

Evaluation of Correctness in Unsupervised Many-to-Many Image Translation

Dina Bashkirova
Boston University
dbash@bu.edu

Ben Usman
Boston University
usmn@bu.edu

Kate Saenko
Boston University and MIT-IBM Watson AI Lab
saenko@bu.edu

Abstract

Given an input image from a source domain and a guidance image from a target domain, unsupervised many-to-many image-to-image (UMMI2I) translation methods seek to generate a plausible example from the target domain that preserves domain-invariant information of the input source image and inherits the domain-specific information from the guidance image. For example, when translating female faces to male faces, the generated male face should have the same expression, pose and hair color as the input female image, and the same facial hairstyle and other male-specific attributes as the guidance male image. Current state-of-the-art UMMI2I methods generate visually pleasing images, but, since for most pairs of real datasets we do not know which attributes are domain-specific and which are domain-invariant, the semantic correctness of existing approaches has not been quantitatively evaluated yet. In this paper, we propose a set of benchmarks and metrics for the evaluation of semantic correctness of these methods. We provide an extensive study of existing state-of-the-art UMMI2I translation methods, showing that all methods, to different degrees, fail to infer which attributes are domain-specific and which are domain-invariant from data, and mostly rely on inductive biases hard-coded into their architectures. Our code can be found at https://github.com/dbash/umi2i_correctness.

1. Introduction

Unsupervised image-to-image translation aims to map an input image from a source domain to a target domain so that the output looks like a valid example of the target domain while preserving the semantics of the input image, given only two sets of source and target images and no supervision. First attempts at the unsupervised image-to-image translation [40, 25, 18, 36, 12] assumed one-to-one relation between the domains, meaning that for each image in the source domain there exists exactly one corresponding image in the target domain. Unfortunately, this assumption does not hold for most applications, as in most cases

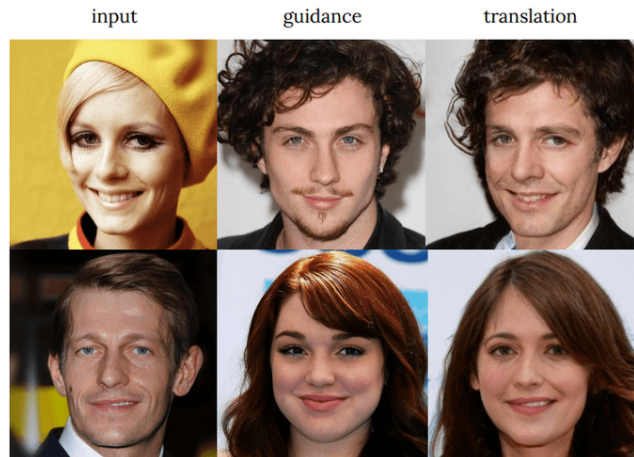


Figure 1. Given an input image from source domain and a guidance image from target domain (e.g. male and female photos), unsupervised many-to-many image translation (UMMI2I) methods aim to generate a new image from the target domain with factors varied in both domains (such as pose, facial expression, hair and skin color, etc.) taken from the input image, while the factors that varied only in the target domain (such as facial hairstyle) taken from the guidance image. We show that current SotA UMMI2I methods fail to infer which attributes are domain-specific and which are domain-invariant from data. For example, StarGANv2 [4] on the CelebaHQ [22] dataset, falsely considers hair color a domain-specific variation and ignores the domain-specific facial hair factor from the guidance image in the top row. We introduce a set of metrics for evaluation of semantic correctness in UMMI2I translation and provide an analysis of current SotA methods.

there are multiple images in the target domain corresponding to each example in the source domain. This makes the problem ill-posed and results in unreliable translation and unpredictable behaviour on the new unseen examples from the source domain.

One natural solution to this problem is to consider many-to-many mappings [14, 28, 2, 23], in which case the translation is guided by an example from the target domain that specifies domain-specific factors of the generated target image. For example, when an image of a female face is translated to the male domain, it is not clear whether the trans-

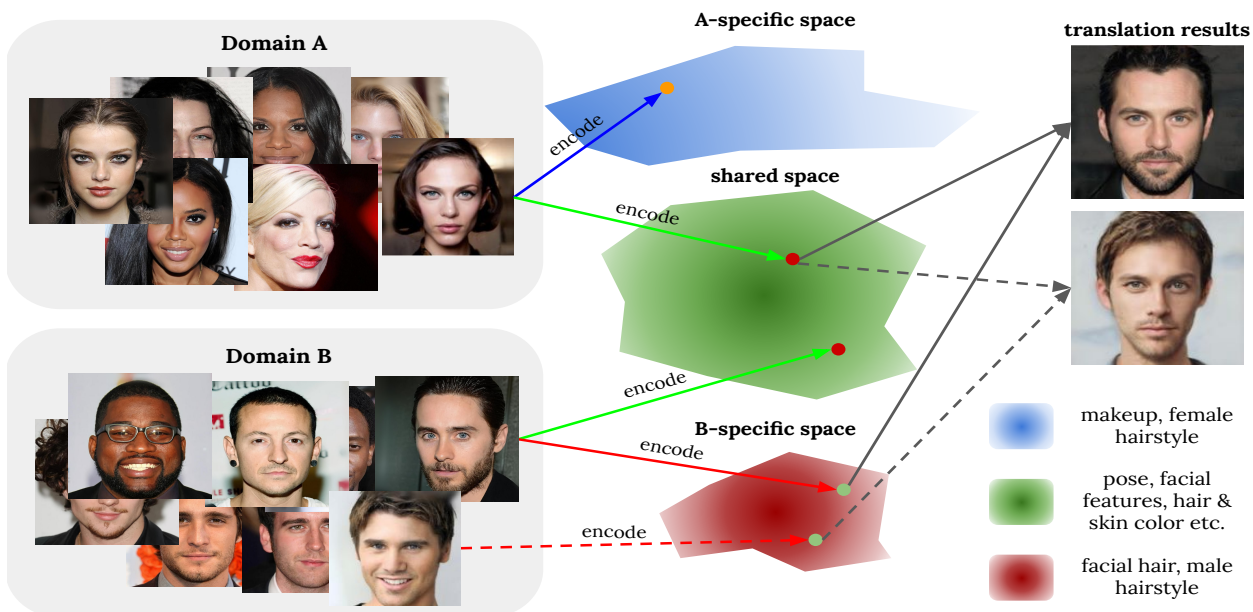


Figure 2. Overview of the common approach to disentanglement in the UMMI2I translation. Given an input image from the source domain and a guidance example from the target domain, the UMMI2I method produce an image from target domain containing content information from the input image and domain-specific attributes (e.g. facial hair) from the guidance example. A common approach to UMMI2I is to encode images to the domain-specific and domain-invariant embedding spaces. All source examples were taken from CelebA [22] dataset; the translation examples are generated with StarGANv2 [4] (*best viewed in color*).

lation result should have a beard or not. In the one-to-one setting, the model would either chose the most likely facial hair option or rely on some unrelated property of the input image (e.g. the intensity of the top-left pixel) to make this decision. In the many-to-many setting, however, the information about the desired facial hairstyle is taken from a guidance image from the target domain. The resulting translation problem for a given pair of input and guidance images has a single correct solution, and therefore lets us reason about the correctness of translation models. Measuring which architecture choices yield semantically correct UMMI2I translations, *i.e.* provide better disentanglement of domain-specific and domain-invariant factors of variation, is crucial, yet the majority of current state-of-the-art UMMI2I translation methods [14, 2, 23] are evaluated using only Frechet inception distance (FID) [10], learned perceptual similarity (LPIPS) [39] and other metrics that measure diversity and realism, and do not take the disentanglement quality into consideration. Unfortunately, existing disentanglement metrics proposed in the representation learning literature [16, 11, 17, 7, 3, 21, 6] are not suitable for evaluation of UMMI2I translation either, since they study the quality of disentanglement of latent embeddings, not generated images, and were designed for single-domain disentanglement, therefore not taking into account which variations are shared and which are domain-specific.

In this paper, we propose a new, data-driven approach for evaluation of unsupervised cross-domain disentanglement quality in UMMI2I methods. We designed three evalua-

tion protocols based on the synthetic 3D Shapes [17] dataset (originally designed for evaluation of single-domain disentanglement), a more challenging synthetic SynAction [33] pose dataset, and a widely used CelebA [27] dataset of faces.

1. To the best of our knowledge, we are the first to propose a set of metrics for evaluation of the semantic correctness of UMMI2I translation. Our metrics evaluate how well the shared attributes are preserved, how reliably the domain-specific attributes are manipulated, whether the translation result is a valid example of the target domain and whether the network collapsed to producing the same most frequent attribute values.

2. We create three evaluation protocols based on 3D-Shapes, SynAction and CelebA datasets, and measure disentanglement quality of the current state-of-the-art UMMI2I translation methods on them.

3. We show that for all tested methods there is a clear trade-off between content preservation and manipulation of the domain-specific variations, leading to subpar performance on at least one dataset. More specifically, all methods we tested poorly manipulated attributes associated with adding or changing certain parts of the objects (e.g. facial hair or smile) and AdaIN-based [13] methods showed an inductive bias towards treating spatial attributes, such as poses and position of objects in the scene, as the domain-invariant factors, and colors and textures as the domain-specific sources of variation, irrespective of which attributes were actually shared between domains and which are domain-specific.

2. Related Work

2.1. Many-to-Many Image Translation Methods

Style transfer. Perhaps one of the first steps towards producing diverse translation results using deep neural networks was the development of neural style transfer methods [8, 34, 9, 35, 15, 38, 20]. Conventionally, these methods perform texture transfer by matching the statistics of output features of the input image at a particular layer of a pretrained CNN with those of the guidance image. While style transfer methods show remarkable results in generating artistically stylized images, they are not suitable for general-purpose UMMI2I, because they, effectively, always treat texture and color palette as domain-specific attributes, which is often not sufficient to minimize the domain shift and does not guarantee content preservation.

Embedding reconstruction. Another line of work applied reconstruction losses to latent embeddings for disentanglement of domain-specific and shared information. For example, Almahairi et al. [2] proposed an augmented version of CycleGAN [40] with two additional domain-specific encoders. The disentanglement is enforced by the reconstruction loss on latent representations of pairs of source and target examples from the corresponding translation results. The disentanglement strategy introduced in DRIT [23], DRIT++ [24] and DMIT [37], involves encoding images from both domains to the shared latent space (content space) in addition to the domain-specific encoding. Additionally, those methods share part of the weights of the content encoders and have an additional content discriminator to force the disentanglement of the shared attributes.

AdaIN-based manipulation. Other UMMI2I translation methods, such as MUNIT [14], FUNIT [26] for few-shot translation, EGSC-IT [28] and StarGANv2 [4], used the Adaptive Instance Normalization (AdaIN) [13] to alter the intermediate representation of a generator with “global” (shared across all spatial locations) shift and scale parameters extracted from a guidance example. The AdaIN was originally used for style transfer, thus those methods share an assumption that the domain-specific factors of variation lie within texture and color variation. In addition to the AdaIN-based manipulation, these methods use separate encoders for the domain-specific information and fully or partially shared encoders for the domain-invariant information and reconstruction losses for the latent representations.

Domain-specific vs style attributes. Contrary to prior work [14], we prefer the term “domain-specific attribute” over “style attribute” to mentally distinguish the task from neural style transfer and highlight that in most applications domain-specific factors go beyond the variation in color palette and texture. We still use “content” and “shared attribute” as a shorthand for “domain-invariant factor”.

2.2. Translation and disentanglement metrics

Realism metrics. The metrics commonly used for the evaluation of the UMMI2I methods include Inception score (IS) [31] and Frechet Inception score (FID) [10] that measure the similarity between the distributions of the translated images and the target domain, and learned perceptual similarity score (LPIPS) [31] that measures their visual diversity. Neither of these metrics reflects whether domain-specific attributes were taken from the guidance image and whether shared attributes were correctly preserved.

Disentanglement metrics. The many-to-many image translation task combines unsupervised disentanglement with pixel-level domain adaptation. Metrics for measuring the quality of unsupervised disentanglement have been extensively explored in the recent literature on representation learning [16, 11, 17, 7, 3, 21, 6]. Perhaps the closest evaluation protocols to ours were used to measure the performance of FactorVAE [17], β -VAE [11] and by Eastwood et al. [7] that used toy data with known generative factors along with the learned factor regressor to infer the target factors from the latent embedding and comparing these predictions to the expected values. Directly inferring the generative factors from the latent representations, however, can be insufficient for evaluation of disentanglement in the UMMI2I task since the disentanglement in the latent space does not guarantee that the translation result is also disentangled. Additionally, these metrics were developed for single-domain disentanglement of *individual* factors of variation and are not directly suitable for the evaluation of disentanglement of two *groups* of factors from each other (shared factors from domain-specific ones). In this paper, we extend this attribute regression-based approach to unsupervised cross-domain many-to-many translation.

Human evaluation. For the majority of the datasets used for unsupervised image translation, it might not be immediately obvious which sources of variation are shared and which are domain-specific from visual inspection of several images alone. Therefore, it is often unclear what the semantically correct translation should look like from visual inspection of image inputs. This makes it difficult - if not impossible - to instruct human subjects to identify a semantically incorrect translation without clear instructions on which attributes to compare across which domains, resulting in subjects evaluating realism, diversity, and quality of texture transfer, and not the semantic correctness. On the other hand, if we use a controlled dataset with known attribute variations and explicitly instruct the subjects on which attributes are supposed to be transferred from the content or guidance image, it is possible to infer the disentanglement quality from their responses. If we have a reliable regression model for these attributes, then this evaluation pipeline can be fully-automated, which is exactly what

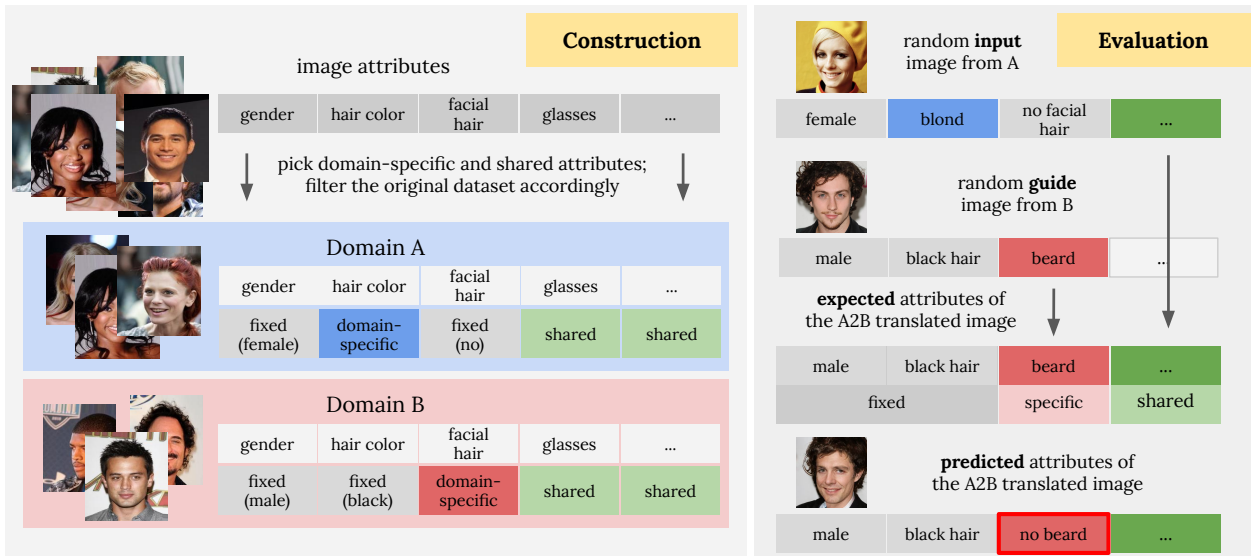


Figure 3. **Left:** Overview of the disentanglement dataset construction process. Given a dataset of images with some ground truth attribute annotation, we first choose a domain-splitting attribute (e.g. gender); then we choose two non-overlapping sets of the domain-specific attributes (e.g. facial hairstyle for male domain and hair color for female domain); we filter the original datasets so that the domain-specific attributes are varied only in one of the domains (e.g. only kept male photos with black hair). The resulting domains have known domain-specific attributes (hair color vs facial hairstyle) and shared attributes (everything else). **Right:** To evaluate the disentanglement of the translation example, we compare predicted attribute of the generated images with those of the perfectly disentangled result. The perfectly disentangled translation result has the same content attributes values as the input image, the target domain-specific attributes same as in the guidance image, and the correct values of the fixed attributes (domain-splitting attribute and the opposite domain-specific attributes).

we propose in this paper. In Section 6, we report results of a user study that indeed matches our automatic evaluation up to the annotator labeling error.

2.3. Limitations shared with prior work

Our evaluation protocol shares some limitations with the existing label-based disentanglement evaluation protocols described above, including being limited to datasets with attribute annotations and only focusing on attributes labeled by humans as “semantic” introducing some subjectivity. On the other hand, if a specific attribute that is important for a particular application is not properly disentangled, and it is not reflected by our label-based metric because this relevant attribute is not labeled in the data, we can either manually label that attribute or, at the very least, label pairs of images as having correct or wrong attribute transfer, and use it for evaluation. This suggests that, while possibly inevitably subjective, our label-based protocol can be applied in all cases in which human annotators are capable of, at the very least, labeling pairs of images as having same or different values for the most important attributes - which covers the absolute majority of useful practical applications.

3. Problem statement

Consider a set of images $A = \{a_i\}_{i=1}^{N_A}$ from the source domain \mathcal{A} , and a set of images $B = \{b_i\}_{i=1}^{N_B}$ from the target domain \mathcal{B} . The main assumption about \mathcal{A} and \mathcal{B} is that

images in these domains share some semantic structure but differ visually, i.e. daytime vs nighttime photos, CCTV vs front-facing videos, etc. Moreover, we assume that each domain has some domain-specific factors of variations that are either entirely absent or do not vary in the other domain, e.g. we can assume that beards and mustache are entirely absent in the female domain. According to the definition provided in MUNIT [14] and later used in other UMMI2I methods, the task of unsupervised many-to-many (also referred to as multimodal) image translation (UMMI2I) is to learn two mappings $M_{A2B} : \mathcal{A}, \mathcal{B} \rightarrow \mathcal{B}$ and $M_{B2A} : \mathcal{B}, \mathcal{A} \rightarrow \mathcal{A}$ that produce images from the target domain with shared factors matching the input image from the source domain and domain-specific factors matching the guidance image from the target domain. For example, assume that females (domain A) in our dataset have variable amount of makeup, and males (domain B) have variable amount of facial hair, and face expression and orientation vary in both. In this case, the correct $M_{A2B}(a, b)$ should be a male face with face expression and orientation as in the input female image a and facial hair as in the male guidance image b , and should not be affected by the amount of makeup in a . The degree to which the learned mappings satisfy the UMMI2I definition can be measured by computing the fraction of attributes that were correctly taken from respective images. In the next section we propose a set of metrics that factor the overall attribute correctness into several metrics that shed light on specific failure modes of different methods.

4. Proposed Evaluation Protocol

In this section, we describe proposed evaluation protocol for the UMMI2I problem defined in Section 3. It comprises two stages: construction of dataset pairs with known dataset-specific and invariant attributes, and estimation of the disentanglement quality on these datasets.

Dataset construction. We assume that the factors of variation of a given dataset can be approximated by a finite set of attributes. Unfortunately, the majority of in-the-wild image-to-image translation datasets (horses2zebras, winter2summer) have no attribute-level annotations beyond domain labels, so it is impossible to quantitatively measure the disentanglement quality using these datasets. For this reason, we propose to take existing complex single-domain datasets such as CelebA [27], 3D-Shapes [17] and SynAction [33], and split each of them into two domains in a way that ensures that domain-specific and domain-invariant attributes are known by design. The overall idea is to choose two sets of attributes that vary in each domain and filter the original single-domain dataset accordingly (see Figure 3). More formally, we consider a single labeled dataset $\mathcal{D} = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_i^N$ of N images \mathbf{x} with M labeled attributes $\mathbf{y} = (y_1, \dots, y_M)$.

We split attribute indices $\mathcal{Z} = (1, \dots, M)$ into four non-intersecting groups: a single *domain-splitting attribute* z_d that is constant within each datasets, but differs across datasets, e.g. gender, a set of *domain-invariant attributes* \mathcal{Z}_c that vary in both datasets, and two sets of *domain-specific attributes* \mathcal{Z}_s^A and \mathcal{Z}_s^B that vary only in respective datasets and are fixed in the other dataset:

$$\mathcal{Z} = \{z_d\} \sqcup \mathcal{Z}_c \sqcup \mathcal{Z}_s^A \sqcup \mathcal{Z}_s^B. \quad (1)$$

By definition, domain-specific attributes must be fixed in one domain and vary in the other, and the domain-splitting attribute should be fixed in both domains, so in order to fully specify the dataset, we also need to pick fixed values q^A, q^B for the domain-splitting attribute, and values $\{t_k^A\}_{k \in \mathcal{Z}_s^B}$ and $\{t_k^B\}_{k \in \mathcal{Z}_s^A}$ for domain-specific attributes in the domain in which these attributes are fixed. This way, the final dataset split can be fully specified as follows:

$$\begin{aligned} A &= \{\mathbf{x} \mid (y_{z_d} = q^A) \wedge (y_k = t_k^A), k \in \mathcal{Z}_s^B, (\mathbf{x}, \mathbf{y}) \in \mathcal{D}\}, \\ B &= \{\mathbf{x} \mid (y_{z_d} = q^B) \wedge (y_k = t_k^B), k \in \mathcal{Z}_s^A, (\mathbf{x}, \mathbf{y}) \in \mathcal{D}\}. \end{aligned}$$

As a result, we can use the subsets A and B as a pair of domains to measure the UMMI2I disentanglement, since their domain-specific attributes \mathcal{Z}_s^A and \mathcal{Z}_s^B and shared attributes \mathcal{Z}_c are known by design. Therefore, an ideal many-to-many image translation network M_{A2B} trained on a pair of datasets A, B and applied to an input image $a \in A$ and a guidance image $b \in B$ with ground truth attribute vectors $\mathbf{y}^a, \mathbf{y}^b$ should yield an image $M_{A2B}(a, b) \in B$ with

an attribute vector matching an correct attribute vector \mathbf{y}^* defined as:

$$\begin{aligned} y_k^* &= y_k^a \cdot \mathbb{1}[k \in \mathcal{Z}_c] + y_k^b \cdot \mathbb{1}[k \in \mathcal{Z}_s^B] + \\ &+ t_k^B \cdot \mathbb{1}[k \in \mathcal{Z}_s^A] + q^B \cdot \mathbb{1}[k = z_d]. \quad (2) \end{aligned}$$

The correct attribute vector for translation in the opposite direction ($M_{B2A}(b, a) \in A$) is defined analogously.

Computing disentanglement metrics. We start by predicting the attribute vector of the generated image using a pre-trained attribute regression network $R(x)$ as $\hat{\mathbf{y}}(a, b) = R(M_{A2B}(a, b))$. In practice, unless stated otherwise, to take into account the error of attribute prediction by R , we compute the perfect disentanglement attributes \mathbf{y}^* using attributes of input images a and b predicted by R . In all experiments, we use CNN-based predictors. A simple agreement metric between $\hat{\mathbf{y}}$ and \mathbf{y}^* fails to differentiate several important failure modes, namely: always generating the input image, or the guidance image, or always generating the same most frequent set of attributes and ignoring input images (*i.e.* mode collapse). To differentiate these failure modes, we introduce following metrics:

1. The *overall translation quality* measures whether the translation result semantically belongs to the target domain, but *not* the visual quality of generated images. It equals to the average probability of the translated image to have the corresponding fixed attribute values of the target domain given that the original input image had different attribute values:

$$Q_{\text{tr}}^{\text{A2B}} = \mathbb{E}_k \text{P}(\hat{y}_k = y_k^* \mid y_k^a \neq y_k^b), k \sim \mathcal{Z}_s^A \cup \{z_d\} \quad (3)$$

2. In order to estimate the *quality of preservation of shared attributes*, we measure the average probability of a shared attribute to be preserved after translation to the target domain given that input images had different values of this attribute:

$$D_c^{\text{A2B}} = \mathbb{E}_k \text{P}(\hat{y}_k = y_k^a \mid y_k^a \neq y_k^b), k \sim \mathcal{Z}_c \quad (4)$$

3. We measure the *quality of transfer of domain-specific attributes* similarly:

$$D_s^{\text{A2B}} = \mathbb{E}_k \text{P}(\hat{y}_k = y_k^b \mid y_k^a \neq y_k^b), k \sim \mathcal{Z}_s^B$$

4. Additionally, we can estimate the *model bias for certain attributes* by computing how likely the translation changes an attribute given that it is same for both input and guidance images:

$$B = \mathbb{E}_k \text{P}(\hat{y}_k \neq y_k^* \mid y_k^a = y_k^b), k \sim \mathcal{Z} \quad (5)$$

Bias metric B , which is the probability of the attribute to change if it was same for both the input and guidance

image, can be high only in two cases: 1) if a mode collapse appeared and the results are biased towards certain attribute values or 2) if the error of the attribute prediction is high on the translated images. We consider attribute predictions for high-bias translations that did not collapse to be low-confidence and report them in gray in tables. These results cannot be used to judge about the disentanglement quality. We also report the average content score $D_c = \frac{1}{2}(D_c^{A2B} + D_c^{B2A})/2$ and overall disentanglement quality $\bar{D} = \frac{1}{4}(D_s^{A2B} + D_s^{B2A} + D_c^{A2B} + D_c^{B2A})$.

5. Experiments

Disentanglement datasets. We modified 3D-Shapes [17], SynAction [33] and CelebA [27] to quantitatively evaluate UMMI2I correctness as described above. We note that most attributes can be grouped into one of three kinds: 1) color and texture; 2) shape and position; 3) presence and absence of objects. We split our attributes in a way that tests various kinds of attributes as domain-specific and domain-invariant, and recommend future researchers to follow the same principle on new datasets. Precise specifications of attribute splits used in each dataset, and values used for fixed attributes in other domains can be found in supplementary Section 8.1 and Table 3. At a high level: 1) **CelebA-D** (Fig. 6) uses males with black hair and variable amount of facial hair, and smiling young females with variable hair color as two domains - we opted for this specific split instead of the usual male-female split to have distinct domain-specific attributes of different kinds to each domain (color of hair in one, presence of beards in the other) while keeping domains large enough - there are, unfortunately, not many females in CelebA that are not young and conventionally attractive; 2) **Shapes-3D** (Fig. 7) contains 3D renders of one of four shapes in one of ten colors in both dataset, with floor and wall color variations in one domain, and size and view angle variation in the other domain; 3) **SynAction** (Fig. 8) contains 3D renders of humans in different poses with background texture variations in one domain and clothing and identity variations in the other.

Automatic evaluation. We present our results on the three disentanglement datasets described above for Augmented CycleGAN [2], DRIT++ [24], FUNIT [26], MUNIT [14], and StarGANv2 [4]. Additionally, we noticed that the original implementation of MUNIT performs poorly when the domain-specific embedding is taken from the real guidance image and not sampled randomly. We fixed this issue by adding the domain-specific embedding reconstruction loss and GAN loss for the non-random embedding and reported these results under the name “MUNITX”. For all methods, we used the official implementations and default hyperparameters. To provide a sense of scale, we added four “naive” baselines: “Random triplets” corresponds to a gen-

Model	$Q_{tr} \uparrow$	$D \uparrow$	$D_s^{A2B} \uparrow$	$D_s^{B2A} \uparrow$	$D_c \uparrow$	$B \downarrow$
DRIT++ [24]	87.8	45.2	49.1	25.1	53.3	12.8
MUNIT [14]	79.7	44.2	43.0	21.5	56.2	11.1
MUNITX	81.4	47.4	66.3	20.9	51.2	12.3
FUNIT [26]	92.1	39.1	31.6	31.4	46.7	16.0
AugCycleGAN [2]	90.9	42.1	38.5	41.3	44.2	14.8
StarGANv2 [4]	94.4	50.0	80.1	24.7	47.7	12.2
Rand. triplets	72.4	37.5	29.5	20.8	49.9	20.5
Rand. target	96.0	41.6	42.3	37.3	43.4	18.4
Content idt	0.0	50.0	0.0	0.0	100.0	0.0
Guidance idt	96.0	50.0	100.0	100.0	0.0	0.0

Table 1. UMMI2I disentanglement metrics on **CelebA-D** subset. The bottom part contains the results on the “naive” baselines described in Section 5. Gray indicates low confidence in the corresponding disentanglement scores.

erator that produces a completely random realistic output image from random domain, “Random target” - random realistic output image from the correct domain, “Content idt” - always producing the input (content) image, and “Guidance idt” - always producing the guidance image. We report full lists of attributes, description of the attribute prediction models architectures and reference performance (Fig. 14-16), per-attribute disentanglement performance, and more qualitative translation examples (Fig. 9-13) in the supplementary.

User study. We ran a blind user study on a random subset of 3D-Shapes asking people to label triplets images as having same or different attribute values (Fig. 17). We used human responses to these questions instead of predicted attribute values to estimate proposed metrics. Most annotated pairs had one generated and one real image, but a small fraction had two real images with known attributes to estimate the annotation error rate. We described our motivation for this rigid human evaluation protocol in favour of generic “does x look like a combined with b ?” in Section 2.2.

6. Results

Results in Tables 1 and 2 as well as the per-attribute results (supplementary Section 8.3) indicate the following:

1. AdaIN-based methods (MUNIT, MUNITX, FUNIT and StarGANv2) tend to treat color- and texture- related information as domain-specific factors, even if these factors are actually domain-invariant and should be preserved, and fail to capture spatially localized domain-specific factors. For example, MUNITX and StarGANv2 perform hair color (domain-specific) change correctly with more than a 80% accuracy on average in Celeb-A, and were able to correctly change the wall and floor hue (domain-specific) in Shapes-3D, and MUNITX correctly changed the identity and background (domain-specific) in SynAction. But on Shapes-3D,

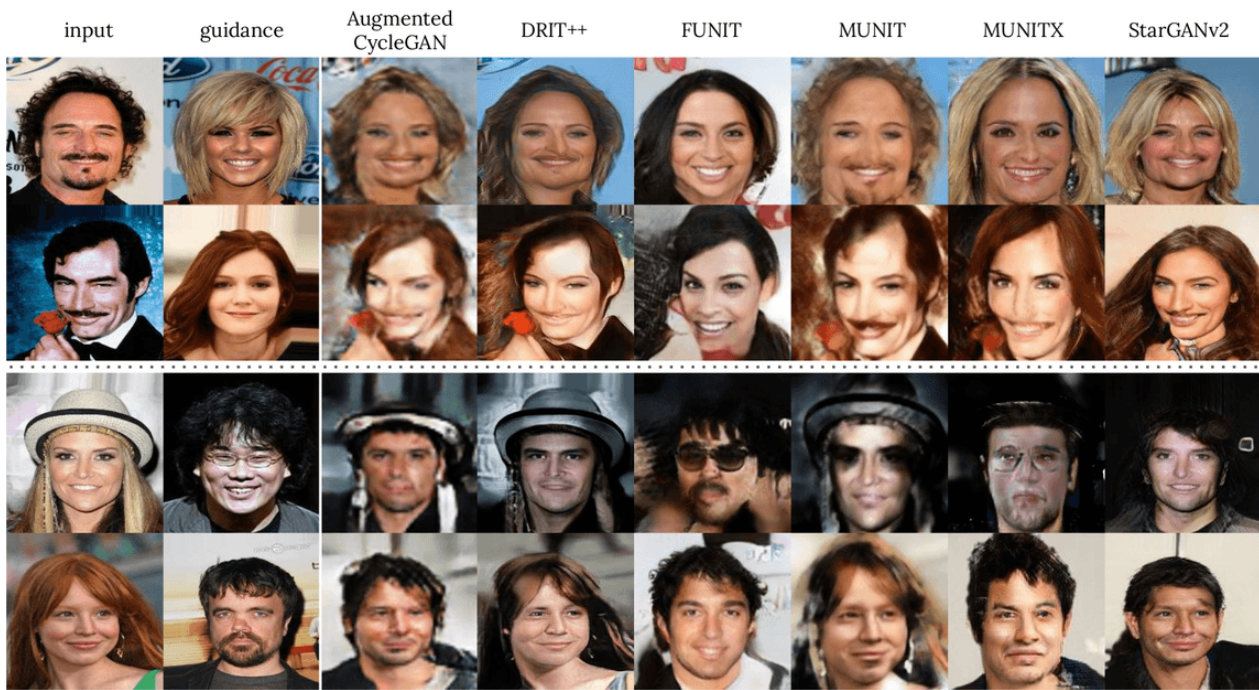


Figure 4. Examples of M→F (top) and F→M (bottom) translations on the proposed **CelebA-D** subset. A correct translation should have domain-specific attributes of the guidance image (hair color in the top two lines; facial hair, smile and age in the bottom two lines), and the rest of attributes (facial features, orientation, etc.) from the input image. More examples can be found in Figures 9-11 in supplementary.

these methods changed the object hue (domain-invariant) that should have been preserved, failed to correctly change the smile with the accuracy 34% and 22% respectively, or add a correct facial hairstyle (below 28% and 16.6%). Also, MUNITX failed to apply the background texture in correct spatial location on SynAction (see Figure 5).

2. Non-AdaIN-based methods (DRIT++, AugCycleGAN) preserved all domain-invariant attributes better across all experiments (e.g. correctly preserves eyeglasses (82%) and hats (81%) from the content image compared to MUNITX’s 54% and 57% respectively), but performed poorly in transferring any domain-specific information including color- and texture- based attributes with the average accuracy of hair color change $\approx 33\%$ and accuracy of correct wall and floor hue change on 3D-Shapes below 16%.

3. All methods preserved domain-invariant spatial information well, such as head orientation in CelebA or body pose in SynAction, with the translation pose being closer to that of the input image than the guidance image $>91\%$ of cases (except FUNIT and AugCycleGAN - see below). The best models in this aspect are MUNIT and MUNITX (98% and 76% on CelebA and SynAction respectively).

4. On all three datasets FUNIT experiments resulted in moderate to severe mode-collapse, resulting in low correctness and high bias scores. Low visual fidelity of images generated by AugCycleGAN prevented reliable automatic evaluation (metrics reported in gray in all tables).

5. The user study reported in Table 9 in supplementary, resulted in same trends as the automatic evaluation reported

in the main paper (Table 2), confirming that the discrepancy between automatic and human evaluation comes from human annotation error (observed human error rate $<10\%$).

More qualitative results for all methods on all datasets can be found in supplementary Figures 9-13. Overall, all methods performed well in preserving spatial domain-invariant features, and some performed better than others in transferring global textures and colors, but no single method was able to perform well across all three kinds of domain-specific and global domain-invariant factors of variation.

7. Conclusion

We believe that one can only improve upon the aspects of the model that can be quantitatively evaluated, and our work provides a much needed way to measure the semantic correctness of UMMI2I translation. In this paper, we formalized the notion of semantic correctness of image translation in terms of shared and domain-specific attributes and proposed three protocols for the evaluation of UMMI2I methods. The results of five SotA UMMI2I translation methods show that all methods, to different degrees, fail to infer which attributes are domain-specific and which are domain-invariant from data, and mostly rely on inductive biases hard-coded into their architectures. Consequently, these methods perform well only on a very limited range of applications that match their inductive biases. We conclude that more effort should be made to create UMMI2I methods that enable actual data-driven disentanglement of domain-specific and shared factors of variation.



Figure 5. Illustration of UMMI2I translation results on **3D-Shapes** (top) and **SynAction** (bottom) subsets. Domain-specific attributes in **3D-Shapes** are wall and floor color (A) / size and view angle (B), and in **SynAction** are background texture (A) / clothing and identity (B)

Model	3D Shapes						SynAction					
	$Q_{tr} \uparrow$	$D \uparrow$	$D_s^{A2B} \uparrow$	$D_s^{B2A} \uparrow$	$D_c \uparrow$	$B \downarrow$	$Q_{tr} \uparrow$	$D \uparrow$	$D_s^{A2B} \uparrow$	$D_s^{B2A} \uparrow$	$D_c \uparrow$	$B \downarrow$
DRIT++	93.5	32.7	10.5	14.9	52.6	43.746	89.6	55.7	27.3	12.4	91.6	25.7
MUNIT	94.9	68.1	76.3	99.7	48.1	5.12	86.8	66.9	52.2	19.0	98.2	23.8
MUNITX	92.9	33.2	18.6	10.2	52.0	38.2	92.1	69.4	58.3	23.7	97.9	20.5
FUNIT	25.0	14.6	18.2	5.39	17.4	66.5	37.4	27.1	6.79	0.5	50.9	15.2
AugCycleGAN	50.2	16.5	14.8	10.0	20.6	63.6	99.9	37.6	14.2	36.2	50.0	33.3
StarGANv2	95.9	29.0	10.3	89.2	8.28	26.0	96.8	31.3	17.6	8.69	49.6	33.1
Rand. triplets	54.7	11.7	7.77	4.55	17.4	59.4	49.9	28.0	5.85	6.31	50.0	36.5
Rand. target	99.9	14.9	13.5	10.1	18.0	49.5	99.9	30.9	11.1	12.7	50.0	18.0
Content idt	0.0	50.0	0.0	0.0	100.0	0.0	0.0	50.0	0.0	0.0	100	0.0
Guidance idt	99.8	50.0	100.0	100.0	0.0	0.0	99.8	50.0	100	100	0.0	0.0

Table 2. The UMMI2I disentanglement evaluation metrics on **3D-Shapes** and **SynAction**. The notation is consistent with in Table 1. AdaIN-based methods tend to treat color and texture-based attributes as domain-specific even if they are domain-invariant, while non-AdaIN methods preserve domain-invariant attributes better, but fail to modify color and texture even if they are domain-specific attributes.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Amjad Almahairi, Sai Rajeswar, Alessandro Sordani, Philip Bachman, and Aaron Courville. Augmented cyclegan: Learning many-to-many mappings from unpaired data. *arXiv preprint arXiv:1802.10151*, 2018.
- [3] Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in Neural Information Processing Systems*, 31:2610–2620, 2018.
- [4] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020.
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- [6] Kien Do and Truyen Tran. Theory and evaluation metrics for learning disentangled representations. *arXiv preprint arXiv:1908.09961*, 2019.
- [7] Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018.
- [8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016.
- [9] Golnaz Ghiasi, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network. *arXiv preprint arXiv:1705.06830*, 2017.
- [10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 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*, pages 6629–6640, 2017.
- [11] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- [12] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018.
- [13] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [14] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [16] Theofanis Karaletsos, Serge Belongie, and Gunnar Rätsch. Bayesian representation learning with oracle constraints. *arXiv preprint arXiv:1506.05011*, 2015.
- [17] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International Conference on Machine Learning*, pages 2649–2658. PMLR, 2018.
- [18] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *International Conference on Machine Learning*, pages 1857–1865. PMLR, 2017.
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [20] Dmytro Kotovenko, Artsiom Sanakoyeu, Sabine Lang, and Bjorn Ommer. Content and style disentanglement for artistic style transfer. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4422–4431, 2019.
- [21] Abhishek Kumar, Prasanna Sattigeri, and Avinash Balakrishnan. Variational inference of disentangled latent concepts from unlabeled observations. In *International Conference on Learning Representations*, 2018.
- [22] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [23] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51, 2018.
- [24] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Kumar Singh, and Ming-Hsuan Yang. Dri++: Diverse image-to-image translation via disentangled representations. *arXiv preprint arXiv:1905.01270*, 2019.
- [25] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. *arXiv preprint arXiv:1703.00848*, 2017.
- [26] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *arxiv*, 2019.
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of*

- International Conference on Computer Vision (ICCV)*, December 2015.
- [28] Liqian Ma, Xu Jia, Stamatios Georgoulis, Tinne Tuytelaars, and Luc Van Gool. Exemplar guided unsupervised image-to-image translation with semantic consistency. *arXiv preprint arXiv:1805.11145*, 2018.
- [29] George Papandreou, Tyler Zhu, Liang-Chieh Chen, Spyros Gidaris, Jonathan Tompson, and Kevin Murphy. Personlab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 269–286, 2018.
- [30] Nataniel Ruiz, Eunji Chong, and James M. Rehg. Fine-grained head pose estimation without keypoints. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018.
- [31] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016.
- [32] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [33] Ximeng Sun, Huijuan Xu, and Kate Saenko. Twostreamvan: Improving motion modeling in video generation. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 2744–2753, 2020.
- [34] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *ICML*, volume 1, page 4, 2016.
- [35] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017.
- [36] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *Proceedings of the IEEE international conference on computer vision*, pages 2849–2857, 2017.
- [37] Xiaoming Yu, Yuanqi Chen, Shan Liu, Thomas Li, and Ge Li. Multi-mapping image-to-image translation via learning disentanglement. In *Advances in Neural Information Processing Systems*, pages 2994–3004, 2019.
- [38] Lvmin Zhang, Yi Ji, Xin Lin, and Chunping Liu. Style transfer for anime sketches with enhanced residual u-net and auxiliary classifier gan. In *2017 4th IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 506–511. IEEE, 2017.
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [40] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [41] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z Li. Face alignment across large poses: A 3d solution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 146–155, 2016.