

Mix and match networks: encoder-decoder alignment for zero-pair image translation

Yaxing Wang, Joost van de Weijer, Luis Herranz
Computer Vision Center, Universitat Autònoma de Barcelona
Barcelona, Spain
`{wang, joost, lherranz}@cvc.uab.es`

Abstract

We address the problem of image translation between domains or modalities for which no direct paired data is available (i.e. zero-pair translation). We propose mix and match networks, based on multiple encoders and decoders aligned in such a way that other encoder-decoder pairs can be composed at test time to perform unseen image translation tasks between domains or modalities for which explicit paired samples were not seen during training. We study the impact of autoencoders, side information and losses in improving the alignment and transferability of trained pairwise translation models to unseen translations. We show our approach is scalable and can perform colorization and style transfer between unseen combinations of domains. We evaluate our system in a challenging cross-modal setting where semantic segmentation is estimated from depth images, without explicit access to any depth-semantic segmentation training pairs. Our model outperforms baselines based on pix2pix and CycleGAN models.

1. Introduction

Image-to-image translations (or simply image translations) are an integral part of many computer vision systems. They include transformations between different modalities, such as from RGB to depth [19], or domains, such as luminance to color images [32], horses to zebras [34], or editing operations such as artistic style changes [9]. These mappings can also include 2D label representations such as semantic segmentations [21] or surface normals [5]. Deep networks have shown excellent results in learning models to perform image translations between different domains and modalities [1, 10, 21]. These systems are typically trained with pairs of matching images between domains, e.g. an RGB image and its corresponding depth image.

Image translation methods, which transfer images from one domain to another, are often based on encoder-decoder

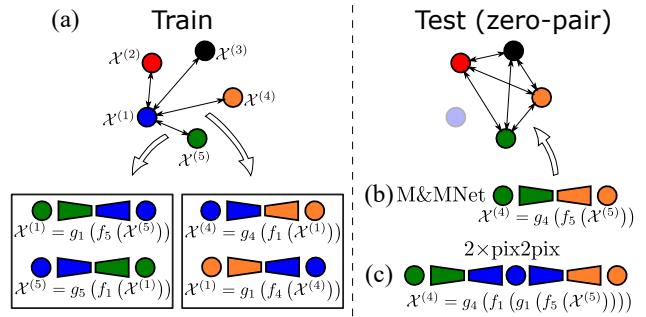


Figure 1: *Zero-pair image translation*: (a) given a set of domains or modalities (circles) for which paired training data is available, the objective is to evaluate zero-pair translations. Translations are implemented as aligned encoder-decoder networks. (b) *Mix and match networks* do not require retraining on unseen transformations, in contrast to unpaired translation alternatives (e.g. CycleGAN [34]). (c) Two cascaded paired translations (e.g. 2×pix2pix [10]) require explicit translation to an intermediate domain. Better seen in color.

frameworks [1, 10, 21, 34]. In these approaches, an encoder network maps the input image from domain A to a continuous vector in a latent space. From this latent representation the decoder generates an image in domain B. The latent representation is typically much smaller than the original image size, thereby forcing the network to learn to efficiently compress the information from domain A which is relevant for domain B into the latent representation. Encoder-decoder networks are trained end-to-end by providing the network with matching pairs from both domains or modalities. An example could be learning a mapping from RGB to depth [19]. Other applications include semantic segmentation [1] and image restoration [23].

In this paper we introduce *zero-pair image translation*: a new setting for testing image translations which involves evaluating on *unseen* translations, i.e. no matching image

or dataset pairs are available during training (see Figure 1a). Note that this setting is different from unpaired image translation [34, 14, 20], which is evaluated on the same paired domains seen during training.

We also propose *mix and match networks*, an approach that addresses zero-pair image translation by seeking alignment between encoders and decoders via their latent spaces. An unseen translation between two domains is performed by simply concatenating the input domain encoder and the output domain decoder (see Figure 1b). We study several techniques that can improve this alignment, including the usage of autoencoders, latent space consistency losses and the usage of pooling indices as side information to guide the reconstruction of spatial structure. We evaluate this approach in a challenging cross-modal task, where we perform zero-pair depth to semantic segmentation translation, using only RGB to depth and RGB to semantic segmentation pairs during training.

Finally, we show that aligned encoder-decoder networks also have advantages in domains with unpaired data. In this case, we show that mix and match networks scale better with the number of domains, since they are not required to learn all pairwise image translation networks (i.e. scales linearly instead of quadratically). The code is available at <http://github.com/yaxingwang/Mix-and-match-networks>

2. Related Work

Image translation Recently, generic encoder-decoder architectures have achieved impressive results in a wide range of transformations between images. Isola et al. [10] trained from pairs of input and output images to learn a variety of image translations (e.g. color, style), using an adversarial loss. These models require paired training data to be available (i.e. *paired* image translation). Various works extended this idea to the case where no explicit input-output image pairs are available (*unpaired* image translation), using the idea of cyclic consistency [34, 14]. Liu et al. [20] show that unsupervised mappings can be learned by imposing a joint latent space between the encoder and the decoder. In this work we consider the case where paired data is available between some domains or modalities and not available between others (i.e. zero-pair), and how this knowledge can be transferred to those zero-pair cases. In concurrent work, Choi et al. [4] also address scaling to multiple domains (always in the RGB modality) by using a single encoder-decoder model. In contrast, our approach uses multiple cross-aligned encoders and decoders. Our cross-modal setting is also requires deeper structural changes and modality-specific encoder-decoders.

Multimodal encoder-decoders Encoder-decoder networks can be extended into multi-way encoder-decoder networks by adding encoders and/or decoders for multiple

domains together. Recently, joint encoder-decoder architectures have been used in multi-task settings, where the network is trained to perform multiple tasks (e.g. depth estimation, semantic segmentation, surface normals) [5, 13], and multimodal settings, where the inputs data can be from different modalities or even combine several ones [25].

Training a multimodal encoder-decoder network was recently studied in [16]. They use a joint latent representation space for the various modalities. In our work we consider the alignment and transferability of pairwise image translations to unseen translations, rather than joint encoder-decoder architectures. Another multimodal encoder-decoder network was studied in [2]. They show that multimodal autoencoders can address the depth estimation and semantic segmentation tasks simultaneously, even in the absence of some of the input modalities. All these works do not consider the zero-pair image translation problem addressed in this paper.

Zero-shot recognition In conventional supervised image recognition, the objective is to predict the class label that is provided during training [18, 8]. However, this poses limitations in scalability to new classes, since new training data and annotations are required. In zero-shot learning, the objective is to predict an unknown class for which there is no image available, but a description of the class (i.e. *class prototype*). This description can be a set of attributes (e.g. has wings, blue, four legs, indoor) [18, 11], concept ontologies [6, 27] or textual descriptions [26]. In general, an intermediate semantic space is leveraged as a bridge between the visual features from seen classes and class description from unseen ones. In contrast to zero-shot recognition, we focus on unseen translations (unseen input-output pairs rather than simply unseen class labels).

Zero-pair language translation Evaluating models on unseen language pairs has been studied recently in machine translation [12, 3, 33, 7]. Johnson et al. [12] proposed a neural language model that can translate between multiple languages, even pairs of language where no explicit paired sentences were provided¹. In their method, the encoder, decoder and attention are shared. In our method we focus on images, which are essentially a radically different type of data, with two dimensional structure in contrast to the sequential structure of language.

¹Note that [12] refers to this as *zero-shot* translation. In this paper we refer to this setting as zero-pair to emphasize that what is unseen is paired data and avoid ambiguities with traditional zero-shot recognition which typically refers to unseen samples.

3. Encoder-decoder alignment

3.1. Multi-domain image translation

We consider the problem of image translation from domain $\mathcal{X}^{(i)}$ to domain $\mathcal{X}^{(j)}$ as $T_{ij}: x^{(i)} \mapsto x^{(j)}$. In our case it is implemented as a encoder-decoder chain $x^{(j)} = T_{ij}(x^{(i)}) = g_j(f_i(x^{(i)}))$ with encoder f_i and decoder g_j (see Figure 1). The domains connected during training are all trained jointly, and in both directions. It is important to note that for each domain one encoder and one decoder are trained. By training all these encoders and decoders jointly the latent representation is encouraged to align. As a consequence of the alignment of the latent space we can do zero-pair translation at testing time between the domains for which no training pairs were available. The main aim of this article is to analyze to what extend this alignment allows for zero-pair image translation.

3.2. Aligning for zero-pair translation

Zero-pair translation in images is especially challenging due to the inherent complexity of images, especially in multimodal settings. Ideally, a good latent representation that also works in unseen translations should be not only well-aligned but also unbiased to any particular domain. In addition, the encoder-decoder system should be able to preserve the spatial structure, even in unseen transformations.

Autoencoders One way to improve alignment is by jointly training domain-specific autoencoders with the image translation networks. By sharing the weights between the auto-encoders and the image translation encoder-decoders the latent space is forced to align.

Latent space consistency The latent space can be enforced to be invariant across multiple domains. Taigman *et al.* [30] use L2 distance between a latent representation and the reconstructed after another decoding and encoding cycle. When paired samples $(x^{(i)}, x^{(j)})$ are available, we propose using cross-domain latent space consistency in order to enforce $f_i(x^{(i)})$ and $f_j(x^{(j)})$ to be aligned.

Preserving spatial structure using side information In general, image translation tasks result in output images with similar spatial structure as the input ones, such as scene layouts, shapes and contours that are preserved across modalities. In fact, this spatial structure available in the input image is critical to simplify the problem and achieve good results, and successful image translation methods often use multi-scale intermediate representations from the encoder as side information to guide the decoder in the upsampling process. Skip connections are widely used for this purpose. However, conditional decoders using skip connections expect specific information from a particular domain-specific encoder that would be unlikely to work in unseen encoder-decoder pairs. Motivated by efficiency, pooling indices [1]

were recently proposed as a compact descriptor to guide the decoder in lightweight models. We show here that pooling indices constitute robust and relatively encoder-independent side information, suitable for improving decoding even in unseen translations.

Adding noise to latent space We found that adding some noise at the output of each encoder also helps to train the network and improves the results during test. This seems to help in obtaining more invariance in the common latent representation and better alignment across modalities.

3.3. Scalable image translation

One of the advantages of our mix and match networks is that the system can infer many pairwise domain-to-domain translations when the number of domains is high, without explicitly training them. Other pairwise methods where encoders and decoders are not cross-aligned, such as CycleGAN[34], would require training $N \times (N - 1)/2$ encoders and decoders for N domains. For mix and match networks each encoder and decoder should be involved in at least one translation pair during training in order to be aligned with the others, thereby reducing complexity from quadratic to linear with the number of domains (i.e. $N - 1$ encoders/decoders).

4. Zero-pair cross-modal image translation

We propose a challenging cross-modal setting to evaluate zero-pair image translation involving three modalities²: RGB, depth and semantic segmentation. It is important to observe that a setting involving heterogeneous modalities³ (in terms of complexity, number and meaning of different channels, etc.) is likely to require modality-specific architectures and losses.

4.1. Problem definition

We consider the problem of jointly learning RGB-to-segmentation translation $y = T_{RS}(x)$ with and RGB-to-depth translation $z = T_{RD}(x)$ and evaluating on an unseen transformation $y = T_{DS}(z)$. The first translation is learned from a semantic segmentation dataset $\mathcal{D}^{(1)}$ with pairs $(x, y) \in \mathcal{D}^{(1)}$ of RGB images and segmentation maps, and the second from a disjoint RGB-D dataset $\mathcal{D}^{(2)}$ with pairs of RGB and depth images $(x, z) \in \mathcal{D}^{(2)}$. Therefore no depth image and segmentation map pairs are available to the system. However, note that the RGB images from both datasets could be combined if necessary (we denote this dataset as $\mathcal{X} = \{x | x \in \mathcal{D}^{(1)} \cup \mathcal{D}^{(2)}\}$). The system is evaluated on a third dataset $\mathcal{D}^{(3)}$ with paired depth images and segmentation maps.

²Here the term *modality* has the same role of *domain* in the previous section.

³For simplicity, we will refer to semantic segmentation maps and depth as modalities rather than tasks

4.2. Mix and match networks architecture

The overview of the framework is shown in Fig. 2. As basic building blocks we use three modality-specific encoders $f_R(x)$, $f_D(x)$ and $f_S(x)$ (RGB, depth and semantic segmentation, respectively), and the corresponding three modality-specific decoders $g_R(h)$, $g_D(h)$ and $g_S(h)$, where h is the latent representation in the shared space. The required translations are implemented as $y = T_{RS}(x) = g_S(f_R(x))$, $z = T_{RD}(x) = g_D(f_R(x))$ and $y = T_{DS}(z) = g_S(f_D(z))$.

Encoder and decoder weights are shared across the different translations involving same modalities (same color in Fig. 2). To enforce better alignment between encoders and decoders of the same modality, we also include self-translations using the corresponding three autoencoders $T_{RR}(x) = g_R(f_R(x))$, $T_{DD}(y) = g_D(f_D(x))$ and $T_{SS}(x) = g_S(f_S(x))$.

We based our encoders and decoders on the SegNet architecture [1]. The encoder of SegNet itself is based on the 13 convolutional layers VGG-16 architecture [29]. The decoder mirrors the encoder architecture with 13 deconvolutional layers. All encoders and decoders are randomly initialized except for the RGB encoder which is pretrained on ImageNet.

As in SegNet, pooling indices at each downsampling layer of the encoder are provided to the corresponding upsampling layer of the (seen or unseen) decoder⁴. These pooling indices seem to be relatively similar across the three modalities and effective to transfer spatial structure information that help to obtain better depth and segmentation boundaries in higher resolutions. Thus, they provide relatively modality-independent side information.

4.3. Loss functions

As we saw before, a correct cross-alignment between encoders and decoders that have not been connected during training is critical for zero-pair translation. The final loss combines a number of modality-specific losses for both cross-domain translation and self-translation (i.e. autoencoders) and alignment constraints in the latent space

$$L = \lambda_R L_{RGB} + \lambda_S L_{SEG} + \lambda_D L_{DEPTH} + \lambda_A L_{LAT}$$

RGB-specific We use a combination of L2 distance and adversarial loss $L_{RGB} = \lambda_{L2} L_{L2} + L_{GAN}$. L2 distance is used to compare the estimated and the ground truth RGB image after translation from a corresponding depth or seg-

⁴The RGB decoder does not use pooling indices, since in our experiments we observed undesired grid-like artifacts in the RGB output.

mentation image. It is also used in the RGB autoencoder

$$\begin{aligned} L_{L2} &= \mathbb{E}_{(x,y) \sim p_{\mathcal{D}(1)}(x,y)} [\|T_{SR}(y) - x\|_2] \\ &+ \mathbb{E}_{(x,z) \sim p_{\mathcal{D}(2)}(x,z)} [\|T_{DR}(z) - x\|_2] \\ &+ \mathbb{E}_{x \sim p_{\mathcal{X}}(x)} [\|T_{RR}(x) - x\|_2] \end{aligned}$$

In addition, we also include the least squares adversarial loss [22, 10] on the output of the RGB decoder

$$L_{GAN} = \mathbb{E}_{x \sim p_{\mathcal{X}}(x)} [(C(x) - 1)^2] + \mathbb{E}_{\hat{x} \sim \hat{p}(x)} [(C(\hat{x}))^2]$$

where $\hat{p}(x)$ is the resulting distribution of the combined images \hat{x} generated by $\hat{x} = T_{SR}(y)$, $\hat{x} = T_{DR}(z)$ and $\hat{x} = T_{RR}(x)$. Note that the RGB autoencoder and the discriminator $C(x)$ are both trained with the combined RGB data \mathcal{X} .

Depth For depth we use the Berhu loss [17] in both RGB-to-depth translation and in the depth autoencoder

$$\begin{aligned} L_{DEPTH} &= \mathbb{E}_{(x,z) \sim p_{\mathcal{D}(2)}(x,z)} [\mathcal{B}(T_{RD}(x) - z)] \\ &+ \mathbb{E}_{(x,z) \sim p_{\mathcal{D}(2)}(x,z)} [\mathcal{B}(T_{DD}(z) - z)] \end{aligned}$$

where $\mathcal{B}(z)$ is the average Berhu loss.

Semantic segmentation For segmentation we use the average cross-entropy loss $\mathcal{CE}(\hat{y}, y)$ in both RGB-to-segmentation translation and the segmentation autoencoder

$$\begin{aligned} L_{SEM} &= \mathbb{E}_{(x,y) \sim p_{\mathcal{D}(1)}(x,y)} [\mathcal{CE}(T_{RS}(x), y)] \\ &+ \mathbb{E}_{(x,y) \sim p_{\mathcal{D}(2)}(x,y)} [\mathcal{CE}(T_{SS}(y), y)]. \end{aligned}$$

Latent space consistency We enforce latent representations to remain close independently of the encoder that generated them. In our case we have two consistency losses

$$\begin{aligned} L_{LAT} &= L_{LAT1} + L_{LAT2} \\ L_{LAT1} &= \mathbb{E}_{(x,y) \sim p_{\mathcal{D}(1)}(x,y)} [\|f_R(x) - f_S(y)\|_2] \\ L_{LAT2} &= \mathbb{E}_{(x,z) \sim p_{\mathcal{D}(2)}(x,z)} [\|f_R(x) - f_D(z)\|_2] \end{aligned}$$

5. Experimental Results

To the best of our knowledge there is no existing work which reports results for the setting of zero-pair image translation. In particular, we evaluate the proposed mix and match networks on zero-pair translation for semantic segmentation from depth images (and viceversa), and we show results for semantic segmentation from multimodal data. Finally, we illustrate the possibility to perform zero-pair translations for unpaired datasets, and the advantage of mix and match networks in terms of scalability.

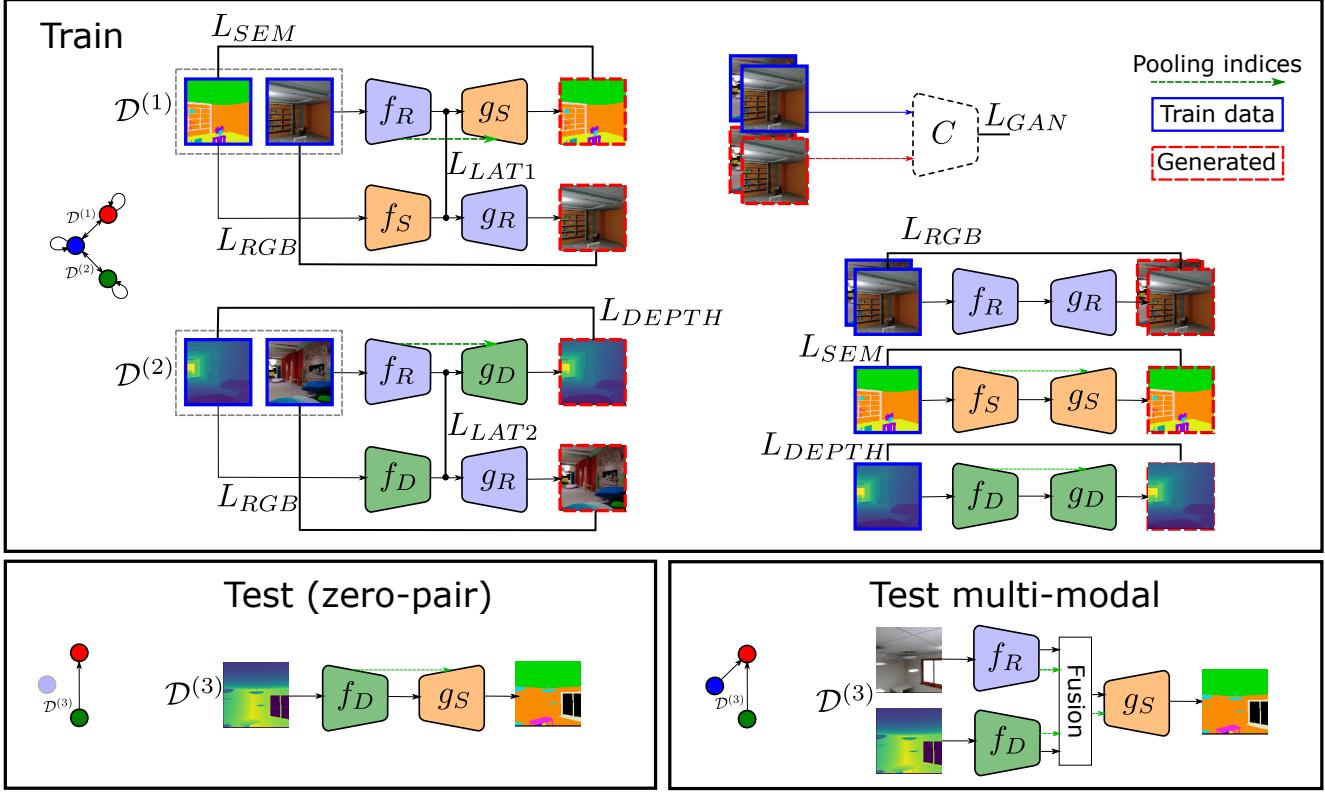


Figure 2: Zero-pair cross-modal and multimodal image translation with mix and match networks. Two disjoint sets $\mathcal{D}^{(1)}$ and $\mathcal{D}^{(2)}$ are seen during training, containing (RGB, depth) pairs and (RGB, segmentation) pairs, respectively. The system is tested on the unseen translation depth-to-segmentation (zero-pair) and (RGB+depth)-to-segmentation (multimodal), using a third unseen set $\mathcal{D}^{(3)}$. Encoders and decoders with the same color share weights. Better viewed in color.

5.1. Datasets and experimental settings

SceneNetRGBD The SceneNetRGBD dataset [24] consists of 16865 synthesized train videos and 1000 test videos. Each of them includes 300 matching triplets (RGB, depth and segmentation map), with a size of 320x240 pixels. In our examples, we use two subsets as our datasets:

- 51K dataset: the train set is selected from the first 50 frames from each of the first 1000 videos in the train set. The test set is collected by selecting the 60th frame from the same 1000 videos. This dataset was used to evaluate some of the architecture design choices.
- 170K dataset: We collected a larger dataset which consists of 150K triplets for the train set, 10K triplets for the validation set and 10K triplets for the test set. The 10K validation set is also from the train set of SceneNetRGBD. For the train set, we select 10 triplets from the first 150000 training triplets. The triplets are sampled from the first frame to last frame every 30 frames. The validation set is from the remaining videos of the train set and the test set is taken from the test dataset.

Each train set is divided into two equal non-overlapping splits from different videos. Although the collected splits contain triplets, we only use pairs to train our model.

Following common practices in these tasks, for segmentation we compute the intersection over union (IoU) and report per-class average (mIoU), and the global scores, which gives the percentage of correctly classified pixels. For depth we also include quantitative evaluation, following the standard error metrics for depth estimation [5].

Network training We use Adam [15] with batch size of 6, using a learning rate of 0.0002. We set $\lambda_R = 1$, $\lambda_S = 100$, $\lambda_D = 10$, $\lambda_A = 1$, $\lambda_{L2} = 1$. For the first 200,000 iterations we train all networks. For the following 200,000 iterations we use $\lambda_R = 10$, $\lambda_A = 10$, $\lambda_{L2} = 10$ and freeze the RGB encoder. We found the network converges faster using a large initial λ_A . We add Gaussian noise to the latent space with zero mean and a standard deviation of 0.5.

5.2. Ablation study

In a first experiment we use the 51K dataset to study the impact of several design elements on the overall per-

Side information	Pretrained	mIoU	Global
-	N	32.2%	63.5%
Skip connections	N	14.1%	52.6%
Pooling indices	N	45.6%	73.4%
Pooling indices	Y	49.5%	80.0%

Table 1: Influence of side information and RGB encoder pretraining on the final results. The task is zero-shot depth-to-segmentation.

AutoEnc	Latent	Noise	mIoU	Global
N	N	N	5.64%	13.5%
Y	N	N	22.9%	52.6%
Y	Y	N	48.9%	78.2%
Y	Y	Y	49.5%	80.0%

Table 2: Impact of several components (autoencoder, latent space consistency loss and noise) in the performance. The task is zero-shot depth-to-segmentation.

formance of the system.

Side information We first evaluate the usage of side information from the encoder to guide the upsampling process in the decoder. We consider three variants: no side information, skip connections [28] and pooling indices [1]. The results in Table 1 show that skip connections obtain worse results than no side information at all. This is due to the fact that skip connections are not domain-invariant and at testing time when we combine an encoder and decoder these connections result in a different input from the one seen under training, resulting in a drop of performance. Fig. 3 illustrates the differences between these three variants. Without side information the network is able to reconstruct a coarse segmentation but without further guidance it is not able to refine it properly. Skip connections provide features that could guide the decoder but instead confuse it, since in the zero-pair case the decoder has not seen the features of that particular encoder. Pooling indices are more invariant as side information and obtaining the best results.

RGB pretraining We also compare training the RGB encoder from scratch and initializing with pretrained weights from ImageNet. Table 1 show an additional gain of around 5% in mIoU when using the pretrained weights.

Given these results we perform all the remaining experiments initializing the RGB encoder with pretrained weights and use pooling indices as side information.

Latent space consistency, noise and autoencoders We evaluate these three factors, with Table 2 showing that latent space consistency and the usage of autoencoders lead to significant performance gains; for both, the performance (in mIoU) is more than doubled. Adding noise to the output of the encoder results in a small performance gain.

5.3. Comparison with other methods

We compare the results of our mix and match networks for depth to segmentation, $D \rightarrow S$, to the following two baselines:

- CycleGAN [34] learns a mapping from depth to semantic segmentation without explicit pairs. In contrast to ours, this method only leverages depth and semantic segmentation, ignoring the available RGB data and the corresponding pairs.
- $2 \times \text{pix2pix}$ [10] learns from paired data two encoder-decoder pairs ($D \rightarrow R$ and $R \rightarrow S$). The architecture uses skip connections and the corresponding modality-specific losses. We use the exact code from [10]. In contrast to ours, it requires explicit decoding to RGB, which may degrade the quality of the prediction.
- $D \rightarrow R \rightarrow S$ is similar as the $2 \times \text{pix2pix}$ but than with a similar architecture as we use in our M&MNet. We train a translation model from depth to RGB and from RGB to segmentation, and obtain the transformation depth-to-segmentation by concatenating them. Note that it requires using an intermediate RGB image.

Table 3 compares the three methods on the 170K dataset. CycleGAN is not able to learn a good mapping from depth to semantic segmentation, showing the difficulty of unpaired translation to solve this complex cross-modal task. $2 \times \text{pix2pix}$ manages to improve the results by resorting to the anchor domain RGB, although still not satisfactory since the first translation network drops information not relevant for the RGB task but necessary for reconstructing depth (like in the Chinese whispers/telephone game).

Mix and match networks evaluated on $(D \rightarrow R \rightarrow S)$ achieve a similar result to CycleGAN, but significantly worse than $2 \times \text{pix2pix}$. However, when we run our architecture with skip connections we obtain similar results as $2 \times \text{pix2pix}$. Note that because in this setting the encoders and decoders are used in the same setting in both training and testing, skip connections function well.

The direct combination $(D \rightarrow S)$ outperforms all baselines significantly. The results more than double in terms of mIoU. Figure 4 illustrates the comparison between our approach and the baselines; our method is the only one that manages to identify the main semantic classes and their general contours in the image. In conclusion, the results show that mix and match networks enable effective zero-pair translation.

In Table 4 we show the results when we test in the opposite direction from semantic segmentation to depth. The conclusions are similar as in previous experiment. Again our method $S \rightarrow D$ outperforms both baseline methods on

Method	Conn.	L_{SEM}	Bed	Book	Ceiling	Chair	Floor	Furniture	Object	Picture	Sofa	Table	TV	Wall	Window	mIoU	Global
Baselines																	
CycleGAN [34]	SC	CE	2.79	0.00	16.9	6.81	4.48	0.92	7.43	0.57	9.48	0.92	0.31	17.4	15.1	6.34	14.2
2×pix2pix [10]	SC	CE	34.6	1.88	70.9	20.9	63.6	17.6	14.1	0.03	38.4	10.0	4.33	67.7	20.5	25.4	57.6
M&MNet $D \rightarrow R \rightarrow S$	PI	CE	0.02	0.00	8.76	0.10	2.91	2.06	1.65	0.19	0.02	0.28	0.02	58.2	3.3	5.96	32.3
M&MNet $D \rightarrow R \rightarrow S$	SC	CE	25.4	0.26	82.7	0.44	56.6	6.30	23.6	5.42	0.54	21.9	10.0	68.6	19.6	24.7	59.7
Zero-pair																	
M&MNet $D \rightarrow S$	PI	CE	50.8	18.9	89.8	31.6	88.7	48.3	44.9	62.1	17.8	49.9	51.9	86.2	79.2	55.4	80.4
Multi-modal																	
M&MNet $(R, D) \rightarrow S$	PI	CE	49.9	25.5	88.2	31.8	86.8	56.0	45.4	70.5	17.4	46.2	57.3	87.9	79.8	57.1	81.2

Table 3: Zero-pair depth-to-semantic segmentation. **SC**: skip connections, **PI**: pooling indexes, **CE**: cross-entropy

Method	$\delta <$			RMSE	RMSE
	1.25	1.25 ²	1.25 ³	(lin)	(log)
Baselines					
CycleGAN [34]	0.05	0.12	0.20	4.63	1.98
M&MNet $S \rightarrow R \rightarrow D$	0.15	0.30	0.44	3.24	1.24
Zero-pair					
M&MNet $S \rightarrow D$	0.33	0.42	0.59	2.8	0.67
Multi-modal					
M&MNet $(R, S) \rightarrow D$	0.36	0.48	0.65	2.48	0.64

Table 4: Zero-pair semantic segmentation-to-depth.

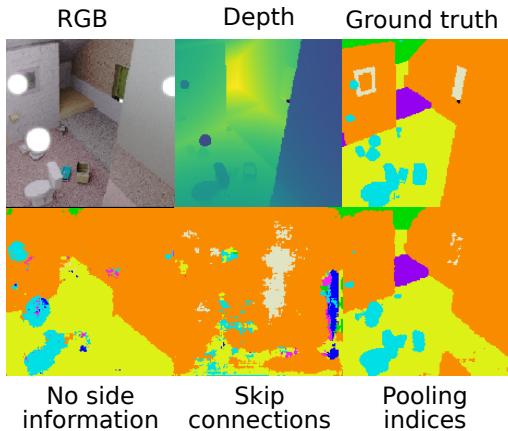


Figure 3: Role of side information in unseen depth-to-segmentation translation.

all five evaluation metrics. Fig. 5 illustrates this case, showing how pooling indices are also key to obtain good depth images, compared with no side information at all.

5.4. Multimodal translation

Next we consider the case of multimodal translation from pairs (RGB, depth) to semantic segmentation. As depicted in Figure 2 multiple modalities can be combined

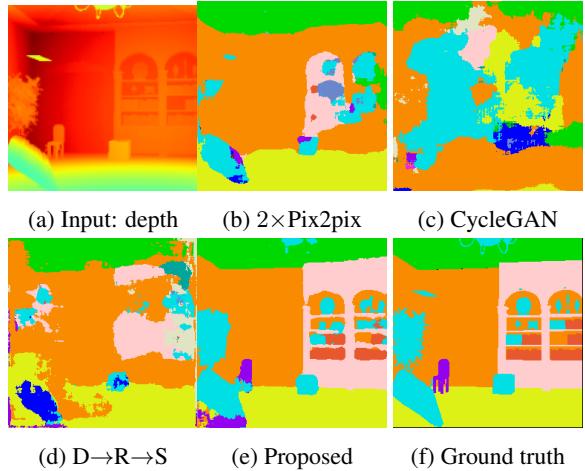


Figure 4: Different methods evaluated on zero-pair depth-to-segmentation.

(since the latent spaces are aligned) at the input of semantic segmentation decoder. To combine the two modalities we perform a weighted average of both RGB and depth latent vectors (the weight α ranges from $\alpha = 0$, only RGB, and $\alpha = 1$, only depth, depending on the case). We set α to 0.2 and use the pooling indices from the RGB encoder (instead of those from the depth encoder, see supplementary material for more details). The results in Table 3 and Table 4 and the example in Figure 5 show that this multimodal combination further improves the performance of zero-pair translation.

5.5. Scalable unpaired image translation

As explained in Section 3.3, mix and match networks scale linearly with the number of domains, whereas existing unpaired image translation methods scale quadratically. As examples of translations between many domains, we show results for object recoloring and style transfer, using mix and match networks based on multiple Cycle-

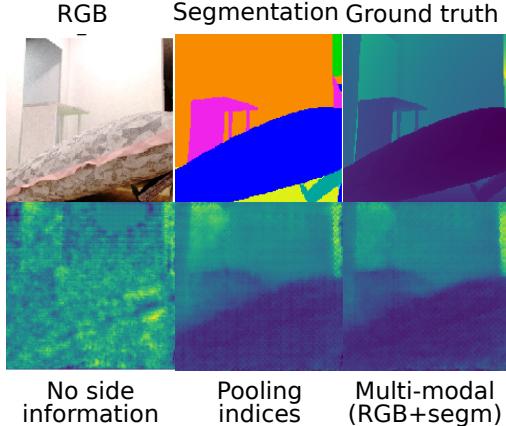


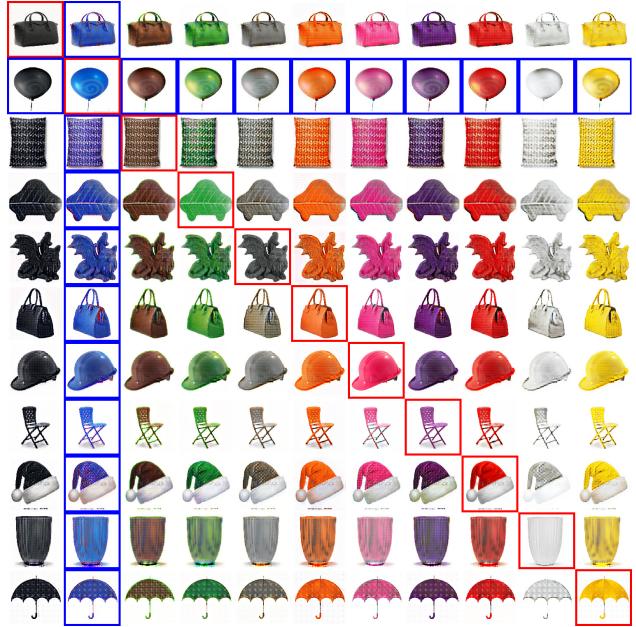
Figure 5: Zero-pair and multimodal segmentation-to-depth.

GANs [34] combined with autoencoders. For the former we use the colored objects dataset [31] with eleven distinct colors ($N = 11$) and around 1000 images per color. Covering all possible image-to-image recoloring combinations requires training a total of $N(N - 1)/2 = 55$ encoders (and decoders) using CycleGANs. In contrast, mix and match networks only require to train $N - 1 = 10$ encoders and eleven decoders, while still successfully addressing the recoloring task (see Figure 6a). Similarly, scalable style transfer can be addressed using mix and match networks (see Figure 6b for an example).

6. Conclusion

In this paper we introduce the problem of zero-pair image translation, where knowledge learned in paired translation models can be effectively transferred and leveraged to perform new unseen translations. The image-to-image scenario poses several challenges to the alignment of encoders and decoders in a way that guarantees cross-domain transferability and without too much dependence on the domain or the modality. We studied this scenario in zero-pair cross-modal and multimodal settings. Notably, we found that side information in the form of pooling indices is robust to modality changes and very helpful to guide the reconstruction of spatial structure. Other helpful techniques are cross-modal consistency losses and adding noise to the latent representation.

Acknowledgements Herranz acknowledges the European Unions H2020 research under Marie Skłodowska-Curie grant No. 6655919. We acknowledge the project TIN2016-79717-R, the CHISTERA project M2CR (PCIN-2015-251) of the Spanish Government and the CERCA Programme of Generalitat de Catalunya. Yaxing Wang acknowledges the Chinese Scholarship Council (CSC) grant No.201507040048. We also acknowledge the generous GPU support from Nvidia.



(a) Color transfer: only transformations from or to blue (anchor domain) are seen. Input images are highlighted in red and seen translations in blue.



(b) Style transfer: trained on four styles + photo (anchor) with data from [34]). From left to right: photo, Monet, van Gogh, Ukiyo-e and Cezanne. Input images are highlighted in red and seen translations in blue.

Figure 6: Zero-pair cross-domain unpaired translations.

References

- [1] V. Badrinarayanan, A. Handa, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for

robust semantic pixel-wise labelling. *arXiv preprint arXiv:1505.07293*, 2015.

[2] C. Cadena, A. R. Dick, and I. D. Reid. Multi-modal auto-encoders as joint estimators for robotics scene understanding. In *Robotics: Science and Systems*, 2016.

[3] Y. Chen, Y. Liu, Y. Cheng, and V. O. Li. A teacher-student framework for zero-resource neural machine translation. *arXiv preprint arXiv:1705.00753*, 2017.

[4] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[5] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2650–2658, 2015.

[6] R. Fergus, H. Bernal, Y. Weiss, and A. Torralba. Semantic label sharing for learning with many categories. *Computer Vision–ECCV 2010*, pages 762–775, 2010.

[7] O. Firat, K. Cho, and Y. Bengio. Multi-way, multilingual neural machine translation with a shared attention mechanism. *arXiv preprint arXiv:1601.01073*, 2016.

[8] Y. Fu, T. Xiang, Y.-G. Jiang, X. Xue, L. Sigal, and S. Gong. Recent advances in zero-shot recognition. *arXiv preprint arXiv:1710.04837*, 2017.

[9] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.

[10] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[11] D. Jayaraman and K. Grauman. Zero-shot recognition with unreliable attributes. In *Advances in Neural Information Processing Systems*, pages 3464–3472, 2014.

[12] M. Johnson, M. Schuster, Q. V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Viégas, M. Wattenberg, G. Corrado, et al. Google’s multilingual neural machine translation system: enabling zero-shot translation. *arXiv preprint arXiv:1611.04558*, 2016.

[13] A. Kendall, Y. Gal, and R. Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. *arXiv preprint arXiv:1705.07115*, 2017.

[14] T. Kim, M. Cha, H. Kim, J. Lee, and J. Kim. Learning to discover cross-domain relations with generative adversarial networks. *arXiv preprint arXiv:1703.05192*, 2017.

[15] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[16] R. Kuga, A. Kanezaki, M. Samejima, Y. Sugano, and Y. Matsushita. Multi-task learning using multi-modal encoder-decoder networks with shared skip connections. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2017.

[17] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab. Deeper depth prediction with fully convolutional residual networks. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 239–248. IEEE, 2016.

[18] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):453–465, 2014.

[19] F. Liu, C. Shen, G. Lin, and I. Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2024–2039, 2016.

[20] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. *arXiv preprint arXiv:1703.00848*, 2017.

[21] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.

[22] X. Mao, Q. Li, H. Xie, R. Y. Lau, and Z. Wang. Multi-class generative adversarial networks with the l2 loss function. *arXiv preprint arXiv:1611.04076*, 2016.

[23] X. Mao, C. Shen, and Y.-B. Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Advances in Neural Information Processing Systems*, pages 2802–2810, 2016.

[24] J. McCormac, A. Handa, S. Leutenegger, and A. J. Davison. Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? 2017.

[25] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011.

- [26] S. Reed, Z. Akata, H. Lee, and B. Schiele. Learning deep representations of fine-grained visual descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–58, 2016.
- [27] M. Rohrbach, M. Stark, and B. Schiele. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1641–1648. IEEE, 2011.
- [28] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [29] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [30] Y. Taigman, A. Polyak, and L. Wolf. Unsupervised cross-domain image generation. *arXiv preprint arXiv:1611.02200*, 2016.
- [31] L. Yu, L. Zhang, J. van de Weijer, F. S. Khan, Y. Cheng, and C. A. Parraga. Beyond eleven color names for image understanding. *Machine Vision and Applications*, 29(2):361–373, 2018.
- [32] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *European Conference on Computer Vision*, pages 649–666. Springer, 2016.
- [33] H. Zheng, Y. Cheng, and Y. Liu. Maximum expected likelihood estimation for zero-resource neural machine translation. *IJCAI*, 2017.
- [34] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision, ICCV*, 2017.