

Unsupervised image translation with distributional semantics awareness

Zhexi Peng¹, He Wang², Yanlin Weng¹ (✉), Yin Yang³, and Tianjia Shao¹

© The Author(s) 2023.

Abstract Unsupervised image translation (UIT) studies the mapping between two image domains. Since such mappings are under-constrained, existing research has pursued various desirable properties such as distributional matching or two-way consistency. In this paper, we re-examine UIT from a new perspective: distributional semantics consistency, based on the observation that data variations contain semantics, e.g., shoes varying in colors. Further, the semantics can be multi-dimensional, e.g., shoes also varying in style, functionality, etc. Given two image domains, matching these semantic dimensions during UIT will produce mappings with explicable correspondences, which has not been investigated previously. We propose *distributional semantics mapping* (DSM), the first UIT method which explicitly matches semantics between two domains. We show that distributional semantics has been rarely considered within and beyond UIT, even though it is a common problem in deep learning. We evaluate DSM on several benchmark datasets, demonstrating its general ability to capture distributional semantics. Extensive comparisons show that DSM not only produces explicable mappings, but also improves image quality in general.

Keywords generative adversarial networks (GANs); manifold alignment; unsupervised learning; image-to-image translation; distributional semantics

1 Introduction

Unsupervised image translation (UIT) has been intensively studied in recent years. Many applications have been inspired by its ability to create mappings between two image domains. Since there can be theoretically an infinite number of mappings between two domains, UIT is by nature an under-constrained problem. Naturally, different approaches have been developed to ensure certain desirable properties, such as shared latent spaces [1], two-way consistency [2], pair-wise distance preservation [3], and image semantics [4]. While existing researches tend to focus on general distributional matching [3, 5], we aim to investigate a rarely examined perspective: the distributional semantics during UIT.

We define distributional semantics as the visually understandable variations between samples (not within a single sample). Shoes vary in color, style (e.g., low/high collars), and functionality (e.g., sneakers/high heels). Similarly, bags also vary in color, style (e.g., with/without handles), and functionality (e.g., purses/backpacks). During UIT, we argue that it is not enough to simply translate images. It is desirable for such distributional semantics to be maintained, e.g., red high-collar high heels map to black purses with handles, while white low-collar sneakers map to blue backpacks. This is exactly the goal of this research.

To maintain distributional semantics during UIT, two critical problems must be addressed. The first question is what semantics should be maintained. As we are considering unsupervised learning, no labeling should be required, which means that the data variations (the distributional semantics) should be characterized without prior knowledge, yet be interpretable by humans. Secondly, distributional semantics are rarely

1 State Key Lab of CAD&CG, Zhejiang University, Hangzhou 310058, China. E-mail: Z. Peng, zhaxipeng@zju.edu.cn; Y. Weng, weng@cad.zju.edu.cn (✉); T. Shao, tjshao@zju.edu.cn.

2 School of Computing, University of Leeds, Leeds, UK. E-mail: H.E.Wang@leeds.ac.uk.

3 School of Computing, Clemson University, Clemson, USA. E-mail: yin5@clemson.edu.

Manuscript received: 2022-03-01; accepted: 2022-05-15



considered in general in deep learning, because data are usually transformed numerous times in the model. To maintain such semantics, we need a mechanism to ensure that the data distribution remains as undistorted as possible, especially in the dimensions of the semantics of interest, during transformations mapping between two domains.

In this paper, we propose a novel deep learning method called *distributional semantics mapping* (DSM). Given two image datasets $\mathbf{A} = \{x_i\}$ and $\mathbf{B} = \{y_j\}$, and a desire to characterize visual semantics in an unsupervised manner, we find that the covariance structure of the data naturally reflects important visual semantics. We choose *principal component analysis* (PCA) to characterize the covariance structure: Härkönen et al. [6] have already demonstrated that PCA applied in feature space can produce interpretable controls for image synthesis. Our approach is agnostic to specific network architectures and consists of three key modules. The first is a semantics-preserving transformation e where the variations of x s in a direction (e.g., the first principal component of \mathbf{A}) must be consistent with the variations of latent vectors $z_x = e(x)$ in its corresponding direction (e.g., the first principal component of $e(\mathbf{A})$) in the latent space. We use two such encoders e_A and e_B to project \mathbf{A} and \mathbf{B} into a shared latent space. The second module aligns the key dimensions of $e_A(\mathbf{A})$ and $e_B(\mathbf{B})$. The last module is a decoder/generative network g which decodes y by $y_{\text{recon}} = g(e_B(y))$ and translates x by $y_{\text{trans}} = g(e_A(x))$.

As far as we know, this is the first approach for UIT that preserves distributional semantics. We identify the importance of preserving distributional semantics, which has a wide range of implications for image translation and beyond. We propose a new approach that helps preserve distribution semantics during image transformations.

2 Related work

2.1 Generative adversarial networks

Generative adversarial networks (GANs) [7] have achieved great success for a fast-growing number of computer vision tasks, including image generation [8, 9], image colorization [10], image inpainting [11], and image super-resolution [12]. Conditional GANs [13] can be used to perform image-to-image translation

[2, 14–16]. Recently, some interactive systems have been proposed using GANs for real-time portrait image editing [17, 18]. Our work also utilizes GANs conditioned on an input image, but it does not rely on any specific GAN model. To validate its generality, we employ two widely used GAN models: LSGAN [19] and NSGAN [7].

2.2 Image-to-image translation

To perform image-to-image translation, early methods such as pix2pix [10, 20] often required the networks to be trained with paired training data. Recently, a variety of approaches [14, 15] have been proposed to learn the image translation from unpaired data. For example, CycleGAN [2] leverages cycle consistency to constrain the mapping. Lu et al. [21, 22] show that optimal transport costs can improve the generative network. UNIT [1] assumes that two image domains can share the same latent space. By decomposing the image into the style (domain-specific) code and content (domain-invariant) code, MUNIT [23] and DRIT [24] can synthesize diverse outputs from an input image. Alami Mejjati et al. [25] and Kim et al. [26] improve translation results using an attention mechanism. Choi et al. [27] propose StarGAN which can perform image translation for multiple domains using a single GAN. More recently, DRIT++ [28] extends DRIT to support multiple domains, while StarGAN v2 [29] extends StarGAN to generate diverse images across multiple domains. FUNIT [30] can work on previously unseen target classes given only a few example images. To enable unsupervised one-sided mapping, Benaim and Wolf [3] present DistanceGAN that maintains the distances between images, and Fu et al. [31] employ other geometric constraints (e.g., orientation). Our method differs from existing methods in that it explicitly preserves and matches the distributional semantics in domains during UIT, which generates explicable mappings between images.

3 Methodology

3.1 Approach

Given two image datasets $\mathbf{A} = \{x_i\}$ and $\mathbf{B} = \{y_j\}$, we aim to compute a mapping $M: \mathbf{A} \rightarrow \mathbf{B}$, so that the distributional semantics of \mathbf{A} are aligned with those of \mathbf{B} ; we use principal components (PCs) to describe the semantics. First, we project \mathbf{A} using an

encoder e_A to $z_A = \{z_{x_i}\}$ where $z_{x_i} = e_A(x_i)$ is a latent vector. We keep the distributional semantics of A and z_A aligned during the encoding process. A similar projection is applied to B using an encoder e_B to get $z_B = \{z_{y_i}\}$. We denote the PCs of the two latent distributions as $V_A \in \mathbb{R}^{P \times P}$ and $V_B \in \mathbb{R}^{P \times P}$ where P is the dimensionality of the latent vectors. Next we ensure the top k PCs of z_A are aligned with those of z_B . Finally, we use a generative network g to reconstruct B using $y_{\text{recon}} = g(e_B(y))$ and translate A by $y_{\text{trans}} = g(e_A(x))$. The model is shown in Fig. 1. Below, we give the details of the components.

3.2 Semantics-preserving transformation

The images need to go through a sequence of transformations during translation, which in deep learning are usually some encoding processes such as convolutions. However, the shape of the latent distribution seems to be rarely considered in terms of its semantic consistency with the data distribution itself. Existing efforts such as imposing a prior distribution [32] or geometric constraints [3] are mostly aimed at encouraging the latent distribution to behave well, rather than to match the data distribution. As a result, current encoders may not be able to preserve the distributional semantics, as we confirm in experiments. We, therefore, introduce a new general autoencoding scheme to preserve the distributional semantics:

$$z_x = e(x), \quad x_{\text{recon}} = d(z_x)$$

subject to $Ux = Vz_x$, U, V have K rows (1)

where x is a data sample, and e and d are encoding and decoding schemes respectively. U and V are the first K PCs of A and z_x . The autoencoder can be trained by, e.g., minimizing $\sum \|x - x'\|_2^2$. Equation (1) states the key difference between our autoencoder and existing autoencoders: it requires the projections of x on the data-space PCs U to be equal to the projections of z_x onto the latent-space PCs V . Note that a standard dimensionality reduction using PCA




Fig. 1 DSM framework. Two image domains are aligned along their data-space PCs via a shared latent space before translation.

is a special case of Eq. (1), when $V = U$ and $e(A) = VA$. Equation (1) is more general because it does not dictate what V is, nor does it require the latent space to have fewer dimensions than the data space. Equation (1) only requires the covariance structure to persist when encoding along the first K PCs in both the data and latent space. One key question is why we do not just set $V = U$. This is because we need the flexibility of encoding data into an arbitrary V during UIT while keeping the general shape of the data distribution, as explained later.

Equation (1) is general but difficult to optimize because it needs to be computed over the whole dataset, with high memory requirements, and V is unknown. Therefore, we propose a local scheme which keeps the global alignment by enforcing local alignments on samples:

$$\left\{ \begin{aligned} L_{\text{localAlign}} &= \frac{1}{N} \sum_{i=0}^{N-2} (L_{\text{angle}} + L_{\text{norm}}) \\ L_{\text{angle}} &= |\cos(x_i, x_{i+1}) - \cos(z_{x_i}, z_{x_{i+1}})| \\ L_{\text{norm}} &= \frac{\|x_i\|_2}{\|x_{i+1}\|_2} - \frac{\|z_{x_i}\|_2}{\|z_{x_{i+1}}\|_2} \end{aligned} \right. \quad (2)$$

where $\cos()$ is the cosine distance and N is the batch size. L_{angle} and L_{norm} are computed on vectorized data samples and corresponding latent vectors respectively. The goal of $L_{\text{localAlign}}$ is that, given any two data samples, their length ratio and angle should remain the same after projection into the latent space. L_{angle} aims to keep the overall shape of the distribution when the data are projected into the latent space, while L_{norm} allows scaling but prefers uniform scaling. $L_{\text{localAlign}}$ has the effect of preserving the covariance structure of the data during transformation. While we only apply local alignment constraints between a certain number of pairs in each batch, we randomly sample various batches in each training epoch, constraining further distinct pairs in the process. Furthermore, the locality of this approach, only considering two samples a time, essentially ensures that the covariance structure is invariant under homogeneous transformations of the basis in the latent space, which allows the autoencoder to choose the optimal V that aids reconstruction while preserving the covariance structure.

3.3 Semantics-based manifold alignment

Given z_A and z_B , we have two latent distributions with their respective covariance structures chara-

cterized by their latent space PCs $V_A \in \mathbb{R}^{P \times P}$ and $V_B \in \mathbb{R}^{P \times P}$. To ensure that the UIT maintains the distributional semantics, we need to align z_A and z_B , e.g., by aligning the direction in which z_A shows the biggest variation with that of z_B by aligning their first PCs. Further, to maintain visual semantics in multiple dimensions, we should align the top K PCs in V_A and V_B . Several alternative methods are possible, such as aligning V_A and V_B directly, or fixing one and aligning the other to it. Since both V_A and V_B are unknown, directly aligning V_A and V_B corresponds to minimizing

$$L_{\text{latentA}} + L_{\text{latentB}} + L_{\text{align}}$$

where

$$\begin{cases} L_{\text{latentA}} = \frac{1}{N_A} \sum_{i=0}^{N_A-1} \|z_{x_i} - V_A V_A^T z_{x_i}\|_2^2 \\ L_{\text{latentB}} = \frac{1}{N_B} \sum_{j=0}^{N_B-1} \|z_{y_j} - V_B V_B^T z_{y_j}\|_2^2 \\ L_{\text{align}} = \frac{1}{K} \sum_{k=0}^{K-1} \|V_A^k - V_B^k\|_2^2 \end{cases} \quad (3)$$

N_A and N_B are the numbers of images in A and B respectively. V_A^k and V_B^k are the k th PCs of V_A and V_B , respectively. Experimentally, we have found that allowing both V_A and V_B to change causes the optimization to settle into local minima, so we fix V_A and align V_B with it. This also means that e_A can be pre-trained and we can compute V_A from z_A using PCA.

Aligning V_B to V_A still presents challenges because directly learning V_B requires to simultaneously transform all z_{y_j} s which is again equivalent to operating on the whole of B . This is a similar difficulty to the one in Section 3.2. Again, we operate only on a batch of N samples to ensure the global alignment of V_A and V_B :

$$\begin{aligned} &L_{\text{maniAlign}} \\ &= \frac{1}{N} \sum_{j=0}^{N-1} \|z_{y_j} - V_{A[0:K-1]} V_{A[0:K-1]}^T z_{y_j}\|_2^2 + \\ &ke^{-\alpha \|z_{y_j}\|_2^2} \end{aligned} \quad (4)$$

where $z_{y_j} \in z_B$ and $V_{A[0:K-1]} \in \mathbb{R}^{P \times K}$ contains the first K PCs of interest of z_A . $L_{\text{maniAlign}}$ requires z_{y_j} to be reconstructable after projecting them into the basis of z_A , which essentially encourages the covariance structure of z_B to be similar to that of z_A , and the two bases to be aligned. The term

$k \exp(-\alpha \|z_{y_j}\|_2^2)$ prevents z_{y_j} from shrinking, leading to a trivial solution of Eq. (4). Since the covariance structures of A and B are kept in z_A and z_B via Eq. (2), $L_{\text{maniAlign}}$ completes the semantics-based alignment of two domains.

3.4 Simultaneous decoding and translation

After semantics-based manifold alignment, we transform z_A and z_B into the space of B to finish the UIT. Since z_A and z_B are aligned, we combine the reconstruction and translation tasks using a single network g which serves both as a decoder and a translator: the general shapes of the two latent distributions are similar after alignment. When the decoder has been trained with the reconstruction loss on z_B , it has already to some extent learnt to translate z_A . We reconstruct B using $y_{\text{recon}} = g(z_y)$. Meanwhile, we treat g as a generator in a generative adversarial network using $y_{\text{trans}} = g(z_x)$ and use a discriminator network $h(y_{\text{trans}}, y) = [0, 1]$ to further improve the translation.

Overall, given a pre-trained e_A and z_A , and hence also V_A , we minimize the objective function in Eq. (5):

$$\begin{aligned} L = &\omega_1 \frac{1}{N_B} \sum_{i=1}^{N_B-1} \|y_i - y_{\text{recon}}\|_2^2 + \omega_2 L_g + \omega_3 L_d + \\ &\omega_4 L_{\text{localAlign}}^B + \left(1 - \sum_{i=1}^4 \omega_i\right) L_{\text{maniAlign}} \end{aligned} \quad (5)$$

where L_g and L_d are the GAN loss that depends on the chosen GAN model. N_B is the total number of images in B and ω_i are weights. In $L_{\text{localAlign}}^B$, we apply Eq. (2) to both the y -to- z_y , and z_x -to- y_{trans} mappings. Further details are in the Appendix.

4 Implementation details

We pre-train an autoencoder for dataset A with $L_{\text{localAlign}}$ to get e_A and calculate V_A from z_A using PCA. For e_B and g , we adopt the network architectures from UNIT. In all experiments, we set $\omega_1 = 0.033, \omega_2 = \omega_3 = 0.333$, and other weights according to the experiment. See the Appendix for details. e_A is trained for 100 epochs and the remainder are trained for 300 epochs on all datasets, using Adam [33] with a batch size of 16, a learning rate of 0.0001, and exponential decay rates $(\beta_1; \beta_2) = (0.5; 0.999)$. All experiments were conducted using 2 NVIDIA GTX 1080 Ti GPUs, and PyTorch. Training took 6–18 hours.

5 Experiments

5.1 Data

We employed several benchmark datasets to validate our method, including SummerWinter [2], CatDog [24], and ShoeHandbag [34, 35]. We also built a very challenging dataset, MMISTHandbag, using hand-drawn digits from MNIST [36], and handbags randomly sampled from Ref. [35]; it has two distinctive distributions and distributional semantics. All comparative results in this section were computed using LSGAN. See the Appendix for details and further results.

5.2 Evaluation metrics

In addition to visual evaluation, we employ the Frechet inception distance (FID) [37] as a quantitative measure. However, there is no good metric to measure the faithfulness of the preservation of distributional semantics. We, therefore, propose a new evaluation metric called the *ordering-tolerance curve* (OTC). Given images $x \in \mathbf{A}$ and their translations $y_{\text{trans}} \in \mathbf{B}$, we define the OTC as

$$c = \frac{1}{N_A} \text{card}(\{x | d(x, y_{\text{trans}}) / N_A \leq \beta, x \in \mathbf{A}\}) \quad (6)$$

where $\beta \in [0, 1]$, $\text{card}(\cdot)$ is the cardinality of a set, $d(x, y_{\text{trans}}) = |\text{rank}(y_{\text{trans}}) - \text{rank}(x)|$, $\text{rank}(x)$ is the rank of x in \mathbf{A} along a chosen PC among all data samples in \mathbf{A} and N_A is the total number of data samples in \mathbf{A} , and $\text{rank}(y_{\text{trans}})$ is the rank of y_{trans} along a chosen PC in \mathbf{B} . $d(x, y_{\text{trans}})$ is equal to zero if the rank of x is retained during translation, and non-zero otherwise (the larger the worse). c is the *percentage of correctly ordered* x values whose normalized rank errors are within β , the *ordering error tolerance*.

5.3 Semantics-preserving transformation

A key contribution of our work is a straightforward but extremely effective transformation scheme that preserves distributional semantics. To show the necessity of such transformations in UIT, we trained an autoencoder on cat images in CatDog then computed the 1st PCs of A and z_A . The autoencoder is based on the encoder and decoder of UNIT. Please see the Appendix for further details. We then ranked all images along the two PCs: see Fig. 2.

In the original data, the variation on the 1st PC is mainly a color transition from light to dark (see Fig. 2(top)). However, without semantics-preserving




Fig. 2 Top: original data. Middle: latent samples of DSM. Bottom: latent samples of a standard autoencoder. Left to right: images ranked 4th, 19th, 24th, 37th, 44th, 52nd, 94th, and 97th on the first PCs of the data (top) and latent space (middle, bottom).

transformations, the shape of the latent distribution is changed, and unable to preserve the visual semantics (Fig. 2(bottom)). In contrast, DSM preserves the distributional semantics (Fig. 2(middle)). We also show the OTCs in Fig. 3(left): DSM can contain the rank error to under 3% while a standard autoencoder fails systematically. Although we only show the results from a specific autoencoder, our preliminary experiments showed this to be a common problem of autoencoders.

5.4 Architectural studies

Two main adjustable components of DSM are the GAN architecture and the number of PCs, K . We first evaluated DSM on the first K PCs on all datasets; we present the FID scores in Table 1. While we expected that a larger K would result in a harder optimization problem and hence lower quality, the results show that it depends on the dataset. After considering the data, this is understandable because different datasets have different variance distributions over the PCs: some have large variance on the 1st PC while others have variances spread over the first K PCs, which affects the behavior of Eq. (4). Although control can be added, e.g., by adding weights to different PCs, we choose not to do so and let DSM adapt to the data. We only tested DSM up to $K = 3$: while DSM can work for $K \geq 3$, it is typically hard for humans to visually understand the semantics in

Table 1 FID scores on four datasets with the first K PCs aligned. When $K = 0$, DSM is trained without alignment loss

Dataset	$K = 0$	$K = 1$	$K = 2$	$K = 3$
CatDog	64.5	58.0	96.2	84.5
SummerWinter	100.9	90.4	107.6	99.5
ShoeHandbag	123.6	129.7	155.5	128.9
MNISTHandbag	172.7	180.6	149.3	171.8




Fig. 3 OTCs in various experiments. Left: OTC along PC0 with/without semantics-preserving transformation on CatDog. Middle: LSGAN vs. NSGAN along PC0 on ShoeHandbags. Right: OTC along PC0 for four datasets.

PCs where $K \geq 3$. We also compared two GAN architectures LSGAN and NSGAN on ShoeHandbags and show the OTCs in Fig. 3(middle). Both keep the rank error to within around 20%, showing that alignment constraints can improve translation ability visibly for differing GAN architectures. Further results can be found in the Appendix.

5.5 Image quality in translation

Although normally the first K PCs bear visually understandable semantics, the specific value of K depends on the dataset. The first PC is almost always visually interpretable across a dataset, and some datasets have meaningful variations on other PCs. We show the images from the four datasets ranked along their respective meaningful PCs in Fig. 4. For CatDog, SummerWinter, and ShoeHandbags, the first

PC (PC0) shows color variations from light to dark, for MNIST, shape varies from slim to round, while for ShoeHandbags, the second PC (PC1) shows shoe collar height variation for shoes and handle length variation for handbags.

In Fig. 5, we show our results for mappings of cat-to-dog, summer-to-winter, shoes-to-handbags, and digits-to-handbags. PC0s of CatDog and SummerWinter are color variations (light to dark). The major difference is that the color variations in CatDog are clearly separated into foreground (faces) and background while there is no such separation in SummerWinter. In both cases, DSM successfully translates the images with high quality while simultaneously matching the semantics. In cat-to-dog, DSM transforms the faces and maintains



Fig. 4 Images ranked according to PCs. Top to bottom: CatDog (PC0), SummerWinter (PC0), ShoeHandbags (PC0), MNIST (PC0), and ShoeHandbags (PC1).



Fig. 5 Mapping results. Top to bottom: CatDog, SummerWinter, ShoeHandbag, and MNISTHandbag, ordered along PC0. In each pair of rows, the upper row shows the input images and the lower row shows the translated images.

the separation of the foreground and background, while in summer-to-winter, the changes are more heterogeneous depending on the scenes; DSM lays snow on different landscapes. In shoes-to-handbags, unlike CatDog and SummerWinter, the color variation is restricted to the object itself, and DSM faithfully keeps the semantics. Finally, to push DSM further, we tested a digit-to-handbag translation. These two datasets have distinctive distributions. The results show that long slim digits are translated to handbags in light colors while fat round digits are translated to handbags in dark colors, which are consistent with their respective PC0s in Fig. 4. We also show the OTCs in Fig. 3(right).

5.6 Comparisons

We compare our model to UNIT, CycleGAN, DistanceGAN, and DRIT++ on the CatDog and ShoeHandbag datasets, by using the public code shared by the authors of these methods. We first give FID scores for all methods in Table 2 and OTCs in Fig. 6. By aligning two data manifolds based on their semantics, DSM is able to improve the translation quality for both datasets. The OTCs show clearly that DSM can keep the semantics on PC0 better than the other methods by containing the rank error to within around 30% and 18% (see Fig. 6). The second best methods contain it to roughly only 50% and 40%. On PC1 of ShoeHandbag (see Fig. 6(right)),

Table 2 FID scores on CatDog and ShoeHandbag datasets for various methods

Method	CatDog	ShoeHandbag
Ours	58.00	128.93
CycleGAN	95.12	135.55
DistanceGAN	63.60	185.28
DRIT++	61.95	185.01
UNIT	64.62	140.88

CycleGAN is close to DSM. However, we argue that its behavior is inconsistent: see Figs. 6(left) and 6(middle) where the variance is large and contains large amounts of semantic information.

Visual comparisons can be found in Fig. 7. Overall, DSM generates images of higher visual quality. Additionally, other methods are incapable of preserving or matching the distributional semantics. In Fig. 7(left), the major variation of the input images from left-to-right is a color variation, light-to-dark. While DSM obviously keeps the same variations during UIT, it is hard to find similar effects in other methods. Similar observations can also be made in Fig. 7(middle). To further explore semantics in other PCs, we show Fig. 7(right). While the input varies from slippers to sneakers, our results generate handbags varying from those without handles to those with handles. In contrast, other methods struggle to generate consistent semantic variation. Further results can be found in the Appendix.




Fig. 6 OTC scores for various methods. Left: CatDog. Middle: Shoe-to-Handbags on PC0 of Handbags. Right: Shoe-to-Handbags on PC1 of Handbags.



Fig. 7 Left, middle: ordering along PC0 of Dogs, Handbags, respectively. Right: ordering along PC1 of Handbags.

6 Discussions and conclusions

While PCA is a straightforward way of characterizing distributional semantics, our method can incorporate alternative techniques such as kernel PCA, multidimensional scaling, and others that can define the covariance structure by “flattening” the data manifold before performing DSM.

We have only evaluated DSM up to $K = 3$ because the semantics start to lack visual meaning for $K > 3$. However, we argue that DSM is effective and useful for two reasons. Firstly, for almost all datasets, the first PC bears the majority of the variance, and the distributional semantics captured by the variance are always visually interpretable. Secondly, aligning the PCs of the two distributions during translation increases image quality.

In summary, we have proposed DSM, the first UIT method which preserves and matches the distributional semantics of two image domains. It is straightforward and effective, as demonstrated on multiple datasets, and capable of improving translation quality compared to the state-of-the-art. DSM is also general in its capacity to incorporate any GAN and autoencoder model. In future, we will incorporate human-labelling in a semi-supervised setting of DSM where humans can arbitrarily decide the semantics by ranking images. This will enable DSM to encode arbitrary semantics and open it up to many other applications.

Appendix A Implementation details

A.1 Loss functions

Equation (5) contains a local alignment loss $L_{\text{localAlign}}^B$, which applies Eq. (2) to both y -to- z_y , and z_x -to- y_{trans} . Applying Eq. (2) to y -to- z_y is straightforward, and ensures semantics-preserving

transformation between B and z_B , just as does the encoder of A . Here we explain why applying Eq. (2) to z_x -to- y_{trans} is essential. As Section 3.4 explains, we align z_A and z_B , and use a single network g to serve as both the decoder for z_B and the generator for z_A . This mechanism provides an implicit constraint on the translated images $\{y_{\text{trans}}\}$ from z_A so that it has a similar distribution to the reconstructed images $\{y_{\text{recon}}\}$ from z_B . However, as there is an additional GAN loss which modifies the distribution of the translated images, the alignment of $\{y_{\text{trans}}\}$ and $\{y_{\text{recon}}\}$ may be affected, compromising the distributional semantics matching. Thus, we apply Eq. (2) to z_x -to- $\{y_{\text{trans}}\}$ to explicitly preserve the semantics:

$$\omega_4 L_{\text{localAlign}}^B = \omega_4^e L_{\text{localAlign}}^e + \omega_4^g L_{\text{localAlign}}^g \quad (7)$$

where $L_{\text{localAlign}}^e$ and $L_{\text{localAlign}}^g$ are the loss terms for y -to- z_y and z_x -to- y_{trans} . All experiments set $\omega_1 = 0.033, \omega_2 = \omega_3 = 0.333, \omega_4^g = 0.167$ in Eq. (5) of the paper. We set $\omega_4^e = 0.066$ in ShoeHandbag align 3PCs experiments and $\omega_4^e = 0.05$ in all other experiments. We set $k = 15, \alpha = 10^{-6}$ for the regularization term in Eq. (4).

A.2 Network architecture

Our image translation network architecture is based on the one from UNIT [1]. For the encoder, to ensure the latent vector size is smaller than the input image size in the case of resolution 256×256 , we halve the kernel numbers in the second and third convolutional layers and add one further convolutional layer (see Table 3 for details). Furthermore, instead of using instance normalization (IN) [38], we utilize batch normalization (BN) [39] in the encoder and generator. For the discriminator network, we employ spectral normalization [40], and multi-scale discriminators at 3 scales. The network architecture is given in Table 3. We use the following abbreviation for ease of

Table 3 Network architecture for the distributional semantics mapping experiments

Layer	Encoder	Generator	Discriminator
1	CONV(N64,K3,S1),RELU	CONV(N512,K3,S1),RELU	CONV(N64,K4,S2),LeakyReLU
2	CONV(N64,K3,S2),RELU	RESBLK(N512,K3,S1)	CONV(N128,K4,S2),LeakyReLU
3	CONV(N128,K3,S2),RELU	RESBLK(N512,K3,S1)	CONV(N256,K4,S2),LeakyReLU
4	CONV(N128,K3,S2),RELU	RESBLK(N512,K3,S1)	CONV(N512,K4,S2),LeakyReLU
5	RESBLK(N128,K3,S1)	RESBLK(N512,K3,S1)	CONV(N512,K1,S1),LeakyReLU
6	RESBLK(N128,K3,S1)	UP+CONV(N256,K3,S1),RELU	
7	RESBLK(N128,K3,S1)	UP+CONV(N128,K3,S1),RELU	
8	RESBLK(N128,K3,S1)	UP+CONV(N64,K3,S1),RELU	
9		CONV(N3,K7,S1),TanH	

presentation: N: number of kernels number, K: kernel size, S: stride. UP indicates a 2 nearest-neighbor upsampling layer and RESBLK, a residual basic block. The detailed network architecture of the pre-trained autoencoder for \mathbf{A} is given in Table 4.

Table 4 Network architecture for the autoencoder

Layer	Encoder	Decoder
1	CONV(N64,K3,S1),RELU	RESBLK(N128,K3,S1)
2	CONV(N64,K3,S2),RELU	RESBLK(N128,K3,S1)
3	CONV(N128,K3,S2),RELU	RESBLK(N128,K3,S1)
4	CONV(N128,K3,K2),RELU	RESBLK(N128,K3,S1)
5	RESBLK(N128,K3,S1)	UP+CONV(N128,K3,S1),RELU
6	RESBLK(N128,K3,S1)	UP+CONV(N128,K3,S1),RELU
7	RESBLK(N128,K3,S1)	UP+CONV(N64,K3,S1),RELU
8	RESBLK(N128,K3,S1)	CONV(N3,K7,S1),TanH

Appendix B Further experimental results

B.1 Data details

For all datasets, images were resized to 256×256 . In CatDog, 871 cat (birman) and 1364 dog (husky, samoyed) images were randomly divided into 771 (cat) and 1264 (dog) for training and the remainder used for testing. SummerWinter comprises 1540 summer photos and 1200 winter photos, which were randomly divided into 1231 (summer) and 962 (winter) for training and the remainder used for testing. For ShoeHandbag, we randomly sampled images from edges2shoes and edges2handbags, using 3726 (shoe) and 3822 (handbag) for training, and 101 (shoe) and 178 (handbag) for testing. For MNISTHandbag, 1600 MNIST images and 1600 handbag images were randomly selected from MNIST and edges2handbags, with 1500 of each for training and 100 for testing. We show images from the ShoeHandbag dataset along the first 3 PCs in Fig. 8. This dataset is very challenging, as the distributions of MNIST and handbags are very different. The top five variances ratios along PCs in MNIST are [0.095, 0.071, 0.066, 0.052, 0.047], while in handbags these ratios are [0.274, 0.115, 0.065, 0.054, 0.026].

B.2 Our OTCs

Figure 10 shows OTCs for our method on the CatDog, MNISTHandbag, and SummerWinter datasets, with up to 3 PCs aligned. We can see that our method



Fig. 8 Top: Shoe dataset. Bottom: Handbag dataset. Left to right: the images ordered along the corresponding PC.

always keeps semantics best for PC0, and the rank errors become larger on PC1 and PC2. This is mainly because the variance ratio along PC0 is always much larger than along other PCs. For example, in the Handbag dataset, the ratios of variances along the first 3 PCs are about 4:2:1. As a result, the network prefers to perform alignment along PC0 in preference to minimize the total loss. We also note that semantics preservation varies largely across different datasets. In the challenging MNISTHandbag dataset which has two distinctive distributions, our method preserves semantics well along PC0, while in the SummerWinter dataset, our method is capable of preserving semantics on the first 3 PCs. Although control could be used, e.g., via weights to enforce the alignment of multiple PCs, we choose not to do so and make DSM adapt to data, as the distribution of variances on different PCs is an intrinsic property of the data itself which should be respected during translation.

We evaluate semantics preservation only for the first 3 PCs for two main reasons. Firstly, in most image datasets, compared to the variances on the first 3 PCs, the variances on the remaining PCs are very small. For example, in the Summer dataset, even the sum of variances on the 4th–20th PCs is smaller than the variance of PC0. Enforcing alignment along directions with very small variations adds complexity to the optimization while decreasing the explicability of the mapping. Secondly, by investigating various popular

datasets (e.g., Shoes, Handbags, Cars, Animals, Faces, Art works), we discovered that while people can easily perceive semantics on the first PC in all datasets, they can only do so on the second PC in the Shoe, Handbags, and MNIST datasets. People cannot perceive any semantics on the 4th and subsequent PCs. Hence we focus on the first 3 PCs.

B.3 Comparison of LSGAN and NSGAN

Figure 9 and Table 5 compare using LSGAN and NSGAN for image translation on the ShoeHandbag dataset when aligning the first 1–3 PCs. The OTCs show that the two GAN models result in very similar semantics preservation in all cases, demonstrating that our method does not rely on a specific GAN

model and can preserve semantics using different GAN models. We also note that the FIDs of NSGAN are higher than that for LSGAN. This is mainly because the image generation capability of NSGAN is weaker than that of LSGAN. Employing other GAN models such as StyleGAN [9] can improve the FID scores.

B.4 Alignment of PCs

Equation (4) in the paper requires the first K PCs of two domains to be aligned. However, it does not specify the order of alignment. In other words, it does not specify if PC1 of the first domain should be aligned with PC1 of the second domain, etc. We have two choices. The first is to enforce order, PC0-to-PC0, PC1-to-PC1, and so on. However, we find this to be sub-optimal in the sense that Eq. (4) is affected by distribution of variances across PCs. For a dataset with a majority of variance on PC0, the alignment forces due to Eq. (4) have little effect on PC1 and subsequent PCs. We argue that they should be small

Table 5 FID scores on ShoeHandbag with the first K PCs aligned

	$K = 1$	$K = 2$	$K = 3$
LSGAN	129.7	155.6	128.9
NSGAN	172.8	187.0	154.9




Fig. 9 LSGAN/NSGAN.

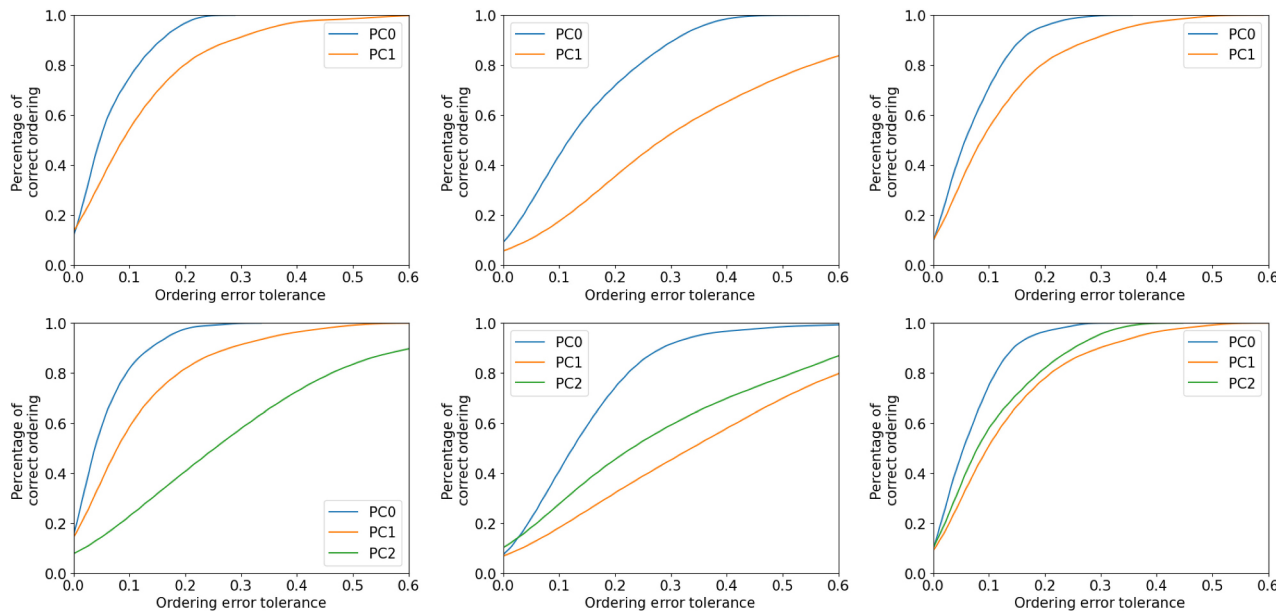


Fig. 10 Above: Aligning PC0 and PC1. Below: Aligning PC0–2. Left to right: CatDog, MNISTHandbag, and SummerWinter.

because the dataset's variance on PC1 and higher PCs is less explicable. Hence, the force provided by Eq. (4) should naturally follow variance distribution. We therefore do not enforce order of alignment of PCs for two datasets. As a result, while PC0 is always mapped to PC0 in all experiments, sometimes PC1 of one dataset can be mapped to PC2 of another. We argue that such a mapping is still valid because it is explicable and reflects the distributional semantics.

Acknowledgements

We thank the anonymous reviewers for their valuable comments. The work was supported by National Natural Science Foundation of China (Grant No. 61772462), and the 100 Talents Program of Zhejiang University.

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

References


- [1] Liu, M. Y.; Breuel, T.; Kautz, J. Unsupervised image-to-image translation networks. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, 700–708, 2017.
- [2] Zhu, J.-Y.; Park, T.; Isola, P.; Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, 2242–2251, 2017.
- [3] Benaim, S.; Wolf, L. One-sided unsupervised domain mapping. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, 752–762, 2017.
- [4] Tomei, M.; Cornia, M.; Baraldi, L.; Cucchiara, R. Art2Real: Unfolding the reality of artworks via semantically-aware image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 5842–5852, 2019.
- [5] Almahairi, A.; Rajeswar, S.; Sordani, A.; Bachman, P.; Courville, A. C. Augmented CycleGAN: Learning many-to-many mappings from unpaired data. In: Proceedings of the 35th International Conference on Machine Learning, 195–204, 2018.
- [6] Härkönen, E.; Hertzmann, A.; Lehtinen, J.; Paris, S. GANSpace: Discovering interpretable GAN controls. In: Proceedings of the 34th International Conference on Neural Information Processing Systems, Article No. 825, 9841–9850, 2020.
- [7] Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. *Communications of the ACM* Vol. 63, No. 11, 139–144, 2020.
- [8] Arjovsky, M.; Chintala, S.; Bottou, L. Wasserstein generative adversarial networks. In: Proceedings of the 34th International Conference on Machine Learning, Vol. 70, 214–223, 2017.
- [9] Karras, T.; Laine, S.; Aila, T. M. A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 4396–4405, 2020.
- [10] Isola, P.; Zhu, J. Y.; Zhou, T. H.; Efros, A. A. Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 5967–5976, 2017.
- [11] Iizuka, S.; Simo-Serra, E.; Ishikawa, H. Globally and locally consistent image completion. *ACM Transactions on Graphics* Vol. 36, No. 4, Article No. 107, 2017.
- [12] Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A. P.; Tejani, A.; Totz, J.; Wang, Z.; et al. Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 105–114, 2017.
- [13] Mirza, M.; Osindero, S. Conditional generative adversarial nets. *arXiv preprint* arXiv:1411.1784, 2014.
- [14] Kim, T.; Cha, M.; Kim, H.; Lee, J. K.; Kim, J. Learning to discover cross-domain relations with generative adversarial networks. In: Proceedings of the 34th International Conference on Machine Learning, Vol. 70, 1857–1865, 2017.
- [15] Yi, Z.; Zhang, H.; Tan, P.; Gong, M. DualGAN: Unsupervised dual learning for image-to-image translation. In: Proceedings of the IEEE International Conference on Computer Vision, 2868–2876, 2017.
- [16] Chen, S. Y.; Su, W. C.; Gao, L.; Xia, S. H.; Fu, H. B. DeepFaceDrawing: Deep generation of face images from sketches. *ACM Transactions on Graphics* Vol. 39, No. 4, Article No. 72, 2020.
- [17] Chen, S. Y.; Liu, F. L.; Lai, Y. K.; Rosin, P. L.; Li, C. P.; Fu, H. B.; Gao, L. DeepFaceEditing: Deep face generation and editing with disentangled geometry and appearance control. *ACM Transactions on Graphics* Vol. 40, No. 4, Article No. 90, 2021.
- [18] Lee, C. H.; Liu, Z. W.; Wu, L. Y.; Luo, P. MaskGAN: Towards diverse and interactive

- facial image manipulation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 5548–5557, 2020.
- [19] Mao, X. D.; Li, Q.; Xie, H. R.; Lau, R. Y. K.; Wang, Z.; Smolley, S. P. Least squares generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, 2813–2821, 2017.
- [20] Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; Catanzaro, B. High-resolution image synthesis and semantic manipulation with conditional GANs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8798–8807, 2018.
- [21] Lu, G. S.; Zhou, Z. M.; Song, Y. X.; Ren, K.; Yu, Y. Guiding the one-to-one mapping in CycleGAN via optimal transport. *Proceedings of the AAAI Conference on Artificial Intelligence* Vol. 33, No. 1, 4432–4439, 2019.
- [22] Lu, G.; Zhou, Z.; Shen, J.; Chen, C.; Zhang, W.; Yu, Y. Large-scale optimal transport via adversarial training with cycle-consistency. *arXiv preprint arXiv:2003.06635*, 2020.
- [23] Huang, X.; Liu, M. Y.; Belongie, S.; Kautz, J. Multimodal unsupervised image-to-image translation. In: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science, Vol. 11207*. Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y. Eds. Springer Cham, 179–196, 2018.
- [24] Lee, H. Y.; Tseng, H. Y.; Huang, J. B.; Singh, M.; Yang, M. H. Diverse image-to-image translation via disentangled representations. In: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science, Vol. 11205*. Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y. Eds. Springer Cham, 36–52, 2018.
- [25] Alami Mejjati, Y.; Richardt, C.; Tompkin, J.; Cosker, D.; Kim, K. I. Unsupervised attention-guided image-to-image translation. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems, 3697–3707, 2018.
- [26] Kim, J.; Kim, M.; Kang, H.; Lee, K. U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*, 2019.
- [27] Choi, Y.; Choi, M.; Kim, M.; Ha, J. W.; Kim, S.; Choo, J. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8789–8797, 2018.
- [28] Lee, H. Y.; Tseng, H. Y.; Mao, Q.; Huang, J. B.; Lu, Y. D.; Singh, M.; Yang, M. DRIT++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision* Vol. 128, Nos. 10–11, 2402–2417, 2020.
- [29] Choi, Y.; Uh, Y.; Yoo, J.; Ha, J. W. StarGAN v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 8185–8194, 2020.
- [30] Liu, M. Y.; Huang, X.; Mallya, A.; Karras, T.; Aila, T. M.; Lehtinen, J.; Kautz, J. Few-shot unsupervised image-to-image translation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 10550–10559, 2019.
- [31] Fu, H.; Gong, M. M.; Wang, C. H.; Batmanghelich, K.; Zhang, K.; Tao, D. C. Geometry-consistent generative adversarial networks for one-sided unsupervised domain mapping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2422–2431, 2019.
- [32] Kingma, D. P.; Welling, M. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [33] Kingma, D. P.; Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [34] Yu, A.; Grauman, K. Fine-grained visual comparisons with local learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 192–199, 2014.
- [35] Zhu, J. Y.; Krähenbühl, P.; Shechtman, E.; Efros, A. A. Generative visual manipulation on the natural image manifold. In: *Computer Vision – ECCV 2016. Lecture Notes in Computer Science, Vol. 9909*. Leibe, B.; Matas, J.; Sebe, N.; Welling, M. Eds. Springer Cham, 597–613, 2016.
- [36] LeCun, Y.; Cortes, C.; Burges, C. J. C. THE MNIST DATABASE of handwritten digits. 1998. Available at <http://yann.lecun.com/exdb/mnist/>.
- [37] Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, 6629–6640, 2017.
- [38] Ulyanov, D.; Vedaldi, A.; Lempitsky, V. S. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [39] Ioffe, S.; Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: Proceedings of the 32nd International Conference on International Conference on Machine Learning, Vol. 37, 448–456, 2015.

- [40] Miyato, T.; Kataoka, T.; Koyama, M.; Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint* arXiv:1802.05957, 2018.



Zhexi Peng received his B.S. degree in information and computing sciences from Beijing University of Posts and Telecommunications. He is currently a Ph.D. student at Zhejiang University. His research interests cover computer vision, SLAM, and machine learning.



He Wang is an associate professor at the School of Computing, University of Leeds, UK. He is the director of High-Performance Graphics and Game Engineering and the Academic Lead of the Centre for Immersive Technology. His current research interests are mainly in computer graphics, vision, and machine learning and applications.




Yanlin Weng is an associate professor in the School of Computer Science and Technology at Zhejiang University. She got her Ph.D. degree in computer science from the University of Wisconsin-Milwaukee, and her master and bachelor degrees in control science and engineering from Zhejiang University. Her research interests include computer graphics and multimedia.



Yin Yang is an associate professor in the School of Computing, Clemson University. Previously, he was a faculty member of the University of New Mexico. He received his Ph.D. degree from the University of Texas at Dallas (with a David Daniel fellowship). He is a recipient of an NSF CRII award (2015)

and a CAREER award (2019). His research aims to develop efficient and customized computing methods for challenging problems in graphics, simulation, machine learning, vision, visualization, robotics, medicine, and many other applied areas.



Tianjia Shao is a ZJU100 Young Professor in the State Key Laboratory of CAD&CG, Zhejiang University. Previously He was a lecturer in the School of Computing, University of Leeds. He received his Ph.D. degree in computer science from the Institute for Advanced Study, Tsinghua University, and his B.S. degree from the Department of Automation, Tsinghua University. During his Ph.D., he spent one year as a visiting researcher in University College London and four years as a research intern in the Internet Graphics Group, Microsoft Research Asia before becoming an assistant researcher in the State Key Laboratory of CAD&CG, Zhejiang University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.