*, which is defined as:

$$Comp.Error = \frac{|\sum_{j=1}^M \sum_{i=1}^N (c_{pred}^{ij} - c_{true}^{ij})|}{\sum_{j=1}^M \sum_{i=1}^N c_{true}^{ij}}$$

where  $N$  and  $M$  are the number of rows and columns,  $c_{pred}^{ij}$  and  $c_{true}^{ij}$  are the predicted and reference values of compliance located in the  $i^{th}$  row and  $j^{th}$  column, respectively.

2. *Mean Squared Error (MSE)*, which is defined as:

$$MSE = \frac{\sum_{j=1}^M \sum_{i=1}^N (y_{pred}^{ij} - y_{true}^{ij})^2}{N \cdot M}$$

where  $y_{pred}^{ij}$  and  $y_{true}^{ij}$  are the predicted and reference values of element located in the  $i^{th}$  row and  $j^{th}$  column, respectively.

3. *Binary Accuracy (BA)*, which is defined as:

$$BA = \frac{TP + TN}{N}$$

where  $TP$  (True Positive) is the number of elements correctly predicted as 1,  $TN$  (True Negative) is the number of elements correctly predicted as 0, and  $N$  is again the total number of elements. Calculating BA is preceded by rounding the element values to the closest integer (0 or 1).

Additional details about these metrics are provided in [5]. All experiments were performed on the UConn HPC facility running Red Hat RHEL7 operating system.




Fig. 5: Boundary conditions: (a) and (b) are used for training and testing; (c) – (e) are unseen BCs only used for testing.

## 4 Experiments and Evaluation

### 4.1 Evaluation with Traditional Cross-Entropy Loss

We used a freely available Matlab Code [4] to generate 11,000 low-resolution cases ( $40 \times 80$ ) used to train the generator and discriminator of the source GAN and 1500 high-resolution cases to train the generator and discriminator of the target GAN. We augmented the training data by mirroring each SIMP-generated case with respect to the  $x$  and  $y$  axes.

**Training:** The source GAN is trained only once, and the trained source network was reused for various target models. For training our source and target GANs, we used two randomly selected boundary conditions corresponding to those shown in Figure 5 (a) & (b). All cases used for training had a single externally applied force with randomized location, orientation, and magnitude within the prescribed ranges.

**Testing:** Given the training mentioned above, the remaining 3 boundary conditions illustrated in Figures 5 (c) – (e) are therefore unseen to both the source and target GANs and were used exclusively for testing. Also, observe that training was completed with only one external force applied, while we tested cases with 1, 2, 5, and 10 external forces. We extensively tested the performance of our predictions and the generalization abilities of our GANTL framework in the following testing scenarios:

Table 1: Testing scenarios for both seen BCs (Fig. 5 (a), (b)) and unseen (Fig. 5 (c–e)). Training was performed with one external force, while testing was completed for 1, 2, 5, and 10 randomized external forces.

Scenario #	BCs from Figure #	# of external forces	Predictions in Figure #	Data in Table #
1	5 (a), (b)	1	6	2
2	5 (c–e)	1	7	3
3	5 (a), (b)	2	8	4
4	5 (c–e)	2	9	5
5	5 (a), (b)	5 & 10	10	-

### 4.2 Discussion

Figures 6–10 show the side by side comparisons for all testing scenarios summarized in Table 1, and the corresponding evaluation data captured by the performance metrics described in section 3.5 was collected in tables 2–5. The low resolution ( $40 \times 80$ ) predictions in all these figures show the prediction of our source GANs and the comparison with the equivalent low resolution ground truth provided by SIMP.

The performance data shows that our network achieves a better prediction performance compared to the state-of-the-art deep learning-based methods.

Specifically, for scenario # 1 (*seen* BCs with one randomly selected external force), the average MSE is 2.12%, and the average binary accuracy is 96.72 with a 1% average compliance error and a 2.5% standard deviation. This implies that our network captures very well the changes in the topological patterns induced by a *continuous* variation in the location, direction, and magnitude of the external force.

For scenario # 2 (*unseen* BCs with one randomly selected external force), the average MSE is 9%, and the average binary accuracy is 89%, with an average compliance error of 7.7% and standard deviation of 11.7%. Given the fact that the GAN was trained on only two different boundary conditions illustrated in 5 (a), (b), this test scenario indicates that the GANTL network does capture some high-level dependence of the topological patterns on the discrete boundary conditions. We note that the errors are primarily coming from incorrect predictions of the thin members of the ground truth, although this is not surprising given the fact that the network learned the topological patterns from boundary conditions that are dissimilar from those considered in this scenario.

The next two test scenarios use an unseen loading case, i.e., two randomly selected external forces and both seen and unseen boundary conditions, as described above. Recall that all training cases considered only one externally applied force. We observe that the performance metrics for both scenarios # 3 & 4 from Table 1, which include partially or totally unseen information, are very similar: average MSE – 7% vs. 11%, binary accuracy – 90% vs 87%, average compliance error – 10.5% vs 11.5%, and standard deviation – 16.2% vs 15.7%.

In order to evaluate the performance of our predictions as the number of external forces increases, we included test cases with five and ten external forces for domains with a  $200 \times 400$  resolution and boundary conditions displayed in Figure 5(a) and (b). The side-by-side comparison between

our predictions and the SIMP ground truths is shown in Figure 10. Although the average binary accuracy of the 250 test cases having five and ten external forces is around 84% and 82%, respectively, so only slightly lower than the data for scenarios # 3 & 4, these figures indicate that as the number of external forces increases, the confidence of our network in predicting the thin members decreases. This behavior follows the known limitations of extrapolating approximations of non-linear functions, including the extrapolation of the neural networks learning outside of the support of the training distribution [49]. In some sense, this is conceptually similar with teaching a student the basic concepts of derivation and integration, then ask the student to solve differential equations. The student may solve simple ODEs, but would have a much harder time with complex ODEs or PDEs.

Collectively, the test scenarios # 1-5 demonstrate that the proposed network is capable of high quality predictions of the overall optimal structure and clearly illustrate the generalization ability of our proposed GANTL. These test scenarios also indicate that the GANTL network has a better performance in predicting larger members of the optimal structure. At the same time, we observe that our network has a somewhat limited confidence in predicting the thinnest members of the optimal structure, and particularly for completely unseen boundary and loading conditions, and that this confidence decreases as the number of forces increases. This, however, is an expected behavior that is consistent with the observations from other machine learning-based predictors, and shows that the topological patterns learned by GANTL from the single external force used in training are changing as the number of external forces increases. Nevertheless, our results also suggest that constructing a training dataset that includes a reasonable sample of the practical loading cases used in practice would rapidly increase the quality of the network’s predictions for design explorations.

### 4.3 A Topology-Aware Loss Function

One common challenge of all published machine learning approaches in topology optimization, which exclusively use for now an image-based approach, is in predicting connected structural members. In fact, all these machine learning-based TO approaches simply fail to consider the topological connectivity of their predictions. Notably, even the best performing such methods attempt to improve the quality of the predictions by merely diversifying the training set. In the context of TO, this strategy increases the number of required training cases and, therefore, the associated computational cost of generating the dataset and of training the network, which may also impact the convergence of the training process.

Motivated by this observation, we propose a promising alternative for stimulating the network to produce connected structural members. Specifically, we construct and evaluate a custom loss function that is based on concepts from persistence homology [21, 22], and specifically on topological persistence diagrams. Informally, these persistence diagrams extract the ‘birth’ and the ‘death’ of the topological features

(e.g., connected components) in the structure as the structure evolves according to predefined filters. By correlating the persistence diagram of the ground truth with that of the prediction and by including this correlation in the custom loss function, we push the network to make predictions that are topologically similar to the ground truth.

For these experiments, we used the bottleneck distance between persistence diagrams, which can be thought of as the shortest ‘distance’  $b$  for which there exist an optimal matching between points of two given persistence diagrams such that any pair of matched points are at distance at most  $b$  [21, 50].

We add the custom loss function to the binary cross-entropy and use the total loss as the GAN loss function:

$$L(y_t, y_p) = L_{bce}(y_t, y_p) + \lambda L_{topology}(y_t, y_p) \quad (6)$$

where  $y_t$  is the ground truth structure,  $y_p$  is the predicted structure,  $L_{bce}$  is the binary cross-entropy,  $L_{topology}$  is the bottleneck distance between the persistent diagrams, and  $\lambda$  controls the weight of the topological loss which can be in  $(0, 1]$ . In our experiments, we chose  $\lambda$  to be 0.1. We extract the persistence diagrams of structures by using Ripser [51] and the bottleneck distance is computed with the GUDHI library [52].

We fine-tuned the target GAN for our highest resolution (200 x 400) domains, which also produced predictions with the highest number of disconnected members in the experiments detailed above, with the new topology-aware loss function. Prior to displaying the predicted structures, we round up or down the density values to the closest integer, either 0 or 1. We tested the performance of the proposed topology-aware loss function for both seen and unseen displacement boundary conditions.

As shown in Figure 11, the proposed topology-aware loss function can significantly reduce the number of disconnected members for both seen and unseen boundary conditions, and also eliminates many small holes in our predictions.

## 5 Conclusions

We proposed a highly efficient and accurate generative design tool that combines the power of generative adversarial networks with the efficiency provided by the knowledge reuse of transfer learning. The proposed method uses 11,000 low resolution cases to train a source generator, whose knowledge is transferred to a target network. Consequently, the training (fine-tuning) of the target network for high-resolution domains needs a much smaller dataset (1,500 cases). The fact that we transferred the knowledge learned from the low resolution to the high resolution cases allowed us to significantly reduce the number of high resolution cases without losing the accuracy of the predictions. We presented a discussion of the associated computational costs of generating the training cases in [5] so we do not repeat it here.



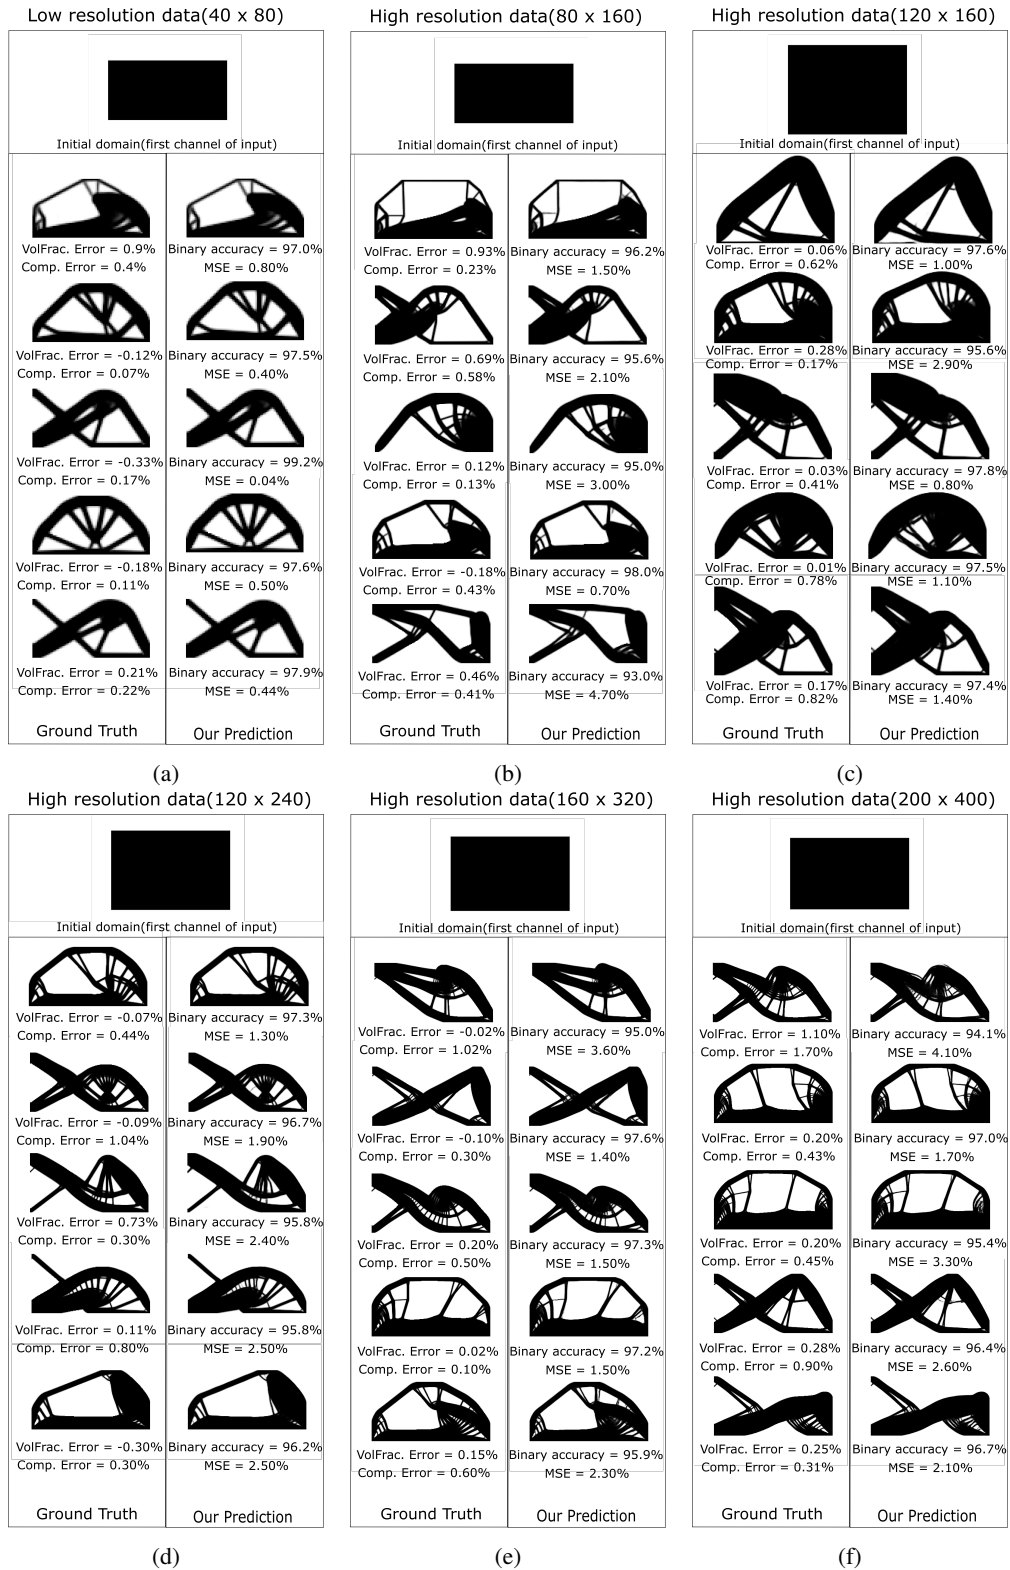


Fig. 6: Predictions vs. ground Truth for *seen* boundary conditions and one external force: (a) the prediction of our source GAN alone; (b-f) optimal structures output by the fine tuned target model. The corresponding quality metrics are presented in Table 2.

Table 2: MSE, binary accuracy, compliance error and standard deviation relative to SIMP for the structures shown in 6 with *seen* BCs and one external force.

Design Domain	Resolution	Number of test cases	MSE	Binary Accuracy	Compliance Error	Compliance Error Std.
from Fig. 6a (source network)	40 x 80	7473	0.48%	98.51%	0.32%	1.7%
from Fig. 6b	80 x 160	350	1.92%	96.66%	1.03%	3.3%
from Fig. 6c	120 x 160	463	2.41%	96.36%	1.08%	3.8%
from Fig. 6d	120 x 240	388	2.28%	96.56%	0.98%	1.9%
from Fig. 6e	160 x 320	350	2.90%	96.06%	1.14%	2.6%
from Fig. 6f	200 x 400	350	2.77%	96.20%	1.45%	1.9%
<i>Average</i>			2.12%	96.72%	1.00%	2.5%

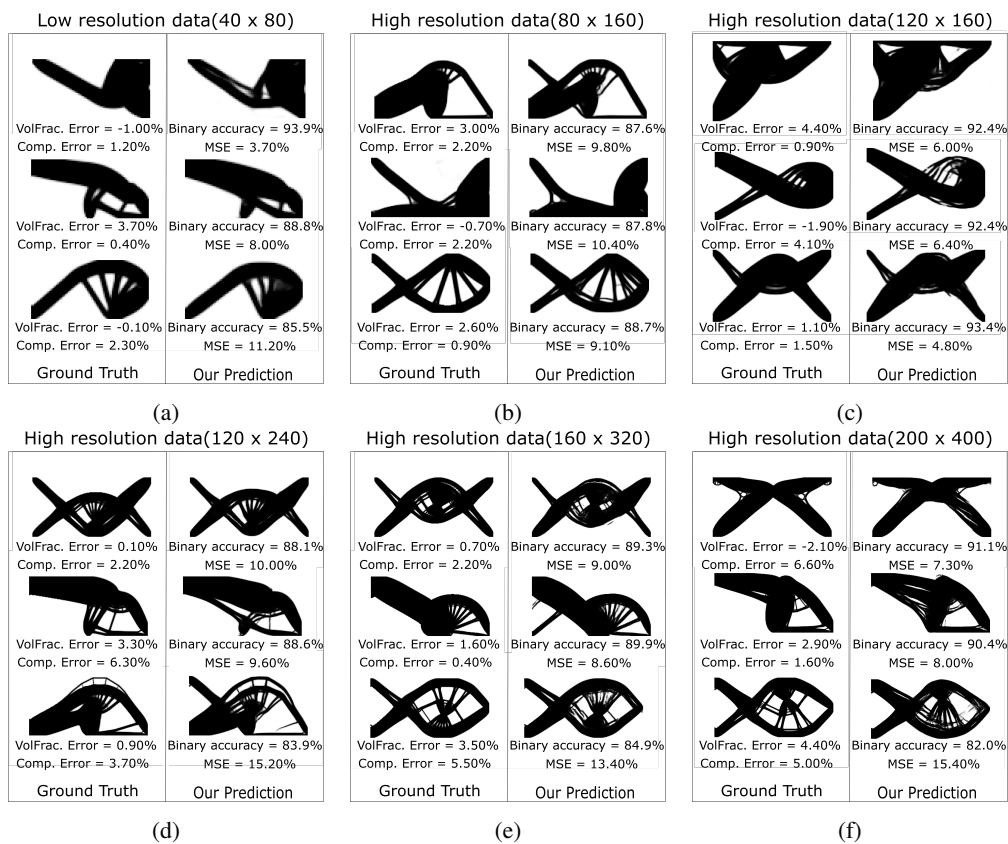


Fig. 7: Predictions vs. ground Truth for *unseen* boundary conditions and one external force: (a) the prediction of our source GAN alone; (b-f) optimal structures output by the fine tuned target model. The corresponding quality metrics are presented in Table 3.

We produced numerous examples to show that the proposed GANTL approach can generate highly accurate predictions for seen boundary conditions with a relatively small performance decrease for domains that use unseen boundary conditions that were not part of the training data. In addition, while all training cases had one single external force applied to the domain, our test data also included external loading with 2, 5, and 10 external loads.

An analysis of these examples points to the fact that the confidence of our network in predicting the thin members decreases as the number of forces increases. This behavior is consistent with the behavior observed by all image-based machine learning algorithms, including all known approaches that use these algorithms in topology optimization, and with the limitations of extrapolating approximations of non-linear functions outside of their sampled do-

Table 3: MSE, binary accuracy, compliance error and standard deviation relative to SIMP for the structures shown in 7 with *unseen* BCs and one external force.

Design Domain	Resolution	Number of test cases	MSE	Binary Accuracy	Compliance Error	Compliance Error Std.
from Fig. 7a (source network)	40 x 80	1000	8.24%	88.52%	8.31%	11.4 %
from Fig. 7b	80 x 160	500	8.81%	88.91%	6.65%	9.43%
from Fig. 7c	120 x 160	500	8.67%	89.64%	8.09%	12.43%
from Fig. 7d	120 x 240	500	8.92%	89.32%	7.03%	11.92%
from Fig. 7e	160 x 320	500	8.97%	89.65%	6.91%	11.80%
from Fig. 7f	200 x 400	250	10.07%	88.53%	9.47%	13.51%
<i>Average</i>			8.94%	89.10%	7.74%	11.74%

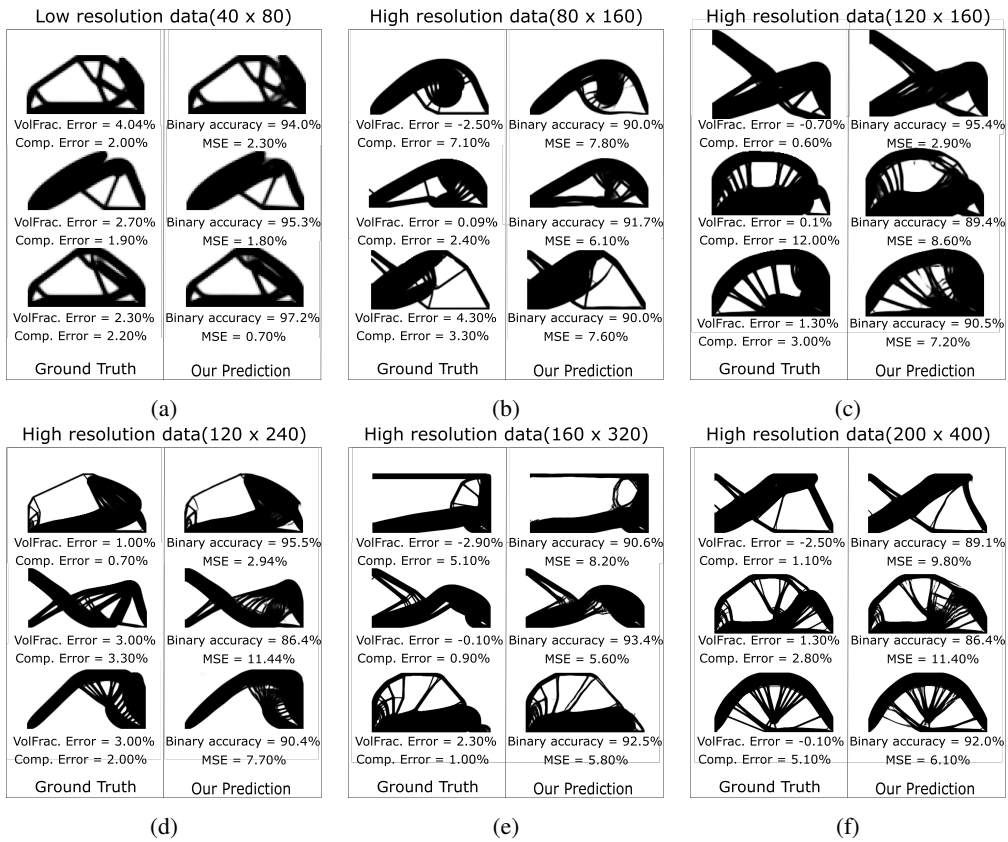


Fig. 8: Predictions vs. ground Truth for *seen* boundary conditions and *two external forces*: (a) the prediction of our source GAN alone; (b-f) optimal structures output by the fine tuned target model. The corresponding quality metrics are presented in Table 4.

main. Specifically, our examples illustrate that the topological patterns learned by GANTL from training data using a single external force are changing as the number of external forces increases.

Motivated by these observations, we also proposed a novel and promising alternative for stimulating the connectedness of the predictions produced by the network. Specifically, we constructed a custom loss function that correlates

the persistence diagram of the ground truth with that of the prediction. We showed that this novel topology-aware loss function adequately promotes those predictions that are topologically similar with the topology of the ground truth, and results in predictions with a significant improvement in terms of the connectedness of the predicted structural members. We posit here that enforcing topological properties of the optimal structures predicted by GANTL can be achieved by de-

Table 4: MSE, binary accuracy, compliance error and standard deviation relative to SIMP for the structures shown in 8 with *seen* BCs and *two external forces*.

Design Domain	Resolution	Number of test cases	MSE	Binary Accuracy	Compliance Error	Compliance Error Std.
from Fig. 8a (source network)	40 x 80	1000	5.84%	90.98%	7.74%	13.2 %
from Fig. 8b	80 x 160	500	7.50%	90.04%	16.12%	20.85%
from Fig. 8c	120 x 160	500	8.91%	90.01%	9.74%	16.88%
from Fig. 8d	120 x 240	500	7.67%	90.49%	9.50%	15.30%
from Fig. 8e	160 x 320	500	8.32%	90.13%	8.86%	14.62%
from Fig. 8f	200 x 400	250	8.42%	90.09%	10.92%	16.52%
<i>Average</i>			7.77%	90.29%	10.48%	16.22%

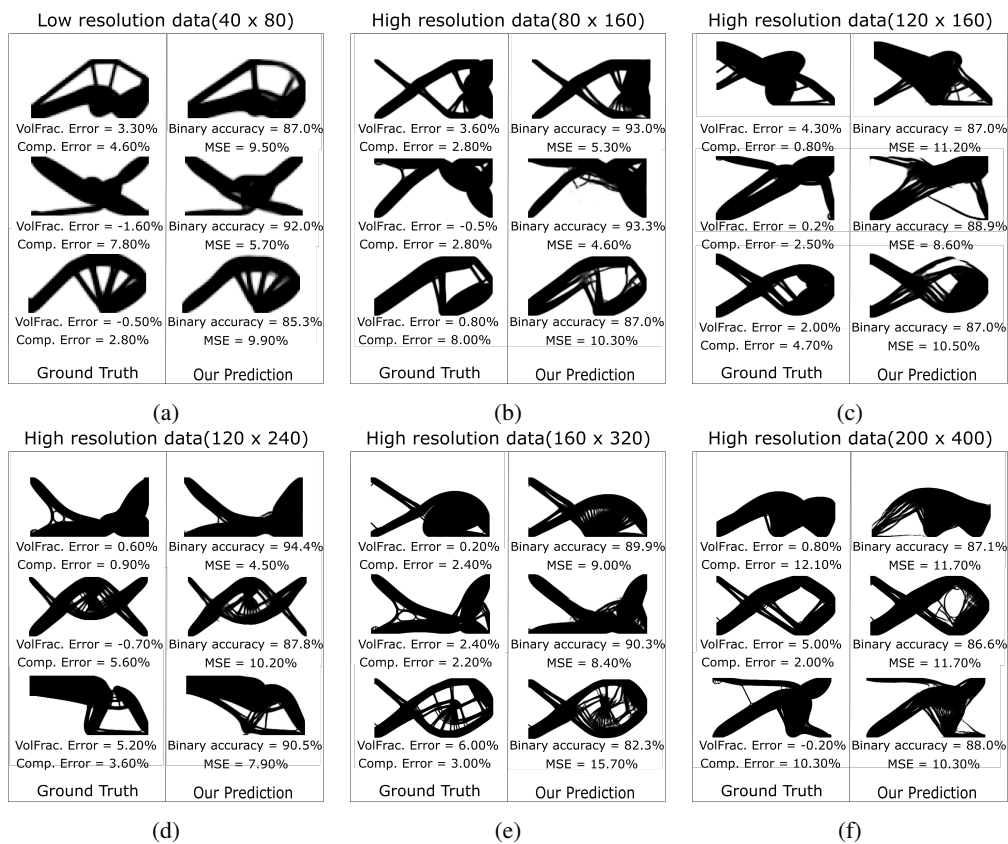


Fig. 9: Predictions vs. ground Truth for *unseen* boundary conditions and *two external forces*: (a) the prediction of our source GAN alone; (b-f) optimal structures output by the fine tuned target model. The corresponding quality metrics are presented in Table 5.

veloping specialized loss functions that integrate topological measures. However, a detailed investigation of these aspects is outside the scope of this paper.

Training GANs can be notoriously difficult, as discussed above. Our network faced the same challenges, which were addressed by employing a binary cross-entropy (BCE) with the logistic unit (logits) instead of using BCE alone. This modification helps the discriminator avoid loss values that

are near zero and consequently improve the stability of the GAN.

One key limitation in promoting predictions with topological properties that are similar to those of the ground truth structures is the added computational overhead of the topology-aware loss function. To compute this loss function, we had to compute the persistence diagram of all the ground truth structures and of the corresponding predicted

Table 5: MSE, binary accuracy, compliance error and standard deviation relative to SIMP for the structures shown in 9 with *unseen* BCs and *two external forces*.

Design Domain	Resolution	Number of test cases	MSE	Binary Accuracy	Compliance Error	Compliance Error Std.
from Fig. 9a (source network)	40 x 80	1000	9.40%	87.95%	10.25%	14.1 %
from Fig. 9b	80 x 160	500	10.60%	86.89%	12.73%	16.98%
from Fig. 9c	120 x 160	500	10.81%	87.20%	11.83%	16.61%
from Fig. 9d	120 x 240	500	10.85%	87.18%	11.21%	14.87%
from Fig. 9e	160 x 320	500	11.68%	86.66%	11.96%	16.11%
from Fig. 9f	200 x 400	250	11.86%	86.47%	11.23%	15.81%
<i>Average</i>			10.86%	87.05%	11.53%	15.74%

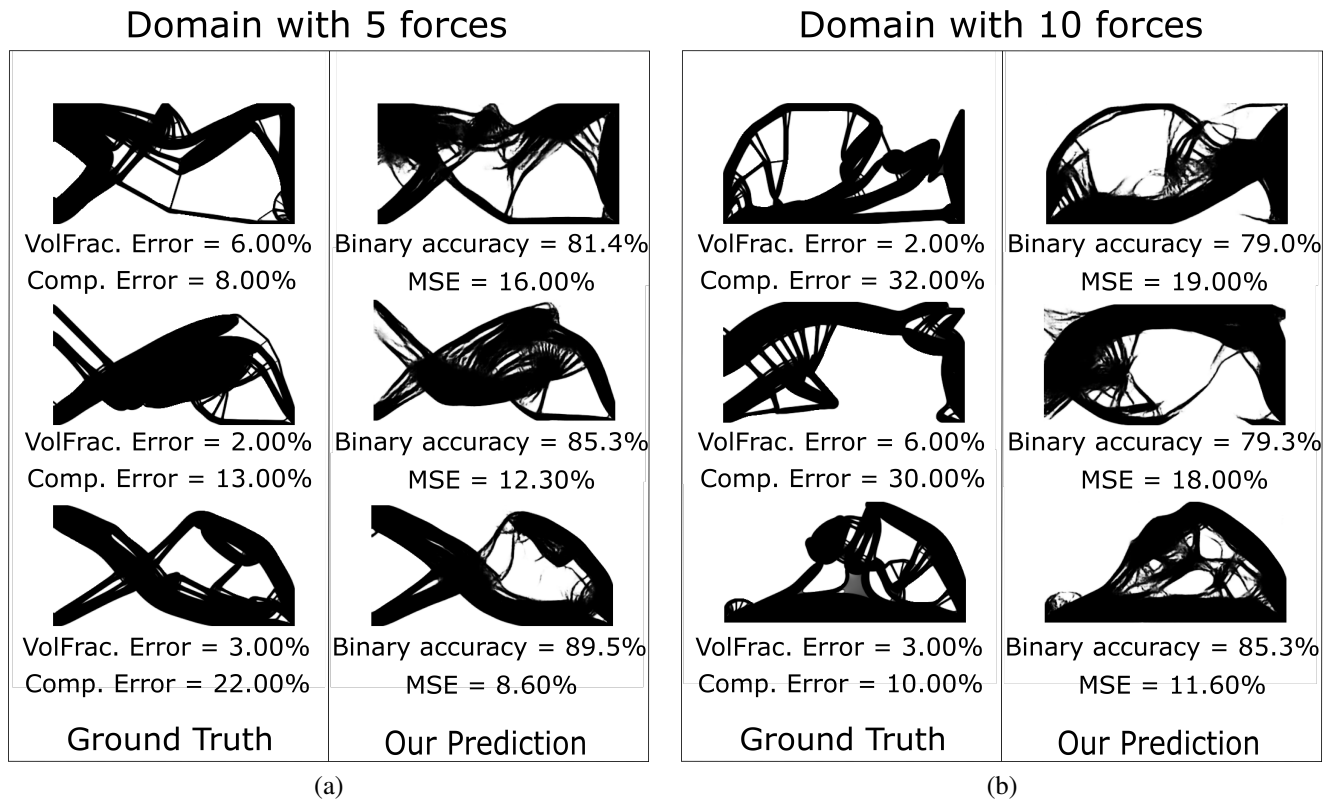


Fig. 10: Predictions vs. ground Truth for *seen* boundary conditions and 5 (a) and 10 (b) external forces for a  $200 \times 400$  domain resolution.

structures, as well as their pairwise bottleneck distance. This, in turn, significantly increased the training time of our GAN model, although it was computed only once for the training set.

As we demonstrated in [5], transfer learning not only significantly reduces the size of the dataset required to train the networks but improves the generalization ability of the deep learning models as long as the target network is fine-tuned with domains and boundary conditions that are not

part the training dataset of the source model. Importantly, by combining the generative power of GANs with the knowledge reuse of transfer learning, we train both the source and target networks only once and therefore eliminate the need to fine-tune the target network every time we want to explore new boundary conditions. This is a critical feature of the proposed GANTL that significantly reduces the training time and directly improves the practicality of the approach for design space explorations.

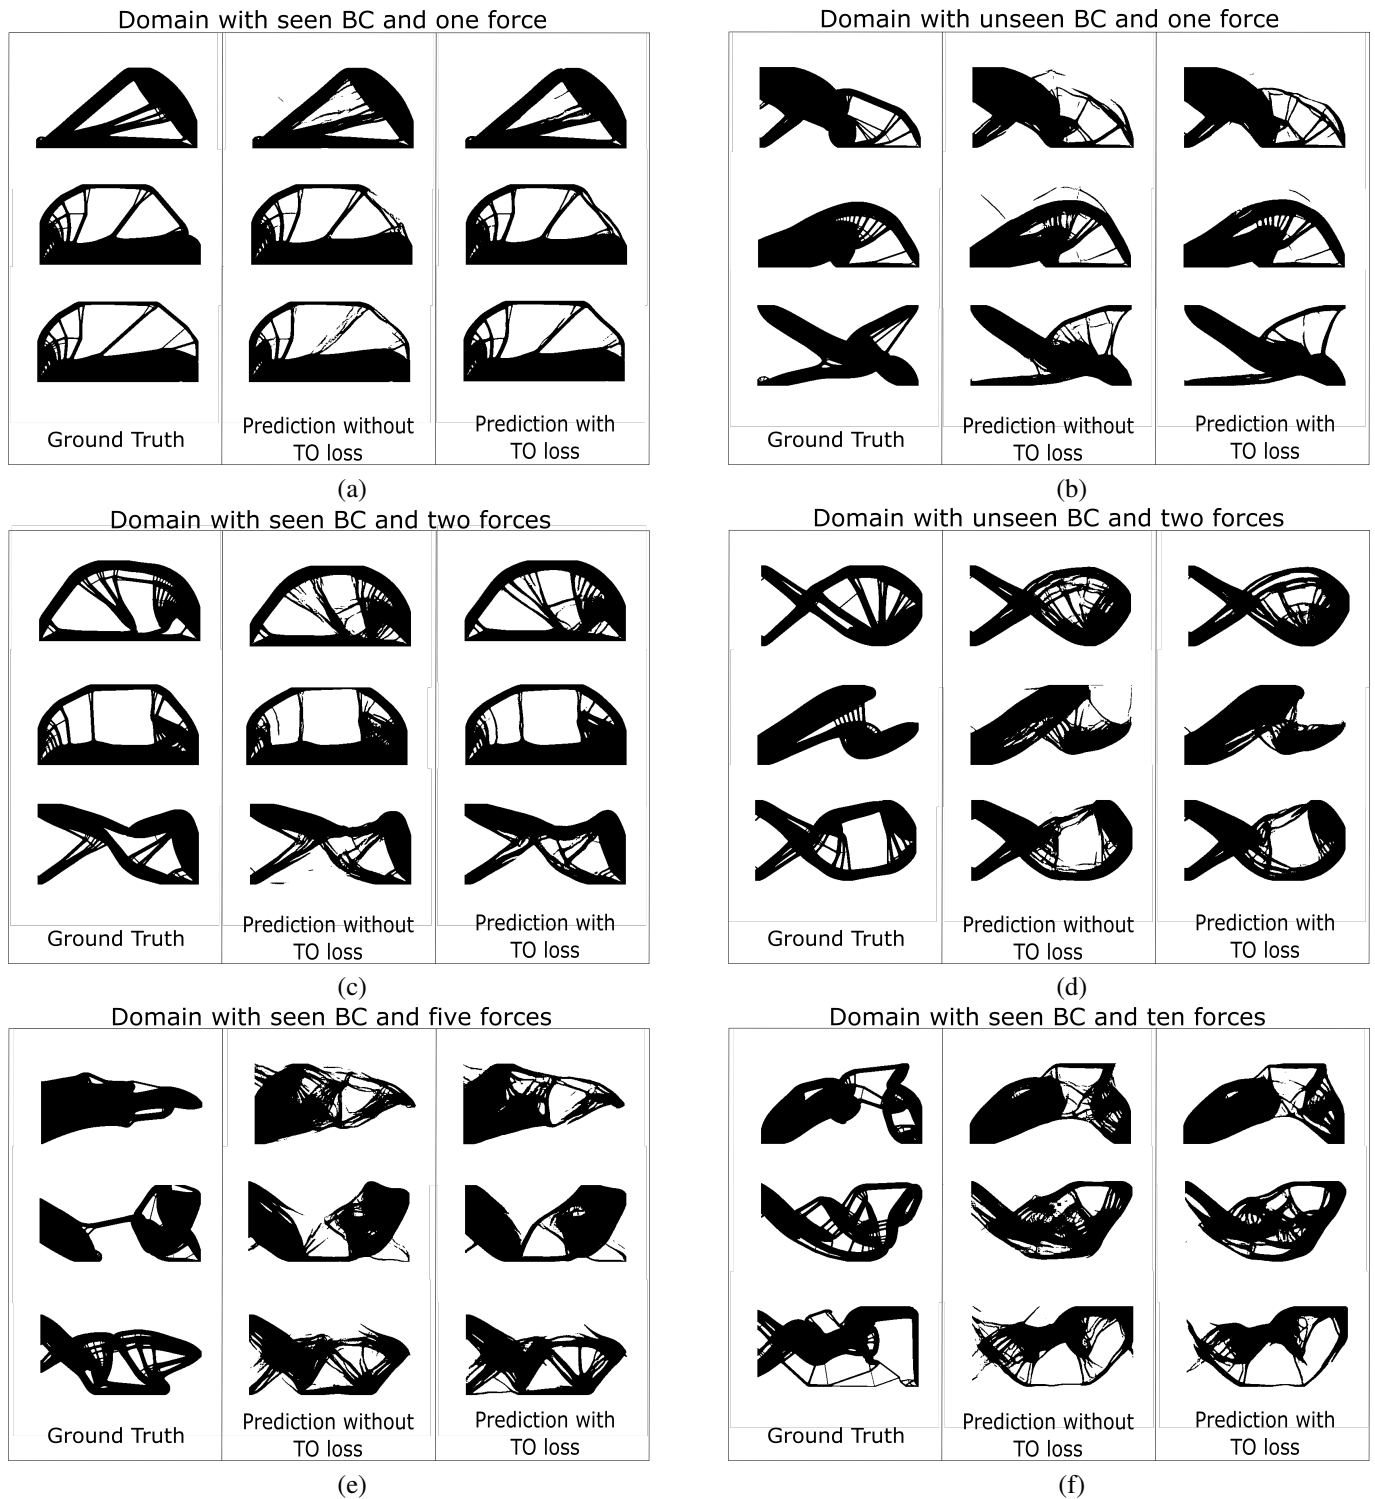


Fig. 11: Comparison between ground truth, and the predictions using standard cross entropy loss (second column in each panel) and those using the proposed topology-aware loss function based on the bottleneck distance between the corresponding persistence diagrams (third column in each panel). (a) seen BCs, 1 external force; (b) unseen BCs and one external force; (c) seen BCs and 2 external forces; (d) unseen BCs and 2 external forces; (e) seen BCs and 5 external forces; (f) seen BCs and 10 external forces.

To the best of our knowledge, this is the first attempt to combine transfer learning with generative adversarial net-

works in the area of topology optimization. Furthermore, our proposed topology-aware loss function pioneers further inte-

grations of powerful topological measures into the TO predictions made by machine learning algorithms, which will eventually lead to practical, powerful, and interactive design exploration tools with topology optimization.

## References

- [1] Bendsoe, M. P., and Sigmund, O., 2013. *Topology optimization: theory, methods, and applications*. Springer Science & Business Media.
- [2] Sigmund, O., and Maute, K., 2013. “Topology optimization approaches”. *Structural and Multidisciplinary Optimization*, **48**(6), pp. 1031–1055.
- [3] Amir, O., and Sigmund, O., 2011. “On reducing computational effort in topology optimization: how far can we go?”. *Structural and Multidisciplinary Optimization*, **44**(1), pp. 25–29.
- [4] Andreassen, E., Clausen, A., Schevenels, M., Lazarov, B. S., and Sigmund, O., 2011. “Efficient topology optimization in matlab using 88 lines of code”. *Structural and Multidisciplinary Optimization*, **43**(1), pp. 1–16.
- [5] Behzadi, M. M., and Ilieş, H. T., 2021. “Real-time topology optimization in 3d via deep transfer learning”. *Computer-Aided Design*, **135**, p. 103014.
- [6] Träff, E. A., Sigmund, O., and Aage, N., 2021. “Topology optimization of ultra high resolution shell structures”. *Thin-Walled Structures*, **160**, p. 107349.
- [7] Liu, H., Tian, Y., Zong, H., Ma, Q., Wang, M. Y., and Zhang, L., 2019. “Fully parallel level set method for large-scale structural topology optimization”. *Computers & Structures*, **221**, pp. 13–27.
- [8] Aage, N., Andreassen, E., and Lazarov, B. S., 2015. “Topology optimization using petsc: An easy-to-use, fully parallel, open source topology optimization framework”. *Structural and Multidisciplinary Optimization*, **51**(3), pp. 565–572.
- [9] Abueidda, D. W., Koric, S., and Sobh, N. A., 2020. “Topology optimization of 2d structures with nonlinearities using deep learning”. *Computers & Structures*, **237**, p. 106283.
- [10] Sosnovik, I., and Oseledets, I., 2019. “Neural networks for topology optimization”. *Russian Journal of Numerical Analysis and Mathematical Modelling*, **34**(4), pp. 215–223.
- [11] Banga, S., Gehani, H., Bhilare, S., Patel, S., and Kara, L., 2018. “3d topology optimization using convolutional neural networks”. *arXiv preprint arXiv:1808.07440*.
- [12] Karras, T., Aila, T., Laine, S., and Lehtinen, J., 2017. “Progressive growing of gans for improved quality, stability, and variation”. *arXiv preprint arXiv:1710.10196*.
- [13] Arjovsky, M., Chintala, S., and Bottou, L., 2017. Wasserstein gan.
- [14] Rawat, S., and Shen, M.-H. H., 2019. “A novel topology optimization approach using conditional deep learning”. *arXiv preprint arXiv:1901.04859*.
- [15] Li, B., Huang, C., Li, X., Zheng, S., and Hong, J., 2019. “Non-iterative structural topology optimization using deep learning”. *Computer-Aided Design*, **115**, pp. 172–180.
- [16] Yu, Y., Hur, T., Jung, J., and Jang, I. G., 2019. “Deep learning for determining a near-optimal topological design without any iteration”. *Structural and Multidisciplinary Optimization*, **59**(3), pp. 787–799.
- [17] Zhang, Y., Peng, B., Zhou, X., Xiang, C., and Wang, D., 2019. “A deep convolutional neural network for topology optimization with strong generalization ability”. *arXiv preprint arXiv:1901.07761*.
- [18] Nie, Z., Lin, T., Jiang, H., and Kara, L. B., 2021. “Topologygan: Topology optimization using generative adversarial networks based on physical fields over the initial domain”. *Journal of Mechanical Design*, **143**(3).
- [19] Chandrasekhar, A., and Suresh, K., 2020. “Tounn: Topology optimization using neural networks”. *Structural and Multidisciplinary Optimization*, pp. 1–15.
- [20] Chandrasekhar, A., and Suresh, K., 2021. “Multi-material topology optimization using neural networks”. *Computer-Aided Design*, **136**, p. 103017.
- [21] Edelsbrunner, H., and Harer, J., 2008. “Persistent homology—a survey”. *Contemporary mathematics*, **453**, pp. 257–282.
- [22] Cohen-Steiner, D., Edelsbrunner, H., and Harer, J., 2007. “Stability of persistence diagrams”. *Discrete & computational geometry*, **37**(1), pp. 103–120.
- [23] Mosinska, A., Marquez-Neila, P., Koziński, M., and Fua, P., 2018. “Beyond the pixel-wise loss for topology-aware delineation”. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3136–3145.
- [24] Hu, X., Li, F., Samaras, D., and Chen, C., 2019. “Topology-preserving deep image segmentation”. *Advances in neural information processing systems*, **32**.
- [25] Munkres, J. R., 2018. *Elements of algebraic topology*. CRC press.
- [26] Sigmund, O., 2001. “A 99 line topology optimization code written in matlab”. *Structural and multidisciplinary optimization*, **21**(2), pp. 120–127.
- [27] Rozvany, G., 2001. “Stress ratio and compliance based methods in topology optimization—a critical review”. *Structural and Multidisciplinary Optimization*, **21**(2), pp. 109–119.
- [28] Bendsoe, M. P., and Sigmund, O., 2007. “Topology optimization”. In *Optimization Of Structural And Mechanical Systems*. World Scientific, pp. 161–194.
- [29] Yang, R., and Chen, C., 1996. “Stress-based topology optimization”. *Structural optimization*, **12**(2-3), pp. 98–105.
- [30] Le, C., Norato, J., Bruns, T., Ha, C., and Tortorelli, D., 2010. “Stress-based topology optimization for continua”. *Structural and Multidisciplinary Optimization*, **41**(4), pp. 605–620.
- [31] Bendsoe, M. P., 1989. “Optimal shape design as a material distribution problem”. *Structural optimization*, **1**(4), pp. 193–202.

- [32] Sigmund, O., 2007. “Morphology-based black and white filters for topology optimization”. *Structural and Multidisciplinary Optimization*, **33**(4-5), pp. 401–424.
- [33] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y., 2014. “Generative adversarial networks”. *arXiv preprint arXiv:1406.2661*.
- [34] Bau, D., Strobel, H., Peebles, W., Zhou, B., Zhu, J.-Y., Torralba, A., et al., 2020. “Semantic photo manipulation with a generative image prior”. *arXiv preprint arXiv:2005.07727*.
- [35] Antoniou, A., Storkey, A., and Edwards, H., 2017. “Data augmentation generative adversarial networks”. *arXiv preprint arXiv:1711.04340*.
- [36] Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A., 2017. “Unpaired image-to-image translation using cycle-consistent adversarial networks”. In Proceedings of the IEEE international conference on computer vision, pp. 2223–2232.
- [37] Zhan, Y., Hu, D., Wang, Y., and Yu, X., 2017. “Semisupervised hyperspectral image classification based on generative adversarial networks”. *IEEE Geoscience and Remote Sensing Letters*, **15**(2), pp. 212–216.
- [38] Gui, J., Sun, Z., Wen, Y., Tao, D., and Ye, J., 2020. “A review on generative adversarial networks: Algorithms, theory, and applications”. *arXiv preprint arXiv:2001.06937*.
- [39] Chollet, F., 2017. “Xception: Deep learning with depthwise separable convolutions”. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1251–1258.
- [40] Pan, S. J., and Yang, Q., 2010. “A survey on transfer learning iee transactions on knowledge and data engineering”. *22 (10): 1345*, **1359**.
- [41] Shin, S., Hwang, K., and Sung, W., 2016. “Generative knowledge transfer for neural language models”. *arXiv preprint arXiv:1608.04077*.
- [42] Wang, Y., Wu, C., Herranz, L., van de Weijer, J., Gonzalez-Garcia, A., and Raducanu, B., 2018. “Transferring gans: generating images from limited data”. In Proceedings of the European Conference on Computer Vision (ECCV), pp. 218–234.
- [43] Noguchi, A., and Harada, T., 2019. “Image generation from small datasets via batch statistics adaptation”. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2750–2758.
- [44] Park, H., Yoo, Y., and Kwak, N., 2018. “Mc-gan: Multi-conditional generative adversarial network for image synthesis”. *arXiv preprint arXiv:1805.01123*.
- [45] Hu, J., Shen, L., and Sun, G., 2018. “Squeeze-and-excitation networks”. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7132–7141.
- [46] He, K., Zhang, X., Ren, S., and Sun, J., 2016. “Deep residual learning for image recognition”. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.
- [47] Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q., 2017. “Densely connected convolutional networks”. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4700–4708.
- [48] Gómez, J., and Guarín-Zapata, N., 2018. Solidspy: 2d-finite element analysis with python.
- [49] Xu, K., Zhang, M., Li, J., Du, S. S., ichi Kawarabayashi, K., and Jegelka, S., 2021. How neural networks extrapolate: From feedforward to graph neural networks.
- [50] Efrat, A., Itai, A., and Katz, M. J., 2001. “Geometry helps in bottleneck matching and related problems”. *Algorithmica*, **31**(1), pp. 1–28.
- [51] Tralie, C., Saul, N., and Bar-On, R., 2018. “Ripser.py: A lean persistent homology library for python”. *The Journal of Open Source Software*, **3**(29), Sep, p. 925.
- [52] Maria, C., Boissonnat, J.-D., Glisse, M., and Yvinec, M., 2014. “The gudhi library: Simplicial complexes and persistent homology”. In International Congress on Mathematical Software, Springer, pp. 167–174.