

Breaking the cycle—Colleagues are all you need

Ori Nizan

Technion, Israel

snizori@campus.technion.ac.il

Ayellet Tal

Technion, Israel

ayellet@ee.technion.ac.il

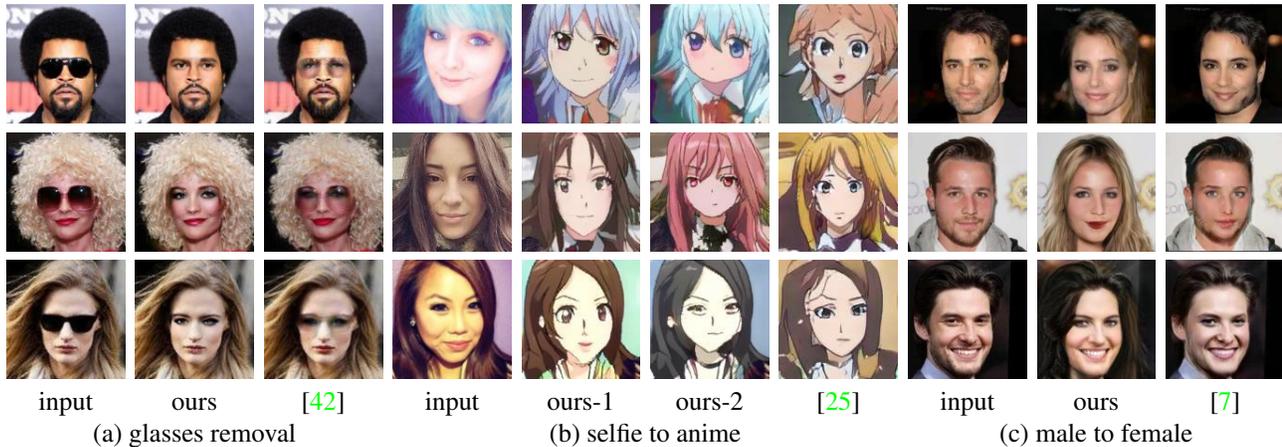


Figure 1: **Translation results for three applications.** Our method is unidirectional (cycles are unnecessary) and multi-modal (multiple results are generated for a given input, e.g. (b)). The results are compared to SOTA results and shown to outperform them. In (a) our method completely removes the glasses; in (b) the shape of the face is well maintained, and in (c) the women look more "feminine", e.g., no beard leftovers. More results & comparisons can be found later in the paper.

Abstract

This paper proposes a novel approach to performing image-to-image translation between unpaired domains. Rather than relying on a cycle constraint, our method takes advantage of collaboration between various GANs. This results in a multi-modal method, in which multiple optional and diverse images are produced for a given image. Our model addresses some of the shortcomings of classical GANs: (1) It is able to remove large objects, such as glasses. (2) Since it does not need to support the cycle constraint, no irrelevant traces of the input are left on the generated image. (3) It manages to translate between domains that require large shape modifications. Our results are shown to outperform those generated by state-of-the-art methods for several challenging applications.

1. Introduction

Mapping between different domains is inline with the human ability to find similarities between features in dis-

tinctive, yet associated, classes. Therefore it is not surprising that image-to-image translation has gained a lot of attention in recent years. Many applications have been demonstrated to benefit from it, yielding beautiful results.

In unsupervised settings, where no paired data is available, shared latent space and cycle-consistency assumptions have been utilized [2, 7, 8, 18, 21, 26, 30, 39, 45, 47]. Despite the successes & benefits, previous methods might suffer from some drawbacks. In particular, oftentimes, the cycle constraint might cause the preservation of source domain features, as can be seen for example, in Figure 1(c), where facial hair remains on the faces of the women. This is due to the need to go back and forth through the cycle. Second, as discussed in [25], sometimes the methods are unsuccessful for image translation tasks with large shape change, such as in the case of the anime in Figure 1(b). Finally, as explained in [42], it is still a challenge to completely remove large objects, like glasses, from the images, and therefore this task is left for their future work (Figure 1(a)).

We propose a novel approach, termed *Council-GAN*, which handles these challenges. The key idea is to rely

on "collegiality" between GANs, rather than utilizing a cycle. Specifically, instead of using a single pair of a generator/discriminator "experts", it utilizes the collective opinion of a group of pairs (the *council*) and leverages the variation between the results of the generators. This leads to a more stable and diverse domain transfer.

To realize this idea, we propose to train a council of multiple council members, requiring them to learn from each other. Each generator in the council gets the same input from the source domain and will produce its own output. However, the outputs produced by the various generators should have some common denominator. For this to happen across all images, the generators have to find common features in the input, which are used to generate their outputs. Each discriminator learns to distinguish between the generated images of its own generator and images produced by the other generators. This forces each generator to converge to a result that is agreeable by the others. Intuitively, this convergence assists to maximize the mutual information between the source domain and the target domain, which explains why the generated images maintain the important features of the source images.

We demonstrate the benefits of our approach for several applications, including glasses removal, face to anime translation, and male to female translation. In all cases we achieve state-of-the-art results.

Hence, this paper makes the following contributions:

1. We introduce a novel model for unsupervised image-to-image translation, whose key idea is collaboration between multiple generators. Conversely to most recent methods, our model avoids cycle-consistency constraints altogether.
2. Our council manages to achieve state-of-the-art results in a variety of challenging applications.

2. Related work

Generative adversarial networks (GANs). Since the introduction of the GAN framework [15], it has been demonstrated to achieve eye-pleasing results in numerous applications. In this framework, a generator is trained to fool a discriminator, whereas the latter attempts to distinguish between the generated samples and real samples. A variety of modifications have been proposed in recent years in an attempt to improve GAN's results; see [3, 10, 11, 20, 24, 33, 36, 38, 40, 43, 46] for a few of them.

We are not the first to propose the use of multiple GANs [12, 14, 17, 23]. However, previous approaches differ from ours both in their architectures and in their goals. For instance, some of previous architectures consist of multiple discriminators and a single generator; conversely, some propose to have a key discriminator that can

evaluate the generators' results and improve them. We propose a novel architecture to realize the concept of a council, as described in Section 3. Furthermore, the goal of other approaches is either to push each other apart, to create diverse solutions, or to improve the results. Our council attempts to find the commonalities between the the source and target domains. By requiring the council members to "agree" on each other's results, they in fact learn to focus on the common traits of the domains.

Image-to-image translation. The aim is to learn a mapping from a source domain to a target domain. Early approaches adopt a supervised framework, in which the model learns paired examples, for instance using a conditional GAN to model the mapping function [22, 44, 48].

Recently, numerous methods have been proposed, which use unpaired examples for the learning task and produce highly impressive results; see for example [9, 13, 21, 26, 28, 30, 42, 47], out of a truly extensive literature. This approach is vital to applications for which paired data is unavailable or difficult to gain. Our model belongs to the class of GAN models that do not require paired training data.

A major concern in the unsupervised approach is the type of properties of the source domain that should be preserved. Examples include pixel values [41], pixel gradients [6], pairwise sample distances [4], and recently mostly cycle consistency [26, 45, 47]. The latter enforces the constraint that translating an image to the target domain and back, should obtain the original image. Our method avoids using cycles altogether. This has the benefit of bypassing unnecessary constraints on the generated output, and thus avoiding to preserve hidden information [8].

Most existing methods lack diversity in the results. To address this problem, some methods propose to produce multiple outputs for the same given image [21, 28]. Our method enables image translation with diverse outputs, however it does so in a manner in which all GANs in the council "acknowledge" to some degree each other's output.

Ensemble methods. These methods use multiple learning algorithms, trained individually [34, 35, 37], whose predictions are combined. They seek to promote diversity among the models they combine. Conversely, we require the council to learn together and "converge" to agreeable solutions.

3. Model

This section describes our proposed model, which addresses the drawbacks described in Section 1. Our model consists of a set, termed a *council*, whose members influence each other's results. Each member of the council has one generator and a couple of discriminators, as described below. The generators need not converge to a specific output; instead, each produces its own results, jointly generating a diverse set of results. During training, they take into account the images produced by the other generators. Intu-

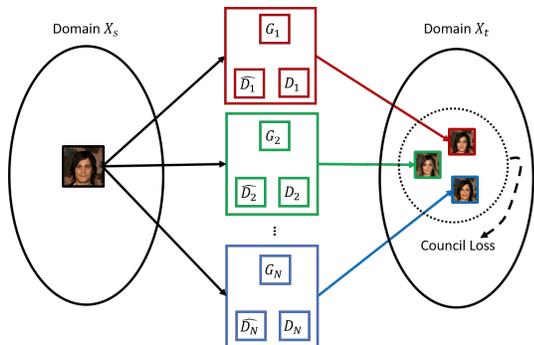


Figure 2: **General approach.** The council consists of triplets, each of which contains a generator and two discriminators: D_i distinguishes between the generator’s output and real examples, whereas \hat{D}_i distinguishes between images produced by G_i and images produced by other generators in the council. \hat{D}_i is the reason that the each of the generators converges to a result that is agreed-upon by all other members of the council.

itively, the mutual influence enforces the generators to focus on joint traits of the images in the source domain, which could be matched to those in the target domain. For instance, in Figure 1, to transform a male into a female, the generators focus on the structure of the face, on which they can all agree upon. Therefore, this feature will be preserved, which can explain the good results.

Furthermore, our model avoids cycle constraints. This means that there is no need to go in both directions between the source domain and the target domains. As a result, there is no need to leave traces on the generated image (e.g. glasses) or to limit the amount of change (e.g. anime).

To realize this idea, we define a council of N members as follows (Figure 2). Each member i of the council is a triplet, whose components are a single generator G_i and two discriminators D_i & \hat{D}_i , $1 \leq i \leq N$. The task of discriminator D_i is to distinguish between the generator’s output and real examples from the target domain, as done in any classical GAN. The goal of discriminator \hat{D}_i is to distinguish between images produced by G_i and images produced by the other generators in the council. This discriminator is the core of the model and this is what differentiates our model from the classical GAN model. It enforces the generator to converge to images that could be acknowledged by all council members—images that share similar features.

The loss function of D_i is the classical adversarial loss of [33]. Hereafter, we focus on the loss function of \hat{D}_i , which makes the outputs of the various generators share common traits, while still maintain diversity. At every iteration, \hat{D}_i gets as input pairs of (input,output) from all the generators in the council. Rather than distinguishing between real & fake, \hat{D}_i ’s distinguishes between the result

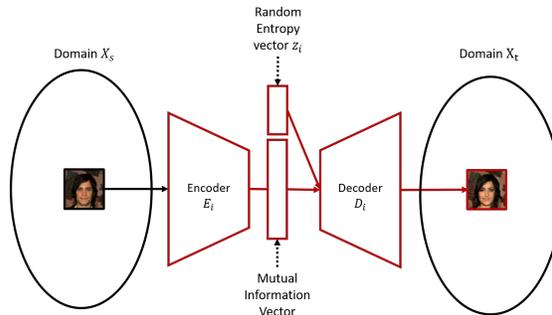


Figure 3: **Zoom into the generator G_i .** Our generator is an auto-encoder architecture, which is similar to that of [21]. The encoder consists of several strided convolutional layers followed by residual blocks. The decoder gets the encoded image (termed the *mutual information vector*), as well as a random entropy vector. The latter may be interpreted as encoding the leftover information of the target domain. The decoder uses a MLP to produce a set of AdaIN parameters for the random entropy vector [19].

of ”my-generator” and the result of ”another-generator”. Hence, during training, G_i attempts to minimize the distance between the outputs of the generators. Note that getting the input and not only the output is important to make the connection, for each pair, between the features of the source image and those of the generated image.

Let X_s be the source domain and X_t be the target domain. In our model we have N mappings $G_i : X_s \rightarrow X_t$. Given an image $x \in X_s$, a straightforward adaptation of the classical adversarial loss to our case would be:

$$Naive_council_loss_i(G_i, \hat{D}_i, \{G_j\}_{j \neq i}, X_s) = \mathbb{E}_{x \sim p(X_s)} \sum_{j \neq i} [\log(1 - \hat{D}_i(G_i(x), x)) + \log(\hat{D}_i(G_j(x), x))], \quad (1)$$

where G_i tries to generate images $G_i(x)$ that look similar to images from domains $G_j(x)$ for $j \neq i$. In analogy to the classical adversarial loss, in Equation (1), both terms should be minimized, where the left term learns to ”identify” its corresponding generator G_i as ”fake” and the right term learns to ”identify” the other generators as ”real”.

To allow multimodal translation, we encode the input image, as illustrated in Figure 3, which zooms into the structure of the generator [21]. The encoded image should carry useful (mutual) information between domains X_s and X_t . Let E_i be the i^{th} encoder for the source image and let z_i be the i^{th} random entropy vector, associated with the i^{th} member of the council, $1 \leq i \leq N$. z_i enables each generator to generate multiple diverse results. Equation (1) is modified so as to get an encoded image (instead of the original input

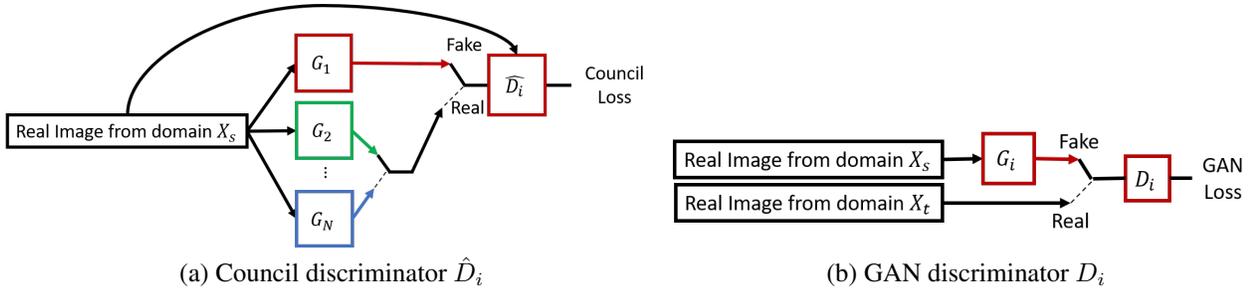


Figure 4: **Differences & similarities between \hat{D}_i and D_i .** While the GAN discriminator distinguishes between "real" and "fake" images, the council discriminator distinguishes between outputs of its own generator and those produced by other generators. Furthermore, while the GAN's discriminator gets as input only the generator's output, the council's discriminator gets also the generator's input. This is because we wish the generator to produce a result that bares similarity to the input image, and not only one that looks real in the target domain.

image) and the random entropy vector. The loss function of \hat{D}_i is then defined as:

$$Council_Loss_i(G_i, \hat{D}_i, \{G_j\}_{j \neq i}, X_s, z_i, \{E_j\}_{1 \leq j \leq N}) = (2)$$

$$\mathbb{E}_{x \sim p(X_s)} \sum_{j \neq i} [\log(1 - \hat{D}_i(G_i(E_i(x), z_i), x)) + \log(\hat{D}_i(G_j(E_j(x), \alpha z_j), x))].$$

Here, the loss function gets, as additional inputs, all the encoders and vector z_i . α controls the size of the sub-domain of the other generators, which is important in order to converge to "acceptable" images.

Figure 4 illustrates the differences and the similarities between discriminators D_i and \hat{D}_i . Both should distinguish between the generator's results and other images; in the case of D_i the other images are real images from the target domain, whereas in the case of \hat{D}_i , they are images generated by other generators in the council. Another fundamental difference is their input: \hat{D}_i gets not only the generator's output, but also its input. This aims at producing a resulting image that has common features with the input image.

Final loss. For each member of the council, we jointly train the generator (assuming the encoder is included) and the discriminators to optimize the final objective. In essence, G_i , D_i , & \hat{D}_i play a three-way min-max-max game with a value function $V(G_i, D_i, \hat{D}_i)$:

$$\min_{G_i} \max_{D_i} \max_{\hat{D}_i} V(G_i, D_i, \hat{D}_i) (3)$$

$$= GAN_Loss_i + \lambda Council_Loss_i.$$

This equation is a weighted sum of the adversarial loss GAN_Loss_i (of D_i), as defined in [33], and the $Council_Loss_i$ (of \hat{D}_i) from Equation (2). λ controls the importance of looking more "real" or more inline with the other generators. High values will result in more similar

images, whereas low values will require less agreement and result in higher diversity between the generated images.

Focus map. For some applications, it is preferable to focus on specific areas of the image and modify only them, leaving the rest of the image untouched. This can be easily accommodated into our general scheme, without changing the architecture.

The idea is to let the generator produce not only an image, but also an associated focus map, which essentially segments the learned objects in the domain from the background. All that is needed is to add a fourth channel, $mask_i$, to the generator, which would generate values in the range $[0, 1]$. These values can be interpreted as the likelihood of a pixel to belong to the background (or to an object). To realize this, Equation (3) becomes

$$\min_{G_i} \max_{D_i} \max_{\hat{D}_i} V(G_i, D_i, \hat{D}_i) (4)$$

$$= GAN_Loss_i + \lambda_1 Council_Loss_i + \lambda_2 Focus_Loss_i,$$

where

$$Focus_Loss_i = \delta \left(\sum_k mask_i[k] \right)^2 (5)$$

$$+ \sum_k \frac{1}{|mask_i[k] - 0.5| + \epsilon}.$$

In Equation (5), $mask_i[k]$ is the value of the 4th channel for pixel k . The first term attempts to minimize the size of the focus mask, i.e. make it focus solely on the object. The second term is in charge of segmenting the image into an object and a background (1 or 0). This is done in order to avoid generating semi-transparent pixels. In our implementation $\epsilon = 0.01$. The result is normalized by the image size. The values of λ_1 and λ_2 are application-dependent and will be defined for each application in Section 5.

Figure 5 illustrates the importance of the various losses. If only the $Focus_Loss$ (jointly with the GAN_Loss) is

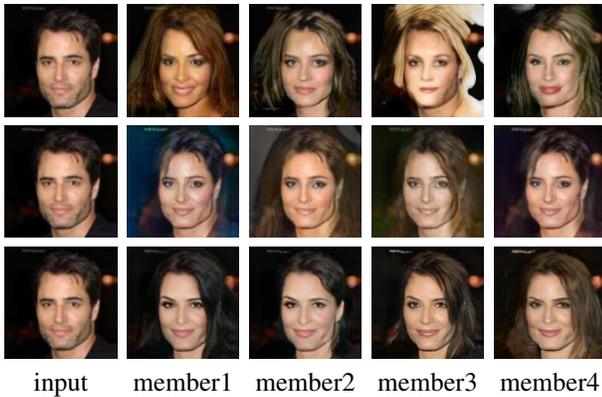


Figure 5: **Importance of the loss function components.** This figure shows the results generated by the four council members for the male-to-female application, after 100K iterations. **Top:** Using the *Focus.loss* (jointly with the classical *GAN.loss*) generates nice images from the target domain, which are not necessarily related to the given image. **Middle:** Using the *Council.loss* instead, relates between the input and the output faces, but might change the environment (background). **Bottom:** Our loss, which combines the above losses, both relates the input and the output faces and focuses only on facial modifications.

used, the faces of the input and the output are completely unrelated, though the quality of the images is good and the background does not change in most cases. Using only the *Council.loss*, the faces of the input and the output are nicely related, but the background might change. Our loss, which combines the above losses, produces the best results.

We note that this idea of adding a 4th channel, which makes the generator focus on the proper areas of the image, can be used in other GAN architectures. It is not limited to our proposed council architecture.

4. Experiments

4.1. Experiment setup

We applied our council GAN to several challenging image-to-image translation tasks (Section 4.2).

Baseline models. Depending on the application, we compare our results to those of some state-of-the-art models, including CycleGAN [47], MUNIT [21], DRIT++ [28, 29], U-GAT-IT [25], StarGAN [7], Fixed-PointGAN [42]. These methods are unsupervised and use cycle constraints. Out of these methods, MUNIT [21] and DRIT++ [28, 29] are multi-modal and generate several results for a given image. The others produce a single result. Furthermore, StarGAN [7] performs translation between multiple domains.

Datasets. We evaluated the performance of our system on the following datasets.

CelebA [31]. This dataset contains 202,599 face images of celebrities, each annotated with 40 binary attributes. We focus on two attributes: (1) the gender attribute and (2) with/without glasses attribute. The training dataset contains 68,261 (/10,521) images of males (/with glasses) and 94,509 (/152,249) images of females (/without glasses). The test dataset consists of 16,173 (/2,672) males (/with glasses) and 23,656 (/37,157) females (/without glasses).

selfie2anime [25]. The size of the training dataset is 3,400 selfie images and 3,400 anime images. The size of the test dataset is 100 selfie images and 100 anime images.

Training. All models were trained using Adam [27] with $\beta_1 = 0.5$ and $\beta_2 = 0.999$. For data augmentation we flipped the images horizontally with a probability of 0.5.

For the selfie/anime dataset, where the number of images is small, we augmented the data also with color jittering with up to $hue = 0.15$, random Grayscale with a probability of 0.25, random Rotation with up to 35° , random translation of up to 0.1 of the image, and with random perspective with distortion scale of 0.35 with a probability of 0.5. On the last 100K iterations we trained only on the original data, without augmentation. We performed one generator update after a number of discriminator updates that is equal to the size of the council. The batch size was set to 3 for all experiments. We trained all models with a learning rate of 0.0001, where the learning rate drops by a factor of 0.5 after every 100,000 iterations. The focus and council losses were added after 10,000 iterations.

Computational cost. The training takes about twice the time comparable to CycleGAN, when the council members run sequentially on the same GPU. The longer time is due to (1) having 4 members (2) a longer iteration of the council-discriminator, and (3) twice as many iterations needed to reach an agreement. It only takes twice as long since we avoid learning of the reverse side (e.g. from anime to selfie). To accelerate the computation, the members could be run in parallel or a smaller council could be used.

Evaluation. We verify our results both qualitatively and quantitatively. For the latter, we use two common measures: (1) the *Frechet Inception Distance score (FID)* [16], which calculates the distance between the feature vectors of the real and the generated images; (2) the *Kernel Inception Distance (KID)* [5], which improves on FID and measures GAN convergence.

4.2. Experimental results

Experimental results for male-to-female translation.

Given an image of a male face, the goal is to generate a female face, which resembles the male face [1, 32]. As explained in [1], three features make this translation task challenging: (i) There is no predefined correspondence in real data of each domain. (ii) The relationship is many-to-many between domains, as many male-to-female mappings

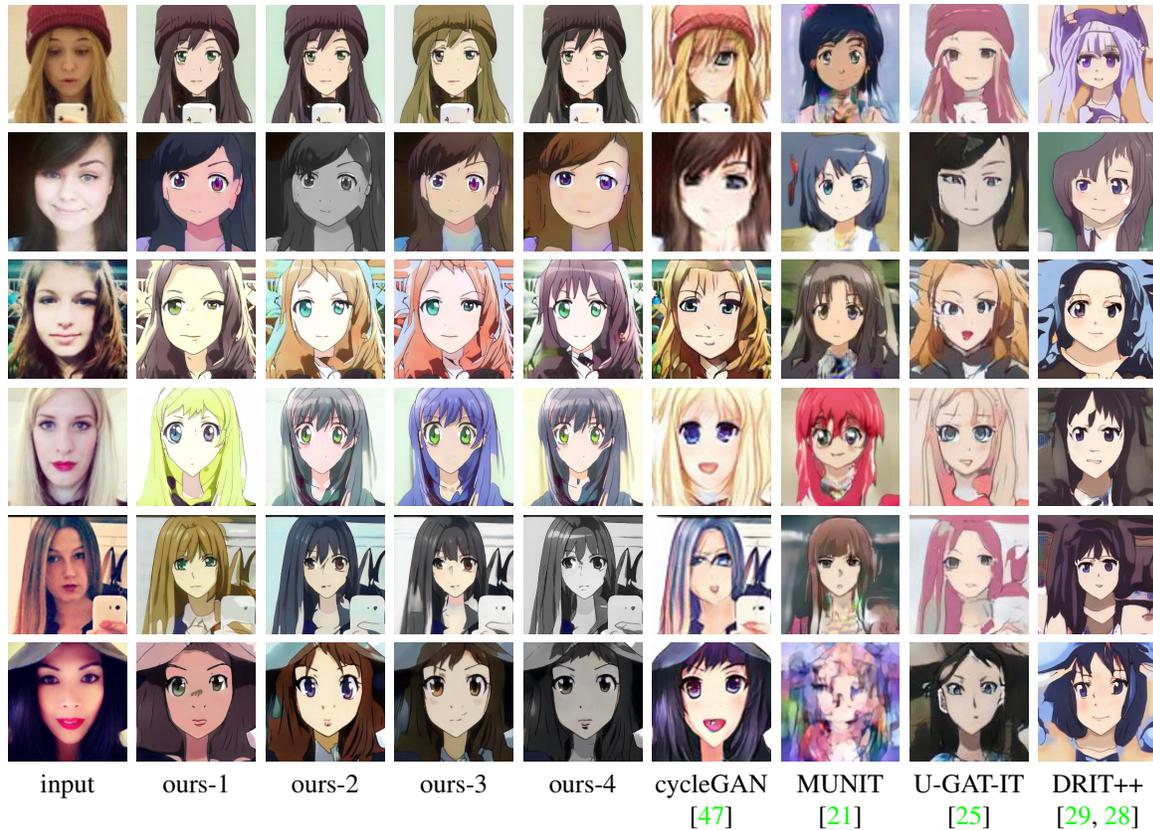


Figure 7: **Selfie-to-anime translation.** Our results preserve the structure of the face in the input image, while generating the characteristic features of anime, such as the large eyes.

	FID	KID
CycleGAN [47]	149.38	0.0056
MUNIT [21]	131.69	0.0057
U-GAT-IT [25]	115.11	0.0043
DRIT++ [29, 28]	109.22	0.0020
Council	101.39	0.0020

Table 2: **Quantitative results for selfie-to-anime translation.** Our results outperform those of other methods when FID is considered and are competitive for KID.

can be explained by the fact that it is easier for the council members to "agree" on features that exist in the input. Table 2 shows quantitative results. It can be seen that our results outperform or are competitive with those of other methods in both evaluation metrics.

Experimental results for glasses removal. Given an image of a person with glasses, the goal is to generate an image of the same person, but with the glasses removed. While in the previous application, the whole image changes, here the challenge is to modify only a certain part of the face and leave the rest of the image untouched.

	FID	KID
cycleGAN [47]	50.72	0.0038
Fixed-point GAN [42]	55.26	0.0041
Council	36.38	0.0026

Table 3: **Quantitative results of glasses removal.** The results of our council outperform state-of-the-art results

Figure 8 compares our results (using a council of four) to those of [42], which shows results for this application, as well as to [47]. Our generated images leave considerably less traces of the removed glasses. Again, this can be attributed lack of the cycle constraint. Table 3 provides quantitative results. For this application as well, our council manages to outperform other methods and address the challenge of removing large objects.

5. Implementation

Our code is based on PyTorch; it is available at <https://github.com/Onr/Council-GAN>. We set the major parameters as follows: α , which controls diversity (Equa-

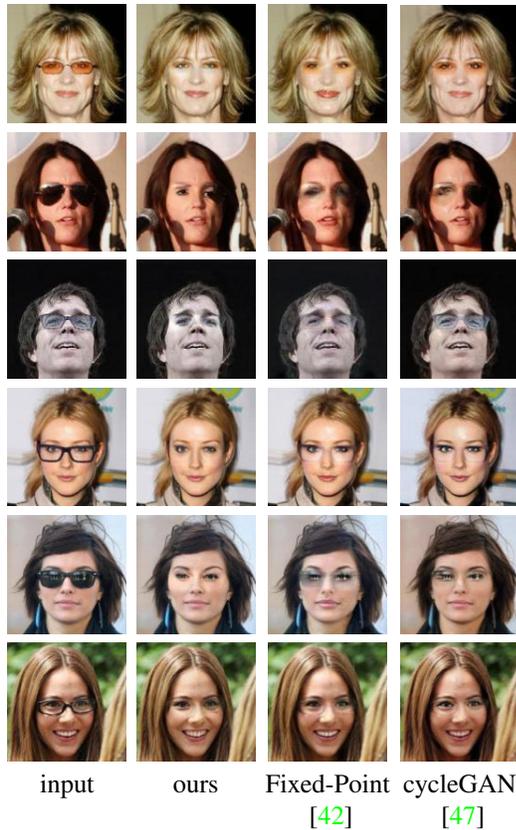


Figure 8: **Glasses removal.** We show a single result per input, since multi-modality is irrelevant for this application. Our generated images remove the glasses almost completely, whereas traces are left in [42]’s and in [47]’s results.

tion (2)), is set to 0.8. δ , which controls the size of the mask (Equation (5)), is set to 0.001. λ_1 and λ_2 from Equation (4) are set according to the applications: in male to female $\lambda_1 = 0.2$ & $\lambda_2 = 0.025$; in selfie to anime $\lambda_1 = 0.5$ & $\lambda_2 = 0$; in glasses removal $\lambda_1 = 0.2$ & $\lambda_2 = 0.2$.

Figure 9 studies the influence of the number of members and the number of iterations on the quality of the results. We focus on the male-to-female application, which is representative. The fewer the number of members in the council, the faster the convergence is. However, this comes at a price: the accuracy is worse. Furthermore, it can be seen that the KID improves with iterations, as expected.

Limitations. Figure 10 demonstrates a limitation of our method. When removing the glasses, the face might also become more feminine. This is attributed to the imbalance inherent to the dataset. Specifically, the ratio of the men to women with glasses is 0.8, whereas the ratio of men to women without glasses is only 0.4. The result of this imbalance in the target domain is that removing glasses also means becoming more feminine. This problem can

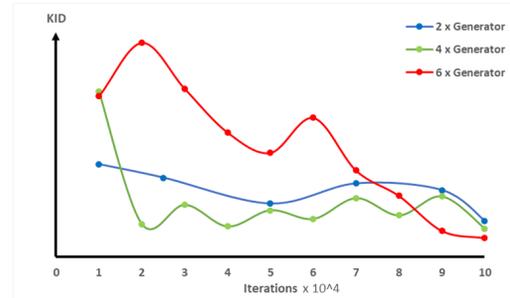


Figure 9: **KID as a function of # iterations.** The more iterations, the better KID. Moreover, with more council members, model converges more slowly, yet the results improve.

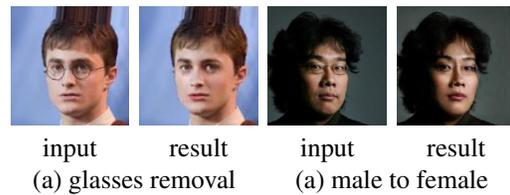


Figure 10: **Limitation.** (a) When removing the glasses, the face also becomes more feminine. (b) Conversely, when transforming a male to a female, the glasses may also be removed. This is attributed to high imbalance of the relevant features in the dataset.

be solved by providing a dataset with an equal number of males and females with and without glasses. Handling feature imbalance without changing the number of images in the dataset, is an interesting direction for future research.

6. Conclusion

This paper introduces the concept of a council of GANs—a novel approach to perform image-to-image translation between unpaired domains. Their key idea is to replace the widely-used cycle-consistency constraint by leveraging collaboration between GANs. Council members assist each other to improve, each its own result.

Furthermore, the paper proposes an implementation of this concept and demonstrates its benefits for three challenging applications. The members of the council generate several optional results for a given input. They manage to remove large objects from the images, not to leave redundant traces from the input and to handle large shape modifications. The results outperform those of SOTA algorithms both quantitatively and qualitatively.

Acknowledgements: We gratefully acknowledge the support of the Israel Science Foundation (ISF) 1083/18, PMRI-Peter Munk Research Institute–Technion, and NVIDIA Corporation with the donation of the GPU.

References

- [1] A. Almahairi, S. Rajeshwar, A. Sordoni, P. Bachman, and A. Courville. Augmented CycleGAN: learning many-to-many mappings from unpaired data. In *35th International Conference on Machine Learning (ICML)*, 2018. 5
- [2] A. Anoosheh, E. Agustsson, and R. Timofte. Combogan: Unrestrained scalability for image domain translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 896–8967, 2018. 1
- [3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *34th International Conference on Machine Learning (ICML)*, 2017. 2
- [4] S. Benaim and L. Wolf. One-sided unsupervised domain mapping. In *Advances in neural information processing systems*, 2017. 2
- [5] M. Binkowski, D. J. Sutherland, M. Arbel, and A. Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations*, 2018. 5
- [6] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *IEEE conference on computer vision and pattern recognition (CVPR)*, 2017. 2
- [7] Y. Choi, M. Choi, M. Kim, J. Ha, S. Kim, and J. Choo. StarGAN: unified generative adversarial networks for multi-domain image-to-image translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, June 2018. 1, 5, 6
- [8] C. Chu, A. Zhmoginov, and M. B. Sandler. CycleGAN, a master of steganography. *CoRR*, 2017. 1, 2
- [9] L. M. David Berthelot, Tom Schumm. BEGAN: boundary equilibrium generative adversarial networks. *CoRR*, 2017. 2
- [10] E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015. 2
- [11] A. Dosovitskiy and T. Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems*, pages 658–666, 2016. 2
- [12] I. P. Durugkar, I. Gemp, and S. Mahadevan. Generative multi-adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [13] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, 2016. 2
- [14] A. Ghosh, V. Kulharia, V. P. Namboodiri, P. H. Torr, and P. K. Dokania. Multi-agent diverse generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8513–8521, 2018. 2
- [15] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2
- [16] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. 5
- [17] Q. Hoang, T. D. Nguyen, T. Le, and D. Phung. MGAN: Training generative adversarial nets with multiple generators. In *International Conference on Learning Representations*, 2018. 2
- [18] X. Hua, D. Rempe, and H. Zhang. Unsupervised cross-domain image generation. In *International Conference on Learning Representations (ICLR)*, 2017. 1
- [19] X. Huang and S. J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1510–1519, 2017. 3
- [20] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie. Stacked generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5077–5086, 2017. 2
- [21] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal unsupervised image-to-image translation. In *The European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 3, 5, 6, 7
- [22] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [23] F. Juefei-Xu, V. N. Boddeti, and M. Savvides. Gang of gans: Generative adversarial networks with maximum margin ranking. *CoRR*, 2017. 2
- [24] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [25] J. Kim, M. Kim, H. Kang, and K. Lee. U-GAT-IT: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *CoRR*, abs/1907.10830, 2019. 1, 5, 6, 7
- [26] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim. Learning to discover cross-domain relations with generative adversarial networks. In *34th International Conference on Machine Learning-Volume 70 (ICML)*, 2017. 1, 2
- [27] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2014. 5
- [28] H.-Y. Lee, H.-Y. Tseng, J.-B. Huang, M. Singh, and M.-H. Yang. Diverse image-to-image translation via disentangled representations. In *The European Conference on Computer Vision (ECCV)*, 2018. 2, 5, 6, 7
- [29] H.-Y. Lee, H.-Y. Tseng, Q. Mao, J.-B. Huang, Y.-D. Lu, M. K. Singh, and M.-H. Yang. Drit++: Diverse image-to-image translation via disentangled representations. *arXiv preprint arXiv:1905.01270*, 2019. 5, 6, 7
- [30] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, 2017. 1, 2
- [31] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 5

- [32] Y. Lu, Y.-W. Tai, and C.-K. Tang. Attribute-guided face generation using conditional CycleGAN. In *European Conference on Computer Vision (ECCV)*, pages 282–297, 2018. 5
- [33] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley. Least squares generative adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 3, 4
- [34] D. Opitz and R. Maclin. Popular ensemble methods: An empirical study. *Journal of artificial intelligence research*, 11:169–198, 1999. 2
- [35] R. Polikar. Ensemble based systems in decision making. *IEEE Circuits and systems magazine*, 6(3):21–45, 2006. 2
- [36] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *4th International Conference on Learning Representations, ICLR*, 2016. 2
- [37] L. Rokach. Ensemble-based classifiers. *Artificial Intelligence Review*, 33(1-2):1–39, 2010. 2
- [38] M. Rosca, B. Lakshminarayanan, D. Warde-Farley, and S. Mohamed. Variational approaches for auto-encoding generative adversarial networks. *CoRR*, 2017. 2
- [39] A. Royer, K. Bousmalis, S. Gouws, F. Bertsch, I. Mosseri, F. Cole, and K. Murphy. Xgan: Unsupervised image-to-image translation for many-to-many mappings. *arXiv preprint arXiv:1711.05139*, 2017. 1
- [40] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, 2016. 2
- [41] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2107–2116, 2017. 2
- [42] M. M. R. Siddiquee, Z. Zhou, N. Tajbakhsh, R. Feng, M. B. Gotway, Y. Bengio, and J. Liang. Learning fixed points in generative adversarial networks: From image-to-image translation to disease detection and localization. *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 5, 7, 8
- [43] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [44] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [45] Z. Yi, H. Zhang, P. Tan, and M. Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2849–2857, 2017. 1, 2
- [46] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5907–5915, 2017. 2
- [47] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 5, 6, 7, 8
- [48] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 465–476, 2017. 2