PHOTONICS Research

Engineering of multiple bound states in the continuum by latent representation of freeform structures

Ronghui Lin, Zahrah Alnakhli, and Xiaohang Li*

Advanced Semiconductor Laboratory, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia
**Corresponding author: xiaohang.li@kaust.edu.sa*

Received 26 November 2020; revised 24 January 2021; accepted 26 January 2021; posted 29 January 2021 (Doc. ID 415655); published 22 March 2021

We demonstrate a neural network capable of designing on-demand multiple symmetry-protected bound states in the continuum (BICs) in freeform structures with predefined symmetry. The latent representation of the freeform structures allows the tuning of the geometry in a differentiable, continuous way. We show the rich band inversion and accidental degeneracy in these freeform structures by interacting with the latent representation directly. Moreover, a high design accuracy is demonstrated for arbitrary control of multiple BIC frequencies by using a photonic property readout network to interpret the latent representation.

https://doi.org/10.1364/PRJ.415655

1. INTRODUCTION

The optical bound states in the continuum (BICs) refer to an exotic class of states that remain perfectly confined despite lying in a continuous spectrum of radiating waves [1,2]. The lack of outgoing radiation means these states have a theoretical infinite lifetime but are unable to be excited by far-field radiation. In practice, perturbations are introduced intentionally or unintentionally to turn the BICs into leaky resonances with finite albeit high Q factors [3]. Engineering these modes to leverage the high quality (Q) factor is benefiting a wide range of applications such as lasers [4,5], sensors [6,7], and nonlinear optics [8,9].

Recent reports show the BICs arise from the vortex centers of the polarization field and they carry quantized topological charges [10]. Merging of multiple BIC points can be used to achieve robust ultra-high Q factor modes immune to out-of-plane scattering losses [11]. Moreover, being able to manipulate multiple BICs might further benefit areas where high Q factors are desirable at multiple frequencies such as nonlinear optics [12] and multi-wavelength sensing [13–15].

The core idea of realizing BICs is the parameter tuning to cancel out the far-field radiation. This is mostly achieved by the sweeping of parametric geometries such as circles, rectangles, ellipses, or the combination of them. These geometrics can be easily described and modified by equations with a limited number of variables and, thus, a limited degree of tunability. They run into problems with more advanced tuning tasks such as the manipulation of multiple BICs simultaneously. So far, arbitrary control of multiple BICs has never been achieved.

The freeform structures optimized by evolutionary algorithms [16,17] and adjoint methods [18,19] show great potential in the topology optimization of photonic structures and promise new methods of tuning BICs. They are not bound by any equations and, thus, offer a limitless degree of tunability, which may yield designs that outperform those by conventional geometries [19,20]. However, these algorithms are generally costly in computational resources. Besides these interactive optimization methods, deep neural networks (DNNs) are viable tools in handling complicated photonic structures [20–22]. There have been some impressive attempts to design freeform photonic structures using generative adversarial networks (GANs) [23–27]. A problem with the GANs is the difficulty in training with the possibility of noisy outputs where extra filtering and smoothing algorithms are needed to refine the geometries [24,25]. Another issue with the GANs is that the desired optical properties are directly linked to the geometrical shapes described by pixels, which have much higher dimensions than that of the optical properties such as the transmission and reflection spectra. This results in converging problems and bad generalization performance. Furthermore, the structures with predefined symmetries, which are a critical quality in the field of photonics, have not been demonstrated by GANs. To make the DNNs learn and generate the symmetry and parity properties of the real-life structures is an active research topic in the machine learning community [28–30]. Special techniques such as symmetry loss [31] and structured GANs [32] are necessary to ensure symmetrical outputs. In this work, we demonstrate a DNN structure based on the variational autoencoders (VAEs) that can handle freeform photonic structures with predefined symmetry. Instead of connecting the property to the pixel representation of the geometries directly, we convert the

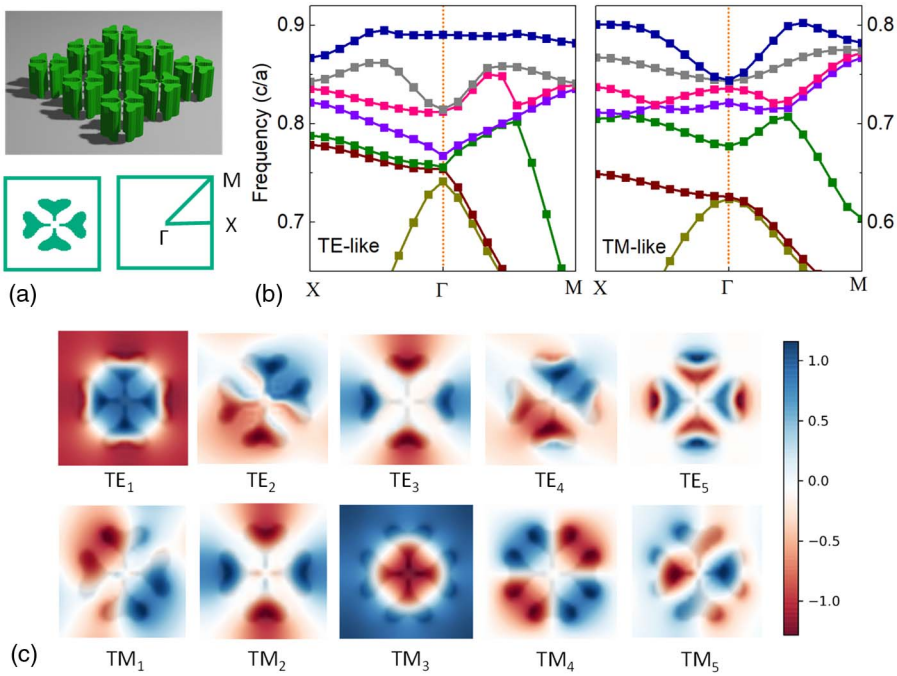

Fig. 1. (a) Top, an artistic rendering of the C_{4v} photonic crystal considered. Bottom, planar view of the C_{4v} unit cell and the definition of high symmetry points. (b) Band diagram for TE-like and TM-like modes. (c) H_z and E_z Bloch mode profiles for TE-like and TM-like modes, respectively.

geometries into the latent representation, which can be linked to the optical properties more easily and, hence, increase the stability of the inverse design. The latent representation also allows small perturbations of the geometries, which in turn allow continuous manipulation of the photonic properties. With a property readout network to interpret the latent representation, we demonstrate arbitrary, on-demand control of multiple BICs with high accuracy. The band inversion and accidental degeneracy arising from these symmetrical freeform structures can also be a platform for further discoveries and innovations.

2. SYMMETRY-PROTECTED BIC IN C_{4v} LATTICE

Consider a photonic crystal with C_{4v} symmetry as shown in Fig. 1(a). At Γ point and below the diffraction limit, the only radiation channel is the zeroth-order diffraction, which has an odd parity under C_2 operation. Therefore, any modes with an even symmetry are decoupled and turned into BICs [1]. As shown in Table 1, the C_{4v} group has four nondegenerate irreducible representations (A_1, A_2, B_1, B_2) and one doubly degenerate irreducible representation (E). It is interesting to see that all the nondegenerate representations are even under C_2 transformation. Thus, it is straightforward to conclude that any nondegenerate modes below the diffraction limit turn into BICs at Γ point.

Table 1. Character Table for C_{4v} Group

	E	$2C_4$	C_2	$2\sigma_v$	$2\sigma_d$
A_1	1	1	1	1	1
A_2	1	1	1	−1	−1
B_1	1	−1	1	1	−1
B_2	1	−1	1	−1	1
E	2	0	−2	0	0

Since these BICs are protected by symmetry, their frequencies can be tuned by geometrical perturbations as long as the symmetry is maintained. The freeform structure with C_{4v} symmetry as shown in Fig. 1(a) is studied. We consider a unit cell of 1 μm × 1 μm × 2 μm with the material refractive index $n = 3.5$, which corresponds to Si in the infrared region. The geometry is assumed to have a fixed height of 550 nm in the air with a mirror symmetry about the z plane. Hence, the photonic modes can be classified as TE- and TM-like modes. The unit cell is discretized with a 100 × 100 × 200 mesh, and the plane wave expansion method is used to determine the eigenmodes and eigenfrequencies along the high symmetry points $X(0, 0.5) - \Gamma(0, 0) - M(0.5, 0.5)$. The calculations are implemented with an open-source package MIT photonic bands (MPB) [33]. The photonic band diagrams and the corresponding Bloch mode profiles are shown in Figs. 1(b) and 1(c). The H_z field is plotted for TE-like modes, and the E_z field is plotted for TM-like modes. The TE_2, TE_4, TM_1, and TM_5 are doubly degenerate modes, as can be verified from the band diagrams. These mode profiles have an odd parity for C_2 transformation. The nondegenerate modes are all even modes for C_2 transformation, and they all lie below the diffraction limit c/na, where c, n, and a are the speed of light, the refractive index of air, and the lattice size, respectively. Hence, they are all symmetry-protected BIC modes. The Bloch modes of the freeform structures largely resemble the field profiles of Mie-type resonators. However, they are distorted and deformed with a lot of local features, which are the key to engineer and fine-tune the BIC frequencies.

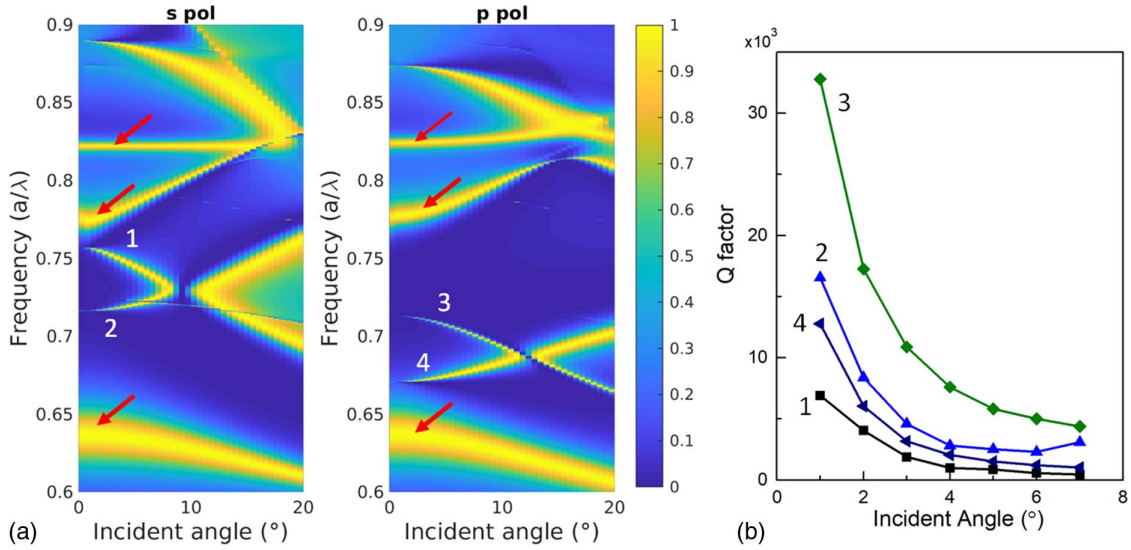

Fig. 2. (a) Transmission spectra at different incident angles for s polarization and p polarization. The red arrows indicate the doubly degenerate modes, and the numbered modes are nondegenerate. (b) The Q factors of the BIC modes shown in (a).

Figure 2(a) shows the transmission spectra at different incident angles for the geometry in Fig. 1 obtained by rigorous coupled-wave analysis (RCWA). The 0° incident angle corresponds to the Γ point. We can observe three modes at frequencies 0.63, 0.77, 0.83 in both s polarization and p polarization [marked by red arrows in Fig. 2(a)]. They correspond to the TM_1, TE_2, and TE_4 doubly degenerate modes because the s polarization and p polarization have a 90° rotation relationship, which matches the symmetries of the doubly degenerate modes shown in Fig. 1. Other modes, such as the numbered ones in Fig. 2(a), are nondegenerate, which appear either in s polarization or p polarization. These modes show vanishing linewidths as the incident angle approaches 0°. The Q factors for the four modes highlighted show diverging trends as the incident angle approaches 0°, which indicates that they are indeed BICs.

3. LATENT REPRESENTATION OF GEOMETRY USING β-VAE

A VAE structure as shown in Fig. 3 is designed to manage the freeform structures with C_{4v} symmetry. In the VAE, the input $\mathbf{x}$ passes through two DNNs sequentially, namely the encoder $q(\mathbf{z}|\mathbf{x})$ and the decoder $p(\mathbf{x}|\mathbf{z})$, to output a probability distribution $\mathbf{x}'$. The latent vector $\mathbf{z}$ has a much lower dimension than the input $\mathbf{x}$. As the data pass through such a bottleneck layer, they are compressed, and the network is forced to learn an efficient representation of the input. Data in a much higher dimension can be represented by a low dimension vector faithfully after training. In this study, the input is 64 × 64 pixels binary images with 1 representing the material and 0 representing the air. The latent vector $\mathbf{z}$ is an array with the form $[z_1, z_2, \ldots, z_{10}]$.

We use the β-VAE [34,35] with a loss function as follows:

$$\mathcal{L} = -\mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}|\mathbf{z})] + \beta D_{\mathrm{KL}}[q(\mathbf{z}|\mathbf{x})\|p(\mathbf{z})]. \quad (1)$$

The first term is the reconstruction loss, which forces the decoder to represent the input as closely as possible. The second term is the Kullback–Leibler (KL) divergence between the prior distribution $p(\mathbf{z})$ and the encoder distribution $q(\mathbf{z}|\mathbf{x})$. It is a regularization term that forces the latent representation to assume the same standard normal distribution as the prior $p(\mathbf{z})$. The resultant latent representation is centered and closely packed in the latent space, and most importantly, continuous and interpolable. The β factor imposes extra weight on the KL divergence, thus increasing the regularization power. Previous reports show an increase in the β factor can promote the disentanglement of latent representation [35], allowing them to be more interpretable. In this work, we choose a β factor of 3. The β-VAE can not only reproduce the training data but also create new data with the same distribution as the training data. As the high dimension data are collapsed to a lower dimension, a lot of the high dimension features and minor details are filtered out. Hence, the output geometries are naturally smooth with less ultrafine structures such as sharp corners and isolated islands. Moreover, the latent representation has similar dimensions with

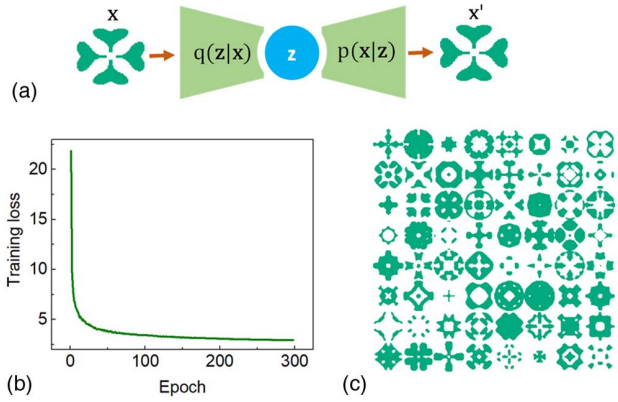

Fig. 3. (a) VAE structure used for geometry management. (b) The training loss for β-VAE. (c) Examples of the β-VAE generated geometries.

our optical response, which is easier to link to the optical properties by DNN.

We use 20,000 randomly generated geometries with C_{4v} symmetry to train the β-VAE. The training geometries are generated by applying symmetry operations to random polygons. The training loss of the β-VAE in Fig. 3(b) shows good convergence, which is an advantage over GANs. Randomly generated latent vectors are used to test the output of the β-VAE after the training. Some examples of the output are shown in Fig. 3(c). The VAE can produce a wide variety of geometries with different topology while maintaining the perfect C_{4v} symmetry. These geometries are obtained without any additional filtering, but they are smooth with little noise. Since the latent representation is continuous, a perturbation of the latent vector $\mathbf{\Delta z}$ results in a perturbation $\mathbf{\Delta x}'$ in the output. This is demonstrated in Visualization 1, where we show the change of the output geometries as the $\mathbf{z}$ vector is varied continuously.

To study how the continuous deformation of such freeform structures influences the photonic band structures, we generate a random initial geometry and then vary one component of the latent vector z_3 continuously with a step of 0.02. A total number of 20 geometries are generated, and geometries 8 and 12 are compared in detail. Geometry 8 is the same as the one depicted in Fig. 1, and geometry 12 is a slightly deformed version. Their shapes are shown in Fig. 4(a). The band frequencies at the Γ point shift smoothly when the geometrical shapes change continuously, as shown in Fig. 4(b). An interesting phenomenon is the band crossing during this process. Comparing with the TE Bloch mode profiles in Fig. 1, we find the mode profiles of TE_2 and TE_3 are swapped. Between these two shapes, there is an intermediate shape (geometry 10) where these two bands cross and form an "accidental degeneracy." TE_2 is a nondegenerate mode with even parity while TE_3 is a doubly degenerate mode with odd parity. Previous studies show that a Dirac cone with linear dispersion [36,37] appears when these bands meet. Such a Dirac cone dispersion plays a vital role in zero-index materials [38,39] and topological photonics [40], which could be another application for these freeform structures.

The reason for this band crossing can be understood from the Bloch mode profiles. In Fig. 4(c), the field of TE_2 is mostly localized at the center of the heart-shaped structures; hence, it is less sensitive as the heart shape deforms. In contrast, TE_3 is more susceptible to such changes because a considerable amount of field intensity is located at the site of deformation. Since the amount of the frequency shift is proportional to the field concentration in the perturbation area [41], the TE_2 band remains mostly unchanged while the TE_3 band shifts more drastically, causing the band crossing. In this geometry, TE_1, TE_2, and TE_5 are symmetry-protected BICs. The shifting and crossing of the bands give us a lot of possible combinations of BICs. Since these shapes are controlled by a latent vector with 10 dimensions, there are 10 dimensions to fine-tune the geometries and, hence, the BIC combinations.

We also study the shift of the photonic bands as the geometry in Fig. 1 is scaled by a factor of 0.5 to 1.2. The band frequencies at Γ point are shown in Fig. 5(a), and the Bloch mode profiles are shown in Fig. 5(b). Comparing the band order to that of Fig. 1, we notice drastic band crossing for TM polarization. Two sets of band crossing $[(TM_2, TM_3)$ and $(TM_4, TM_5)]$ are observed when the scaling factor is varied from 1 to 0.8. TM_4 has odd parity, and TM_5 has even parity while both TM_2 and TM_3 have even parity. The crossing of TM_4 and TM_5 results in a Dirac-like dispersion while the crossing of TM_2 and TM_3 would lead to a double quadratic dispersion as shown in previous reports [36]. These results show great potential for band engineering for photonic applications.

4. ON-DEMAND MULTIPLE BIC DESIGN

The continuous deformation and scaling of the geometry make it possible to engineer a large number of BICs. To achieve the on-demand design of multiple BICs, a property readout

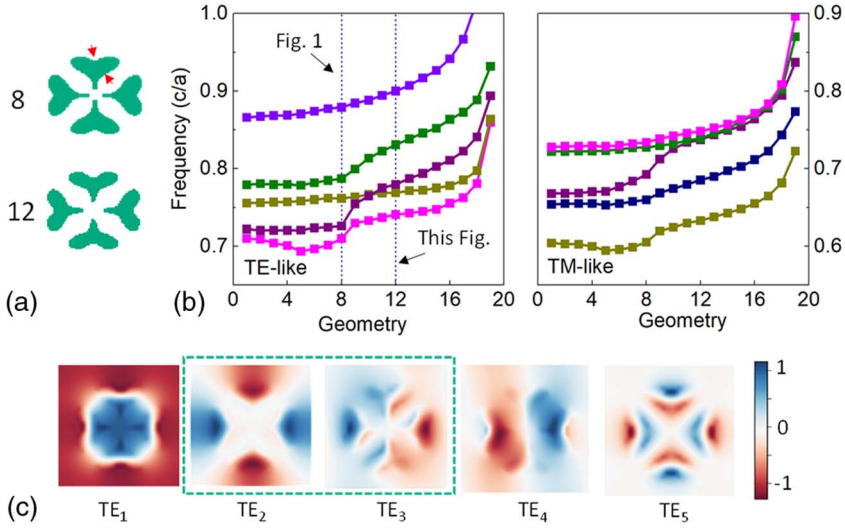

Fig. 4. (a) Shapes of geometry 8 and 12, with the sites of deformation marked by red arrows. (b) The shift of TE-like and TM-like bands at Γ point as the latent vector is varied continuously (see Visualization 1 for the continuous variation of the geometries). (c) H_z field of TE-like modes for geometry 12. The inversed bands are grouped by dashed green boxes.

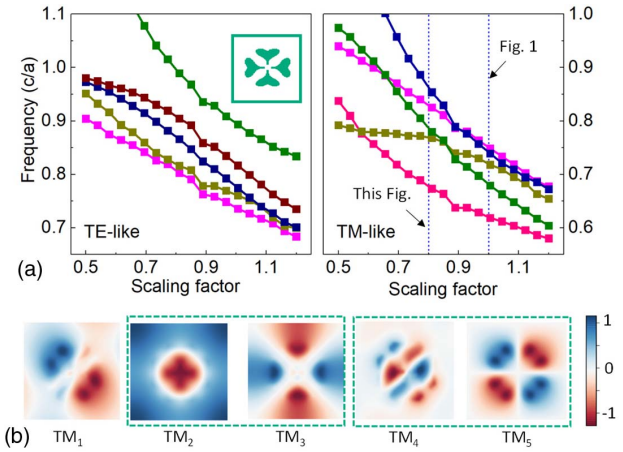

Fig. 5. (a) Shift of TE-like and TM-like bands at Γ point as the scaling factor is varied. The inset in (a) shows the geometry considered. (b) The E_z field of TM-like modes for geometry with a scaling factor of 0.8. The inversed bands are grouped by dashed green boxes.

network is necessary to link the latent representation of geometries to their photonic properties. The whole DNN structure is shown in Fig. 6(a). The property readout network, consisting of two convolutional neural networks (CNNs), is connected to the geometry handling network (β-VAE) by the latent vector **z**. The information flow during training is shown in Fig. 6(b). In process 1, the geometry management network is trained to realize the mapping between the input geometry **x**, latent space **z**, and output geometry **x**′. With this mapping established, a total number of 20,000 **z** vectors are sampled from the latent space, and their corresponding geometries are recorded. After that, their photonic bands at the Γ point are calculated. Then the BIC frequencies are selected by removing the degenerate modes. Three lowest BIC frequencies are picked to form the BIC vector $\mathbf{b} = [b_1, b_2, b_3]$. The **z**–**b** pair is used to train the property readout network. In process 2, the latent vector **z** is set as the input and **b** as the output to train CNN_2. This ensures the forward mapping from the latent vector **z** to its corresponding **b** vector. After that, the weights of CNN_2 are fixed when the BIC vector **b** is fed to the CNN_1-**z**-CNN_2 network, and the weights of CNN_1 are updated. This CNN_1-**z**-CNN_2 structure is similar to the tandem network [42] that is used to solve

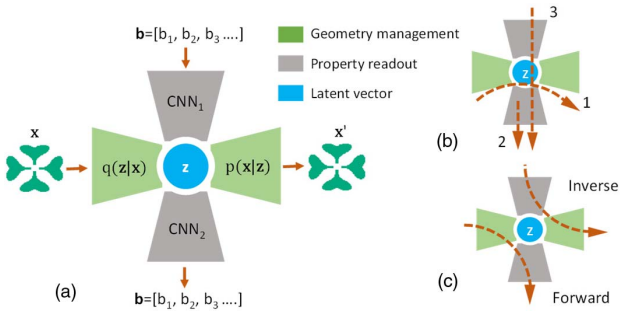

Fig. 6. (a) Whole DNN structure studied in this work. (b) Information flow during training, where 1 is the training of β-VAE, 2 is the training of CNN_2, and 3 is the training of CNN_1. (c) Information flow for the forward modeling and the inverse design.

data inconsistency problems. Since the mapping from physical property **b** to the latent vector **z** is not unique, there might be converging problems if CNN_1 is trained directly. The idea is that the mapping from **z** to **b** is always one to one so that CNN_2 can always converge to one possible solution. After that, this solution is used as the criteria to train CNN_1, which ensures the convergence. After all the networks are trained, the whole DNN can be used for the forward modeling or the inverse design, as shown in Fig. 6(c). In the former case, a geometry is fed to the network, and the output is the BIC vector **b**. In the latter case, a target **b** vector is fed to the network, and the output is the geometry. The detailed parameters of the whole DNN can be found in the last section of this work. We use 18,000 sets of **z**–**b** pairs for training, 1000 for testing, and the final 1000 for validation. The mean squared error (MSE) is used for the training of both CNN_1 and CNN_2. The training and testing losses are shown in Fig. 7. It can be seen that the training and testing losses converge after 200 epochs for both CNN_1 and CNN_2. We use both the validation data set and randomly generated data to test the performance of the network. In Fig. 8(a), the 1000 sets of reserved validation data are fed to the CNN_1-**z**-CNN_2 network, and the output is compared to the target value. In Fig. 8(b), we generate 1000 random **b** vectors to test the output of the network. In both figures, the x axis is the target value, and the y axis is the output of the DNN. The green bands are defined by $y = x \pm 15$ nm. We can see a strong linear correlation for the target value and the network output with 99% of the points lying within the bands of $y = x \pm 15$ nm, indicating good design accuracy, even for randomly generated input.

Next, we show a few specific examples of the inverse design of multiple BICs. Three sets of random **b** vectors are generated and fed to the DNN through the inverse design pathway. The obtained geometries are simulated to determine their BIC frequencies. The target value, the value inferred by DNN, and the value obtained by simulations agree well, as shown in Fig. 9(a). The insets show the output geometries. Figure 9(b) shows the band diagrams for the output geometries with the target BIC frequencies marked by red arrows. Figure 9(c) shows the corresponding mode profiles at the BIC frequencies. The BIC frequencies are all nondegenerate modes with an even parity for the C_2 transformation, which means they are indeed BIC states. We observe the maximum discrepancy for b_1 and b_3 is around 30 nm, which is within the ±15 nm range shown earlier. This discrepancy is around 2.5% of the shortest operation wavelength (from 1200 nm to 1700 nm), which is

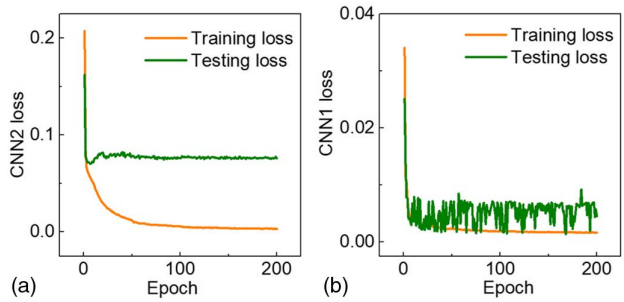

Fig. 7. Training and testing losses for (a) CNN_2 and (b) CNN_1.

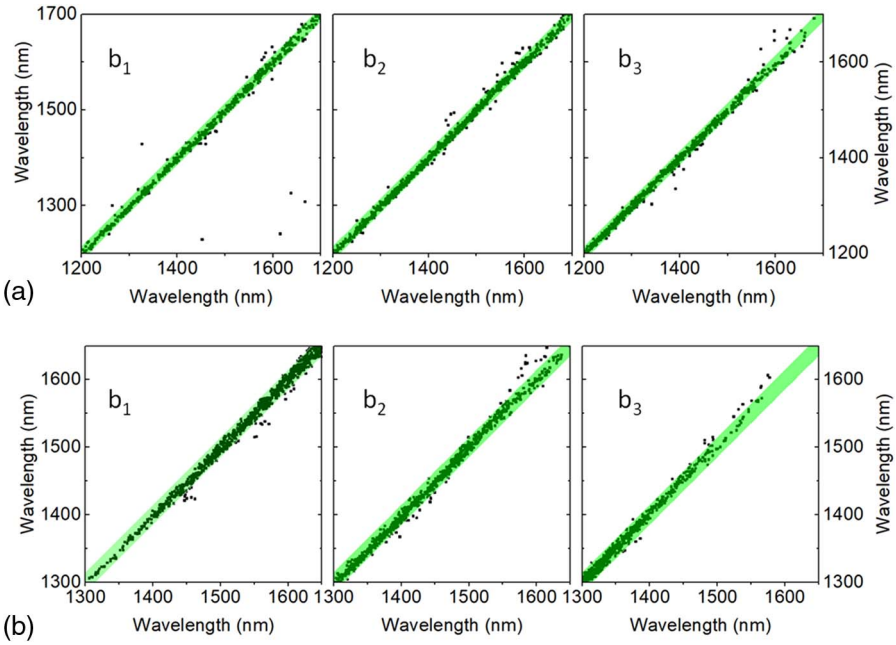

Fig. 8. (a) Correlation between target BIC wavelengths in the validation data set and the output of the DNN. (b) Correlation between randomly generated target BIC wavelengths and the output of the DNN. The x axis is the target value, and the y axis is the DNN output in both (a) and (b).

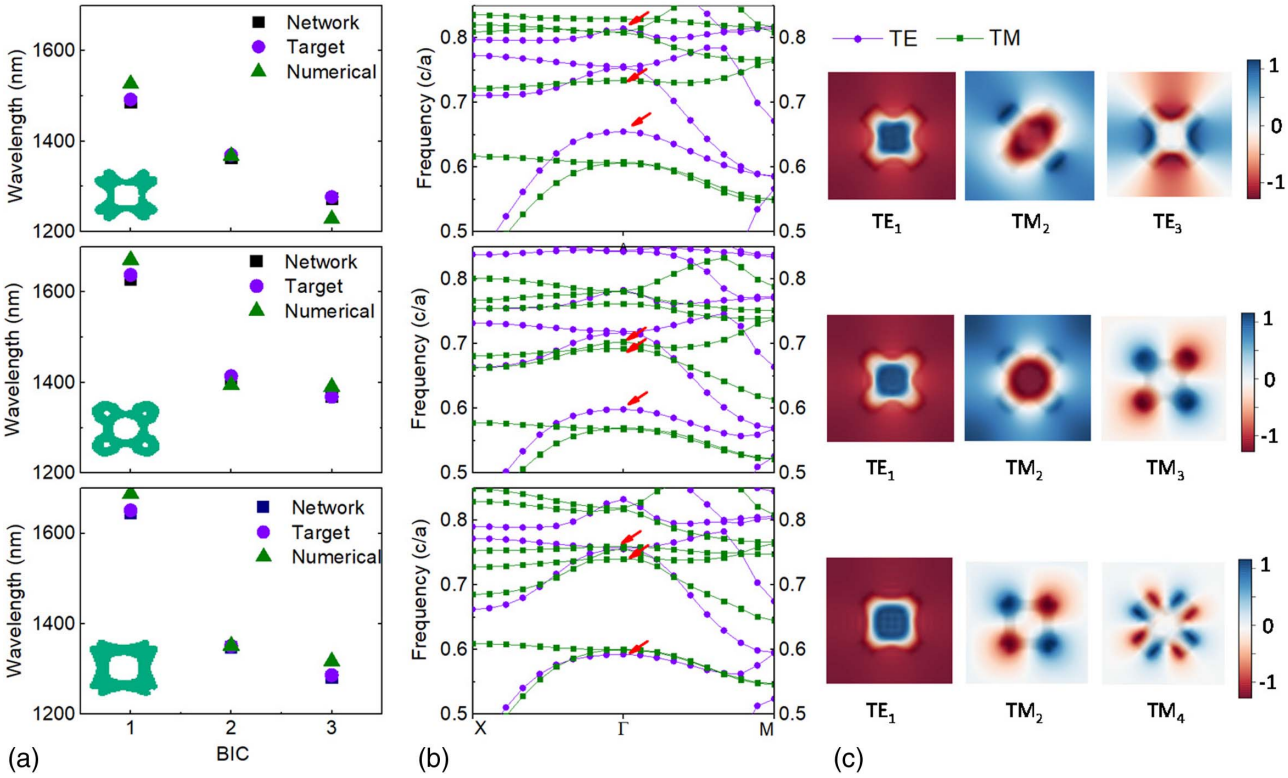

Fig. 9. Demonstration of the multiple BIC inverse design. (a) The comparison of the random target BIC wavelengths, DNN output, and numerical simulation. The inset shows the output geometry of the DNN. (b) Band diagrams for the designed structures. (c) The H_z and E_z field profiles for the TE and TM BIC states, respectively.

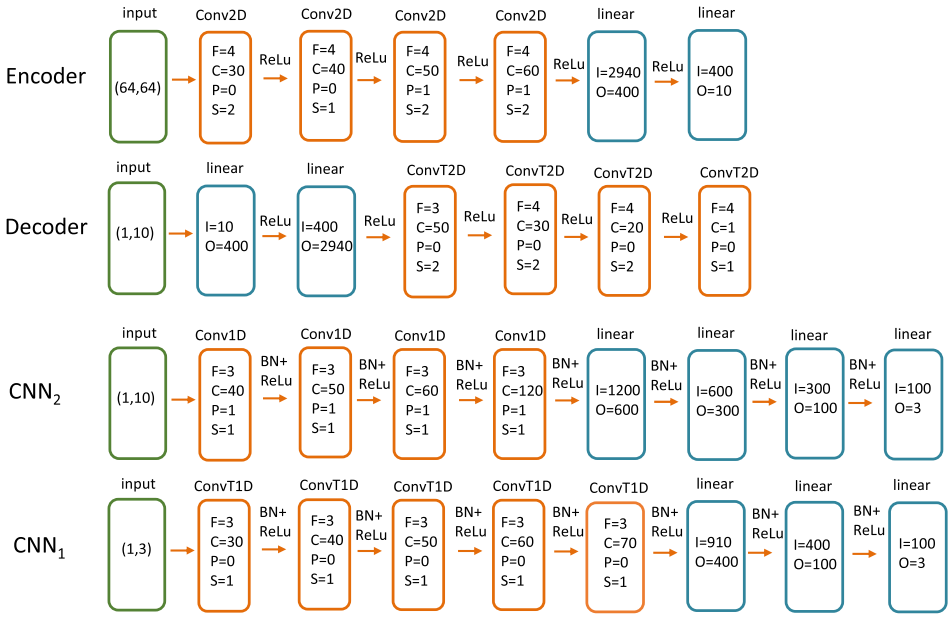

Fig. 10. Detailed parameters of the network structure.

adequate in actual applications. Moreover, a higher design accuracy can be obtained by higher pixel resolution in the image representation.

5. PARAMETERS OF THE DNN

Figure 10 shows the parameters of the DNN architecture. The encoder contains four convolutional (Conv2D) layers and two linear layers. Its input is a 64×64 pixels image with a pixel value of 0 or 1. For convolutional layers, F is the filter size, C is the number of channels, P is the padding size, and S is the stride. I and O are the input and the output for linear layers. ReLu stands for rectified linear unit activation function. The decoder consists of two linear layers and four transposed convolution layers (ConvT2D). Both the decoder and encoder are based on CNN structures developed for 2D image classification [43]. The CNN_2 has an input dimension of (1, 10), and it consists of four convolutional layers (Conv1D) and four linear layers. A batch normalization layer (BN) is applied before the ReLu to increase the convergence and generalization of the network. The CNN_1 has an input dimension of (1, 3), and it consists of five ConvT1D layers and three linear layers. The network is implemented in Python using the Pytorch machine learning library. The parameters of CNN_1 and CNN_2 are based on an early work with similar data structures [22].

6. CONCLUSION

We demonstrate a DNN structure that can design and engineer multiple symmetry-protected BICs by manipulating freeform structures with predefined symmetry. The geometries are represented by latent vectors, which can then be mapped to the photonic property by a property readout network. We demonstrate the on-demand design of three arbitrary BIC frequencies with high accuracy. We also analyze the nature of complicated band inverse and accidental degeneracy when such freeform structures are tuned and scaled continuously, which shows the potential for further discovery and application.

Funding. King Abdullah University of Science and Technology Baseline Fund (BAS/1/1664-01-01); Competitive Research Grants (URF/1/3437-01-01, URF/1/3771-01-01).

Acknowledgment. The authors would like to thank Prof. Andrea Fratalocchi from KAUST for his advice.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. C. W. Hsu, B. Zhen, A. D. Stone, J. D. Joannopoulos, and M. Soljačić, "Bound states in the continuum," Nat. Rev. Mater. **1**, 16048 (2016).
2. C. W. Hsu, B. Zhen, J. Lee, S.-L. Chua, S. G. Johnson, J. D. Joannopoulos, and M. Soljačić, "Observation of trapped light within the radiation continuum," Nature **499**, 188–191 (2013).
3. K. Koshelev, S. Lepeshov, M. Liu, A. Bogdanov, and Y. Kivshar, "Asymmetric metasurfaces with high-Q resonances governed by bound states in the continuum," Phys. Rev. Lett. **121**, 193903 (2018).
4. A. Kodigala, T. Lepetit, Q. Gu, B. Bahari, Y. Fainman, and B. Kanté, "Lasing action from photonic bound states in continuum," Nature **541**, 196–199 (2017).
5. C. Huang, C. Zhang, S. Xiao, Y. Wang, Y. Fan, Y. Liu, N. Zhang, G. Qu, H. Ji, J. Han, L. Ge, Y. Kivshar, and Q. Song, "Ultrafast control of vortex microlasers," Science **367**, 1018–1021 (2020).
6. M. Meudt, C. Bogiadzi, K. Wrobel, and P. Görrn, "Hybrid photonic–plasmonic bound states in continuum for enhanced light manipulation," Adv. Opt. Mater. **8**, 2000898 (2020).
7. F. Yesilkoy, E. R. Arvelo, Y. Jahani, M. Liu, A. Tittl, V. Cevher, Y. Kivshar, and H. Altug, "Ultrasensitive hyperspectral imaging and biodetection enabled by dielectric metasurfaces," Nat. Photonics **13**, 390–396 (2019).
8. Z. Liu, Y. Xu, Y. Lin, J. Xiang, T. Feng, Q. Cao, J. Li, S. Lan, and J. Liu, "High-Q quasibound states in the continuum for nonlinear metasurfaces," Phys. Rev. Lett. **123**, 253901 (2019).

9. L. Xu, K. Zangeneh Kamali, L. Huang, M. Rahmani, A. Smirnov, R. Camacho-Morales, Y. Ma, G. Zhang, M. Woolley, D. Neshev, and A. E. Miroshnichenko, "Dynamic nonlinear image tuning through magnetic dipole quasi-BIC ultrathin resonators," Adv. Sci. **6**, 1802119 (2019).
10. B. Zhen, C. W. Hsu, L. Lu, A. D. Stone, and M. Soljačić, "Topological nature of optical bound states in the continuum," Phys. Rev. Lett. **113**, 257401 (2014).
11. J. Jin, X. Yin, L. Ni, M. Soljačić, B. Zhen, and C. Peng, "Topologically enabled ultrahigh-*Q* guided resonances robust to out-of-plane scattering," Nature **574**, 501–504 (2019).
12. Y. W. Lee, F. C. Fan, Y. C. Huang, B. Y. Gu, B. Z. Dong, and M. H. Chou, "Nonlinear multiwavelength conversion based on an aperiodic optical superlattice in lithium niobate," Opt. Lett. **27**, 2191–2193 (2002).
13. S. Diaz, "Stable dual-wavelength erbium fiber ring laser with optical feedback for remote sensing," J. Lightwave Technol. **34**, 4591–4595 (2016).
14. Y. Zuo, Y. Gao, S. Qin, Z. Wang, D. Zhou, Z. Li, Y. Yu, M. Shao, and X. Zhang, "Broadband multi-wavelength optical sensing based on photothermal effect of 2D MXene films," Nanophotonics **9**, 123–131 (2019).
15. I. Malkiel, M. Mrejen, A. Nagler, U. Arieli, L. Wolf, and H. Suchowski, "Plasmonic nanostructure design and characterization via deep learning," Light Sci. Appl. **7**, 60 (2018).
16. W.-K. Lee, S. Yu, C. J. Engel, T. Reese, D. Rhee, W. Chen, and T. W. Odom, "Concurrent design of quasi-random photonic nanostructures," Proc. Natl. Acad. Sci. USA **114**, 8734–8739 (2017).
17. T. Feichtner, O. Selig, M. Kiunke, and B. Hecht, "Evolutionary optimization of optical antennas," Phys. Rev. Lett. **109**, 127701 (2012).
18. Z. Shi, A. Y. Zhu, Z. Li, Y.-W. Huang, W. T. Chen, C.-W. Qiu, and F. Capasso, "Continuous angle-tunable birefringence with freeform metasurfaces for arbitrary polarization conversion," Sci. Adv. **6**, eaba3367 (2020).
19. A. Y. Piggott, J. Lu, K. G. Lagoudakis, J. Petykiewicz, T. M. Babinec, and J. Vučković, "Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer," Nat. Photonics **9**, 374–377 (2015).
20. K. Yao, R. Unni, and Y. Zheng, "Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale," Nanophotonics **8**, 339–366 (2019).
21. R. S. Hegde, "Deep learning: a new tool for photonic nanostructure design," Nanoscale Adv. **2**, 1007–1023 (2020).
22. R. Lin, Y. Zhai, C. Xiong, and X. Li, "Inverse design of plasmonic metasurfaces by convolutional neural network," Opt. Lett. **45**, 1362–1365 (2020).
23. T. Christensen, C. Loh, S. Picek, D. Jakobović, L. Jing, S. Fisher, V. Ceperic, J. D. Joannopoulos, and M. Soljačić, "Predictive and generative machine learning models for photonic crystals," Nanophotonics **9**, 4183–4192 (2020).
24. Z. Liu, D. Zhu, S. P. Rodrigues, K.-T. Lee, and W. Cai, "Generative model for the inverse design of metasurfaces," Nano Lett. **18**, 6570–6576 (2018).
25. J. Jiang, D. Sell, S. Hoyer, J. Hickey, J. Yang, and J. A. Fan, "Free-form diffractive metagrating design based on generative adversarial networks," ACS Nano **13**, 8872–8878 (2019).
26. F. Wen, J. Jiang, and J. A. Fan, "Progressive-growing of generative adversarial networks for metasurface optimization," arXiv:1911.13029 (2019).
27. S. So and J. Rho, "Designing nanophotonic structures using conditional deep convolutional generative adversarial networks," Nanophotonics **8**, 1255–1261 (2019).
28. T. Cohen and M. Welling, "Group equivariant convolutional networks," in *International Conference on Machine Learning* (2016), pp. 2990–2999.
29. S. Dieleman, J. De Fauw, and K. Kavukcuoglu, "Exploiting cyclic symmetry in convolutional neural networks," arXiv:1602.02660 (2016).
30. I. Nachum and A. Yehudayoff, "On symmetry and initialization for neural networks," arXiv:1907.00560 (2019).
31. R. Huang, S. Zhang, T. Li, and R. He, "Beyond face rotation: global and local perception GAN for photorealistic and identity preserving frontal view synthesis," in *IEEE International Conference on Computer Vision* (2017), pp. 2439–2448.
32. I. Peleg and L. Wolf, "Structured GANs," in *IEEE Winter Conference on Applications of Computer Vision (WACV)* (2018), pp. 719–728.
33. S. G. Johnson and J. D. Joannopoulos, "Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis," Opt. Express **8**, 173–190 (2001).
34. I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. M. Botvinick, S. Mohamed, and A. Lerchner, "Beta-VAE: learning basic visual concepts with a constrained variational framework," in *International Conference on Learning Representations* (2017), pp. 1–13.
35. C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, "Understanding disentangling in β-VAE," arXiv:1804.03599 (2018).
36. K. Sakoda, "Universality of mode symmetries in creating photonic Dirac cones," J. Opt. Soc. Am. B **29**, 2770–2778 (2012).
37. Y. Li, Y. Wu, X. Chen, and J. Mei, "Selection rule for Dirac-like points in two-dimensional dielectric photonic crystals," Opt. Express **21**, 7699–7711 (2013).
38. X. Huang, Y. Lai, Z. H. Hang, H. Zheng, and C. T. Chan, "Dirac cones induced by accidental degeneracy in photonic crystals and zero-refractive-index materials," Nat. Mater. **10**, 582–586 (2011).
39. M. Minkov, I. A. D. Williamson, M. Xiao, and S. Fan, "Zero-index bound states in the continuum," Phys. Rev. Lett. **121**, 263901 (2018).
40. L. Lu, J. D. Joannopoulos, and M. Soljačić, "Topological photonics," Nat. Photonics **8**, 821–829 (2014).
41. J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light*, 2nd ed. (Princeton University, 2011).
42. D. Liu, Y. Tan, E. Khoram, and Z. Yu, "Training deep neural networks for the inverse design of nanophotonic structures," ACS Photon. **5**, 1365–1369 (2018).
43. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," Commun. ACM **60**, 84–90 (2017).