


# Re:Draw - Context Aware Translation as a Controllable Method for Artistic Production

João Libório Cardoso  
TU Wien  
Austria

Francesco Banterle  
CNR-ISTI  
Italy

Paolo Cignoni  
CNR-ISTI  
Italy

Michael Wimmer  
TU Wien  
Austria



**Figure 1: Teaser.** As shown, the proposed context-aware translation is capable of automatically redrawing parts of images according to any provided design, without the need for fine-tuning. Unlike image-to-image translation [32], which neglects surrounding context, our approach considers the entire frame. Unlike inpainting [28], which lacks artistic control by ignoring the original content, our method honors the artist’s input. This facilitates the production of more consistent artwork and allows for more complex design choices.

## Abstract

We introduce context-aware translation, a novel method that combines the benefits of inpainting and image-to-image translation, respecting simultaneously the original input and contextual relevance – where existing methods fall short. By doing so, our method opens new avenues for the controllable use of AI within artistic creation, from animation to digital art.


As an use case, we apply our method to redraw any hand-drawn animated character eyes based on any design specifications – eyes serve as a focal point that captures viewer attention and conveys a range of emotions, however, the labor-intensive nature of traditional animation often leads to compromises in the complexity and consistency of eye design. Furthermore, we remove the need for production data for training and introduce a new character recognition method that surpasses existing work by not requiring fine-tuning to specific productions. This proposed use case could help maintain consistency throughout production and unlock bolder and more detailed design choices without the production cost drawbacks. A user study shows context-aware translation is preferred over existing work 95.16% of the time.

## CCS Concepts

• Applied computing → Media arts; • Computing methodologies → Neural networks; Image processing.

## 1 Introduction


Alfred Yarbus’s influential work [58] quantified a long-standing intuition: using an eye tracker, he noted that observers spend a surprisingly large fraction of time fixated on the eyes in a picture. The eyes of others are important to humans because they convey subtle information about a person’s mental state (e.g., attention, intention, emotion) and physical state (e.g., age, health, fatigue).



**Figure 2: Industry Survey.** Average reported time spent in the different stages of production of a cut, in percentage.

This significance has translated into the realm of hand-drawn animation, where the eye designs have often become increasingly complex and expressive to capture these nuances. However, this complexity comes at a cost. Despite its massive resurgence in the last decade [34], traditional animation has struggled to benefit from advances in computer graphics: techniques used in production remain largely the same, with productions relying on repetitive manual labor from a large workforce. As a result, the eyes, being the most time-consuming and intricate to draw, are often the first elements to be simplified, leading to compromises in both expression and artistic consistency. Our aim is to introduce a computational method that can alleviate some of these challenges without sacrificing the artistic integrity of the medium.

We further conducted a survey among 17 professional animators, of which 29.4% (5/17) work at established studios, and the rest either



**Figure 3: Industry Survey.** Reported ranking of common elements in animation from most to least time consuming, in percentage.

freelanced or worked for smaller studios. We asked them multiple questions about time consumption, (answering optional), of which the full breakdown is shown in Figs. 2, 3. Character faces were reported to be the most complex part of animation, with 50% (8/16) reporting it as the element they spend the most time on. Of the remaining animators, 75% (6/8) voted for either anatomy or hair. Drawing was estimated to constitute the vast majority of the work (84.8%), and doing it at the highest level of detail was estimated to take 1.7 times the amount of work than on average, for a total of 66m of additional human effort per key frame from 1st key through coloring. Sadly, the fact that 52.9% (9/17) still use paper drawings in their studios, despite 100% preferring to draw digitally, indicates that using computational tools during the early drawing stages might not be possible yet in practical terms.

### 1.1 Problem

Existing deep learning methods present significant limitations within artistic applications. Inpainting, while capable of generating detailed art that fits within existing content, offers little control over the generated content, making it unsuitable for most precise artistic endeavors [1]. Image-to-image translation, while being able to take artistic input, is constrained by only being applicable to entire images, as it does not take into account the context surrounding target regions.

We propose context-aware translation as the solution to these limitations. We then apply it in a novel pipeline that automates increasing consistency and amount of detail in the eyes of hand-drawn animation characters. It effectively mimics the work of cleanup animators, who redraw frames to fix mistakes and better match the character color guides – despite the misleading name, color guides, also known as model sheets, depict all the information an artist would need to draw a character while remaining true to its intended design and the art style of the production (see Fig. 4 for an example). We also tackle an additional problem this use-case raises: the lack of training datasets of anime production, which we address by proposing methods to negate the need for production data entirely, including a novel character recognition method.

In summary, these are our key contributions:

- (1) Context-aware translation, a novel general deep-learning method that avoids the limitations of both inpainting and image-to-image translation.



**Figure 4: Color Guide.** Also known as a model sheet, this type of document depicts all the information an artist is expected to need to know how to draw a character in accordance with the production direction and style. Example from *Aquarion Evol* [24]

- (a) A dual discriminator structure and novel adversarial losses that enforce simultaneous respect for input content, translation requirements, and context constraints.
- (b) A triple-reconstruction loss that yields greater generation capabilities than traditional loss.
- (2) A character design recognition network that outperforms existing work by using a production-style-aware latent space.
- (3) A novel pipeline that takes advantage of the aforementioned contributions to automatically increase the consistency and amount of detail in the eye region of characters, and without the need of production data during training.

Furthermore, we present an ablation study in Section 4 that scrutinizes the benefits of each of our novel components, contribution of each loss used, compares both our context-aware translation and style-aware clustering against existing work, and assesses the robustness and temporal coherence of our method. We also present a user study with 63 participants in Section 5 that tests three key properties: the absence of detectable artifacts, the enhancement of artwork detail, and the overall aesthetic preference when compared to existing methods, all of which our pipeline successfully validated.

## 2 Related Work

In this Section we describe the minimum animation production background necessary to frame our use-case, and analyze relevant deep-learning existing work.

### Anime Production Background

The production of limited animation is complex and requires the combined efforts of many professional technicians and artists. This results in the need for a pipeline with precisely defined steps, which can be roughly divided into drawing, finishing, and compositing [12]. In the first step, frame sequences are planned out and frames are drawn, going through multiple revisions and artists. As per our



**Figure 5: Proposed Use-Case.** Context-aware translation can be used to automatically redraw the eyes of any character according to a provided color-guide, without fine-tuning. These characters (from *RWBY: Ice Queendom* [41]) were not part of the training set.

survey in Section 1, paper is still prevalent. Finishing involves cleaning and coloring. Each frame is redrawn and in-between frames are added, both in a strict digital format that contains info for the coloring team. Starting with coloring, all operations are digital. Finally, compositing assembles the drawn animations with other elements, such as background or digital effects.

Thus, for a computational method to be usable in this setting, it must: 1) fit within the existing pipeline, preferably after cleaning; 2) not require manual input for each frame in a sequence, defeating its purpose; 3) have a high level of artistic control.

### Editing Content with Style

The seminal work by Gatys et al.[13] showed how to edit an image by varying its overall style by transferring the style of a target image using deep learning optimizations. This work has generated a prolific field [16, 20–23].

Compared to early image-to-image translation methods[18, 64], Liu et al.[32] introduced FUNIT a few-shot unsupervised image-to-image translation framework based on MUNIT [17]. This approach has the ability to translate from an unseen domain with unpaired data. FUNIT achieves this by coupling an adversarial training scheme with a novel multi-task adversarial discriminator. This work has already stemmed variants and improvements [27, 30, 38, 45], but maintain the main drawback of previously mentioned works to modify globally the entire image without guaranteeing that characters in the image will maintain the same pose or expression and other elements won’t be modified. In our work, we want to edit/modify the drawing of a character, but only the design should be varied locally to certain parts; e.g., the eyes. A crucial aspect is to maintain the original pose, facial expressions (e.g., eyes’ looking direction), and non-targeted local elements of the original image.

### Artistic Methods

Deep learning techniques have started to be applied to comics [3] and illustration editing; most literature focused on the colorization of either sketches or shaded manga drawings and main lines extraction. Regarding this latter topic, Simo-Serra et al.[53] proposed one of the first deep learning methods based on a simple encoder-decoder CNN architecture for sketches. Li et al.[31] generalized

it for patterns in drawings using a U-Net with skip connections. Lines extraction methods can be improved when they are paired with user inputs [52] or adversarial training [51] or unpaired data with the synthesis of paired ones [29].

Concerning colorization, Yuanzheng et al.[7] employed a conditional GAN to color illustration line art using scribbles from the user. This approach was improved by Zhang et al.[63] who proposed a two-stage sketch colorization for illustration: First the user splashes colors on a sketch, then, the system generates a draft using a GAN that the user can correct mistakes with edits propagated using a refinement GAN. The illustration dataset by Zhang et al.[2] was used for training, and it has also become a staple for many other methods. In a different approach for manga colorization was proposed in [48, 50], where the user, instead of scribbling or splashing colors, provides an input image with basic colors. In particular, Shimizu et al.[48] showed that providing a flat colored image of the sketch image can generate high-quality colorization with few training data. Note that this flat filling colorization can also be automated using deep learning [62] and user inputs for complex line art. In this domain, researchers have also focused on specific parts to colorize using user inputs [1], or combined with text tags [26]. While at first, this might appear appropriate for animation, as they do not require input per frame, they suffer from limitations when it comes to artistic control. Both remove control of lighting conditions, and the first removes control over pose and expression by virtue of being an inpainting method, while the second entirely removes control over shading. Akinobu et al.[33] tackle anime colorization using a few-shot strategy and an ad-hoc sampling method for patches tailored for anime.

An emerging topic is manga generation using adversarial training and some character parameters [19], photos [54, 56], and sketches [60]. Although they can generate high-quality results, they lack precise artistic control.

Recently researchers have applied classic topics such as segmentation [61] and clustering [37] to illustration and cartoon art. Nir et al.[37] proposed a method to learn a style-specific semantic representation for animated content using self-supervision. However, it needs to be trained separately on each production and is more suited for tracking characters and not their designs. In our work,

we introduce a more flexible semantic clustering that decouples the style of the anime from the content. The content is clustered in an art-style normalized Euclidean space, where distances between these portraits correspond to a measure of character design similarity in a production.

### 3 Method

Our approach for enhancing animated eyes takes as input the animation frames to be improved and a character color guide; see Fig. 6. We use an unsupervised convolutional network trained alongside classification networks, capable of telling designs apart, as its adversaries. Using such a model requires artists to manually associate regions to redraw and color guides manually, which is not practical. Even more problematic, to train this type of adversarial structure, one normally uses pairs of these images, labeled into different classes (character designs). In particular, to ensure our model is capable of generalizing to new designs, we need to train on a large enough variety of them. Yet, art direction is not easily available and generally not created in high enough quantities that would be needed for a robust training. Moreover, manually tagging and cropping this data would be extremely labor intensive and hard to replicate. To address these issues we propose a novel character design clustering method, and use it to automatically infer training data from random frames, thus solving the association problem. As such, Re:Draw does not require internal production data for training, instead it only requires a set of random sampled frames from different productions; see Fig. 9.

Image in-painting has shown to be capable of completing missing regions, yet predictions based only on the surrounding of the area to be redrawn do not allow artists to finely control the output results using art or style direction examples. Image-to-image translation and style transfer are capable of using both of these inputs, yet existing work is incapable of generating art that fits and correctly matches within the actual context of the drawing; they can be very unreliable in preserving the artwork pose. For these reasons, we introduce context-aware translation. We make use of two adversarial discriminators built using partial convolutions, allowing them to weight images differently and independently, and a novel triple reconstruction loss based on the concept of the generation of image triplets.

#### 3.1 Dataset Generation

We will now describe how we avoid the need for production data, by automatically clustering art by character design and then further splitting it into low- and high-levels of detail. This categorized dataset is required during the training phase of our context-aware translation model, which involves solving multiple adversarial classification tasks simultaneously.

##### Object Detection

We first train an object-detection network – we use the well-established Faster R-CNN network [44] – to identify character faces and details in them (such as eyes) and run it on the randomly sampled frames. This results in a dataset of character faces in a variety of poses, split by the sources they were sampled from.

##### Style-Aware Clustering

Although re-identification of human faces is a long studied topic [4], we found existing work to be ineffective at automatically identifying animated characters not seen during training. We attribute this to the fact that, while human faces have a consistent and predictable structure, animated characters are not restrained by the laws of reality and thus present a much higher variance: in a given production, characters might have a very similar look and feel, while in another production they might vary widely in structure and shape (see Fig. 7).

To solve this issue, we improve upon the state of the art of character recognition by combining ideas from facial recognition and image-to-image translation. We propose a supervised network that, unlike existing work, maps character portraits to an art-style normalized Euclidean space, where distances between these portraits correspond to a measure of character design similarity within its production. It takes as input character portraits to be mapped and a collection of random portraits from the same production for normalization estimation.

As shown in Fig. 8, latent representations of both inputs are estimated: we compute the content representation using a ResNet [14] encoder – which is well established for object recognition – and the production representation using a convolutional encoder. The latter is done using only the lightness in  $l\alpha\beta$  color space [43], as we found that the normalization input works better if it only contains the main shape information, so we use a color space to decorrelate it from color variation. This style-latent representation is then used to compute a set of affine transformations, with the goal of mapping the encoding from an absolute Euclidean-space representation of portraits to the style-normalized one. This mapping is done using Adaptive Instance normalization [16] on the content-input latent representation. This finally results in 32 parameters per portrait thanks to the linear layers, which are then clustered using traditional hierarchical clustering, with unweighted pair group method, arithmetic mean and Euclidean distance. These methods and parameters were chosen by testing the rate of correct clustering across a validation dataset.

To train this network  $E$  and ensure the content-encoding output respects the desired intra and inter-class properties of the normalized Euclidean space, we use the option of Triplet Margin Loss [5, 15] – that is, given a pair of portraits from the same design  $\{P_1, P_2\}$  and one from another  $P_3$  but from the same production  $\mathbb{P}$ , we minimize the distance from the first two, while maximizing the distance of the third:

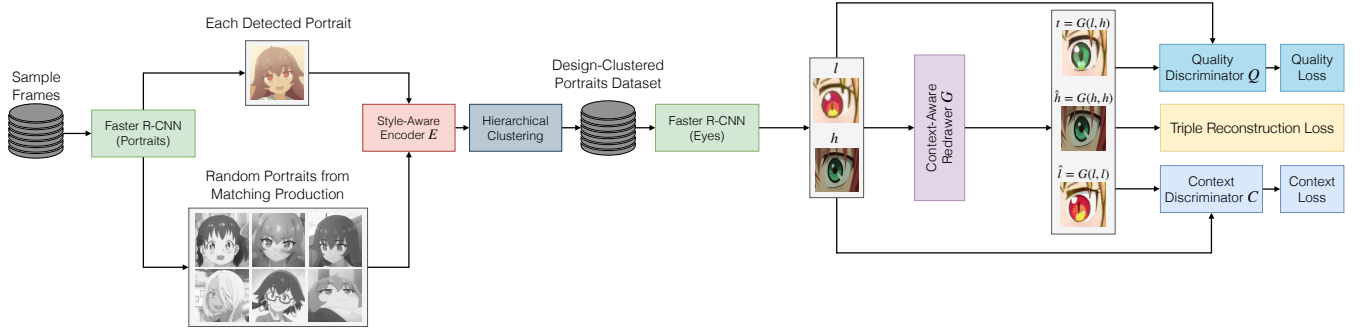
$$\operatorname{argmin}_E \max \left\{ \|E(P_1, \mathbb{P}) - E(P_2, \mathbb{P})\|^2 - \|E(P_1, \mathbb{P}) - E(P_3, \mathbb{P})\|^2 + 1, 0 \right\} \quad (1)$$

This means that, during training, character portraits must be provided in sets of three. We also ensure that the total training weight of each production style and of each character design within each style is the same, to further help with generalization. While it is technically possible to train  $E$  in conjunction with the context-aware translation, it is more computationally efficient to train  $E$  first and then freeze it to train the remainder networks.

##### Level-of-Detail Split

Having the portraits organized by character design, we extract





**Figure 9: Context-Aware Training.** Detected portraits are clustered using our style-aware encoder (left), with additional portraits as style guides. Content and style images are extracted from portraits of different designs (center). These are used to generate 3 redrawings from different combinations of these input pairs, which are judged by multiple losses, including two multi-class discriminators (right).


and high-detail images are ignored as well, allowing us to create a reconstruction loss in low-detail images.

### Adversarial Discriminators


We address our aforementioned conflicting goals by using two independent image multi-task classifiers instead: a *quality discriminator*  $Q$  judging whether the output is high detail and matches the intended design, and a *context discriminator*  $C$  judging whether it fits within the original artwork and its own design, irrespective of detail-level.

To achieve this purpose, as we show in Section 4, we need to train each discriminator differently, despite sharing many commonalities: both have the same partial convolutional structure and are trained using hinge loss with R1 regularization [35], to prevent over-fitting and mode collapse. This results in the following losses for penalizing wrong classifications of positive  $\mathcal{L}_P$  and negative  $\mathcal{L}_N$  examples:

$$\begin{aligned} \mathcal{L}_P(x, \mathbb{S}) &= [\max(0, 1 - D(x)_{\mathbb{S}}) + \gamma \|\nabla D(x)_{\mathbb{S}}\|^2]_1 \\ \mathcal{L}_N(x, \mathbb{S}) &= [\max(0, 1 + D(x)_{\mathbb{S}})]_1 \end{aligned} \quad (3)$$



**Figure 10: Low-Pass Filter.**  $\hat{l}$  is computed during training from enhancing  $l$ . Right shows the result of our low pass filter.



### Post-Processing

To further ensure image regions generated fit within the original image, we apply a few image processing operations: after re-sampling the network output to the original resolution, we apply color transfer [43] and place it into the original image using Poisson image editing [42].

## 4 Ablation Experiments

### Clustering

We compared our style-aware clustering approach with FaceNet [47], trained on the same labeled animated character faces. We statistically analyzed how effective the latent representations learned from either method work on a validation dataset of productions styles not seen during training: we measured the ratio of the average squared norm distance between each point of the same character, and the average squared norm distance between the mean points of each character. The lower this value, the better is the latent space representation in principle. Our method measures a ratio of  $1.212e^{-4}$  in the validation data, which outperforms FaceNet’s  $1.494e^{-4}$  ratio. Interestingly, our method performs similarly to a FaceNet network trained on the validation dataset, whose ratio was  $1.209e^{-4}$ . This shows our method presents better generalization to unseen data. Furthermore, we show in Fig. 12 how well split the character faces from validation production styles are. The colors of the points represent their ground-truth labels, and there is a visible improvement with our method.

### Redrawing

We evaluated our redrawing approach against neural cross-domain image translation with content and style inputs. While multiple variations of these networks exist [40, 45], they mostly compete in translation ability and do not address our problems. Thus, we chose as a baseline for comparison Liu et al.[32] method (FUNIT). We progressively introduced each of our novelties to show the importance of each one (Fig 13): we introduced our style-aware clustering by training FUNIT only on the few manually labeled faces versus the generated dataset. We then added our double-discriminator method, while maintaining FUNIT’s traditional reconstruction loss. Finally, we introduced the proposed triplet reconstruction, followed by the post-processing step. FUNIT fails to respect pose/expression and context. Without our large-scale dataset, it often fails to generate realistic art and cannot generalize to designs not seen during training. Our dataset and double discriminator approach solve all of these issues, but reduce the ability of the network to generate highly detailed art. Our novel triplet reconstruction fixes that.

Our method is also very stable and demonstrates temporal coherence, shown in Fig. 20, despite not specifically incorporating it in the loss. We believe the reason for this consistency is its emphasis on preserving the intended pose and context of the drawing: as shown in Fig. 14, removing the regularizer from the adversarial losses results in mode-collapse, and removing the quality loss results in a poor auto-encoder as expected. Yet, removing the reconstruction loss does not result in unpredictable output as expected, just an inability to learn useful transformations, and removing the context loss does not result in output reminiscent of FUNIT. Our explanation is that the network, by following the context loss, is

being incentivized to exhibit spatial and temporal consistency in its output. This stability is what enables our less explicit form of reconstruction loss to direct the training toward useful results.

We critically evaluate the performance of existing techniques – namely style-transfer, image-to-image translation, and text-to-image diffusion—against our proposed context-aware translation. Figs. 1 and 15 clearly illustrate the limitations of these existing approaches, thereby reinforcing why our method is a more apt solution for this particular type of problem. Furthermore, Figs. 5, 17, 18 and 19 provide compelling evidence of our method’s robustness and versatility in handling a variety of challenging scenarios, further substantiating its suitability for this application.

The limitations of the network we found boiled down to two cases: uncommon occlusions and strong rotations. As most occlusion to anime eyes is hair and the vast majority of are drawn up-right, the network can generate artifacts outside of those expected conditions. As shown in Fig. 16, the shape of occluders below the eyeline might be distorted as if was a skin tattoo. Eyes will be drawn upright if a character is reversed. Adding a network capable of estimating eye rotation to the pipeline would automate this process and further improve the generated dataset. But outside of these two unusual spaces, artifacts we did find were created by Poisson-blending, not our models, which leads us to conclude post-processing is the current main limitation of the method.

## 5 User Study

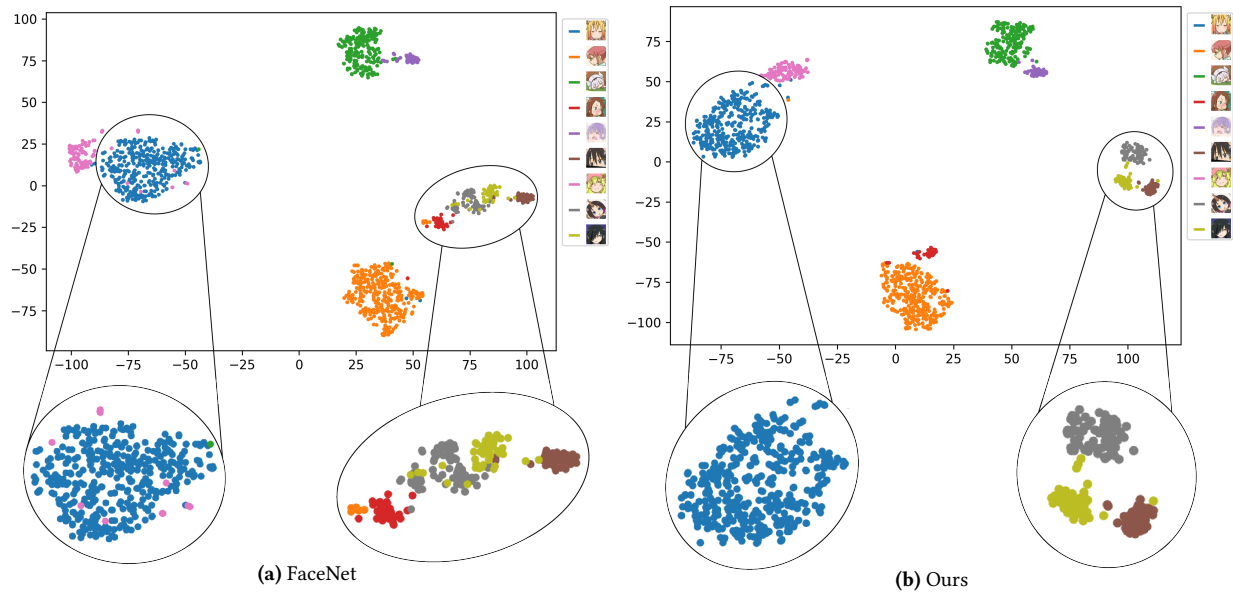
To validate our work, we conducted a user study with 63 participants and three different tasks with images coming from an anime production. Participants applied voluntarily to the study and the majority of them classified themselves as anime experts. Three visual tasks were presented:

### Realness (T1)

The user had to watch eight images and pick the ones with drawing problems. In this test, four were original un-retouched images (*Market, Tent, Cold Night, Field of Leaves*), and four had the eyes replaced by our approach (*Dormitory, Bar, Bathroom, Crying*). Images were presented in random order, and the user did not know how many could have problems. This study aimed to determine if our method generates visual artifacts or if its result is not distinguishable from a production. To analyze these results, we used the  $\chi^2$  test with Yates correction [49] and null hypothesis  $H_0$ : “Original and retouched images are equally probable to be considered as wrong”. The number of times users decided that an image was wrong has a very similar distribution for images generated by our approach and original ones (see Tabs. 3 and 4).  $\chi^2(1, N=504)=0.0405, p=0.840$ , therefore, with a  $p$ value larger than 0.05, we can conclude independence between groups (*Our vs Real*) and categories (*Wrong vs Right*), or in simpler words, our generated eyes do not have evident artifacts and are indistinguishable from real production.

### Level of Detail (T2)

Participants were asked to choose between two images: one unedited (U) and one with eyes enhanced by our approach (O), focusing on which image displayed more detailed eyes. Each user evaluated




**Figure 12: Clustering Validation** 2D visualization of characters faces from a production not seen during training [8], generated using T-Distributed Stochastic Neighbor Embedding [55] on the latent spaces learned using FaceNet [47] and our proposed encoder network. Colors correspond to ground truth labels. As shown, character differentiation is clearer in our learned space.

eight pairs of images. This study aimed to determine if we effectively produce images with higher detail than the original production. To understand if the expressed preferences were statistically significant, we employed the formula for the multiple comparison test [10]: if difference between the score of one method against another exceeded a critical value  $R$ , the results were deemed significant.  $R$  was calculated according to specified alpha significance level of 0.05. We also tested if the agreement among observers  $u \in [-1, 1]$  (where  $u = -1$  strong disagreement and  $u = 1$  strong agreement) with the following null hypothesis  $H_0$ : “There is no agreement among participants on one category being better”. The failure of this hypothesis implies that categories are perceptually equivalent causing difficulties in judging (e.g. we check if users allotted their preferences at random). The overall result (Table 5) is that our method effectively produces images with higher detail than the original production artwork with statistical significance.

### Preference (T3)

The user had to select the better looking image from a pair. For each of the 8 scenes, there were three pairs cross-comparing 3 methods: FUNIT[32] (F), FUNIT with our clustering (FC), and our method (O); in total, 24 pairs for each user. This study aims to determine if our method is more effective in producing results better looking than F and FC. In short, our method was preferred to (F) 95.16% of times. In Tab. 6 we used the same analysis of T2, showing that images generated by our method are more preferred than the ones generated by the previous work (F), even when it is enhanced with our dataset (FC).

**Table 2: Domain Knowledge.** More than half of the users classified themselves on the high end of an anime (viewer) knowledge spectrum.



## 6 Conclusions


We have presented context-aware translation: a novel unsupervised image-to-image network, trained with two adversarial classification networks. We built these classifiers using partial convolutions, allowing them to weight generated images differently and independently. We introduced novel loss functions for these discriminators and a novel reconstruction loss based on image triplets to achieve context-aware translation. We proposed a deep-learning approach with a novel character design clustering to automatically collect training data from animation frames and input data during inference.

We have shown the method is capable of automatically redrawing the eyes of an anime character according to a provided character design direction. In this way, the shape and color of the iris can be changed, features, like reflexes and shades, can be added and the level of detail can be increased under the precise artist’s control. We







**Figure 13: Ablation of Contributions.** We compared our method with traditional image translation by progressively introducing our contributions (rows 4 to 6). The synthetic data extrapolated using our style-aware clustering prevents over-fitting (a,e) and allows generalization to unseen designs (g,h,i). Further introducing our dual discriminators allows redrawn areas to maintain artwork fit, but sacrifices detail (a,b,c,e,i) and ability to redesign shape or color (g,h). Finally, introducing the triplet reconstruction loss brings that expressiveness back.



**Figure 14: Ablation of Losses.** We tested removing each loss during training of redrawer  $G$  on examples from Fig. 13. It shows  $\mathcal{L}_Q$  and  $\mathcal{L}_R$  are crucial for useful results. Most notably, it reveals  $\mathcal{L}_C$  makes output much more predictable than otherwise expected, and this stability is what enables the use of our less explicit  $\mathcal{L}_R$ .




**Figure 15: Ablation of Related Work.** We compared our method with style transfer and image-to-image translation with and without post-processing.



**Figure 16: Redrawing Limitations.** Our method vastly improves upon existing work on handling occlusions; however, it still can struggle with uncommon scenarios (top). As it was trained mostly with upright eyes, fails if input is not pre-rotated to upright position (bottom).

**Table 3: Realness Test (T1):** The number of times that our redrawn art frames ■ and original unmodified ones ■ have been picked as fake is not significantly different.



**Table 4: Realness Test (T1).** We tested the independence between groups (Our vs Real) and categories (Wrong vs Right) and our generated eyes were indistinguishable from real artwork.

	Wrong	Right	Marginal Row Totals
Ours	66	186	252
Real	69	183	252

have made the method easily replicable, by removing the need for internal production data or labeled data. Given the general nature of our method, we expect it to be usable or extendable to other elements in and outside of animation. Only our frequency threshold value and the quality-split criteria could be specific to our use case of animation, but such is the case of many meta parameters in deep learning methods.

The obtained results also indicate the model might be presenting emergent behavior of precise image segmentation, despite never having seen segmented data, which we plan to explore in future work. As the models run nearly in real-time, Re:Draw could be used interactively as artists draw or color line-art. Only post-processing,

**Table 5: Level-of-Detail Test (T2).** U and O refer to unedited and our generated images, respectively. For this test, the minimum score difference  $R$  to determine statistical difference is  $R = 45$  for the overall comparison and  $R = 16$  for each single image [10], with results not significantly different being circled together. We also performed the  $\chi^2$  test with Yates correction [49], setting a lower bound of  $\chi^2 = 3.84$  for this test.

Image	Agree. Coeff. $u$	$\chi^2$	Sign. $u$	Groups
1	0.230	15.25	OK	(U - 16) (O - 47)
2	0.041	3.57	NO	(U - 24; O - 39)
3	0.119	8.40	OK	(U - 20) (O - 43)
4	0.334	21.73	OK	(U - 13) (O - 50)
5	0.373	24.14	OK	(U - 12) (O - 51)
6	0.297	19.44	OK	(U - 14) (O - 49)
7	0.058	4.59	OK	(U - 23) (O - 40)
8	0.144	9.92	OK	(U - 19) (O - 44)
Overall	0.192	97.79	OK	(U - 141) (O - 363)

**Table 6: Preference Test (T3).** F, FC and O refer to images generated by FUNIT, FUNIT with our clustering and our method, respectively. For this test, the minimum score difference  $R$  to determine statistical difference is  $R = 65$  for the overall comparison and  $R = 24$  for each single image [10]. We also performed the  $\chi^2$  test with Yates correction [49], setting a lower bound of  $\chi^2 = 7.82$  for this test.

Image	Agree. Coeff. $u$	$\chi^2$	Sign. $u$	Groups
1	0.460	73.38	OK	(FC - 32; F - 40) (O - 117)
2	0.808	126.61	OK	(F - 5) (FC - 63) (O - 121)
3	0.856	134.00	OK	(F - 2) (FC - 68) (O - 119)
4	0.830	129.92	OK	(F - 4) (FC - 63) (O - 122)
5	0.396	63.53	OK	(F - 20) (FC - 67) (O - 102)
6	0.307	49.92	OK	(F - 32) (FC - 59) (O - 98)
7	0.526	83.53	OK	(F - 17) (FC - 56) (O - 116)
8	0.413	66.23	OK	(F - 23; FC - 52) (O - 114)
All	0.530	802.93	OK	(F - 143) (FC - 460) (O - 909)

which is significantly more computationally intensive, prevents the entire method from being run interactively. Given blending is also the main culprit behind artifacts, the focus of future work will be to further improve the method, likely within the reconstruction loss, to remove the need for post-processing altogether.

We have substantiated the need for such a style-driven enhancement with a professional user survey that reported the impact in the time of high quality drawing of the face details. Finally, we have validated our approach and results with ablation and a user study showing that our style-normalized latent space pushes the state of the art regarding the identification of non-photorealistic imagery, our approach is preferred over traditional image-to-image translation 95.16% of the time and the images generated are not discernible with respect to images drawn by artists with traditional techniques. Please look at the additional materials for videos showing frame coherence.



**Figure 17: Increasing Detail.** The method behaves as intended when prompted to increase the amount of detail in character eyes, despite challenges such as unusual face proportions, adverse color balance due to lighting effects, head tilts, minor hair occlusions and character designs from productions that have not been seen during training.

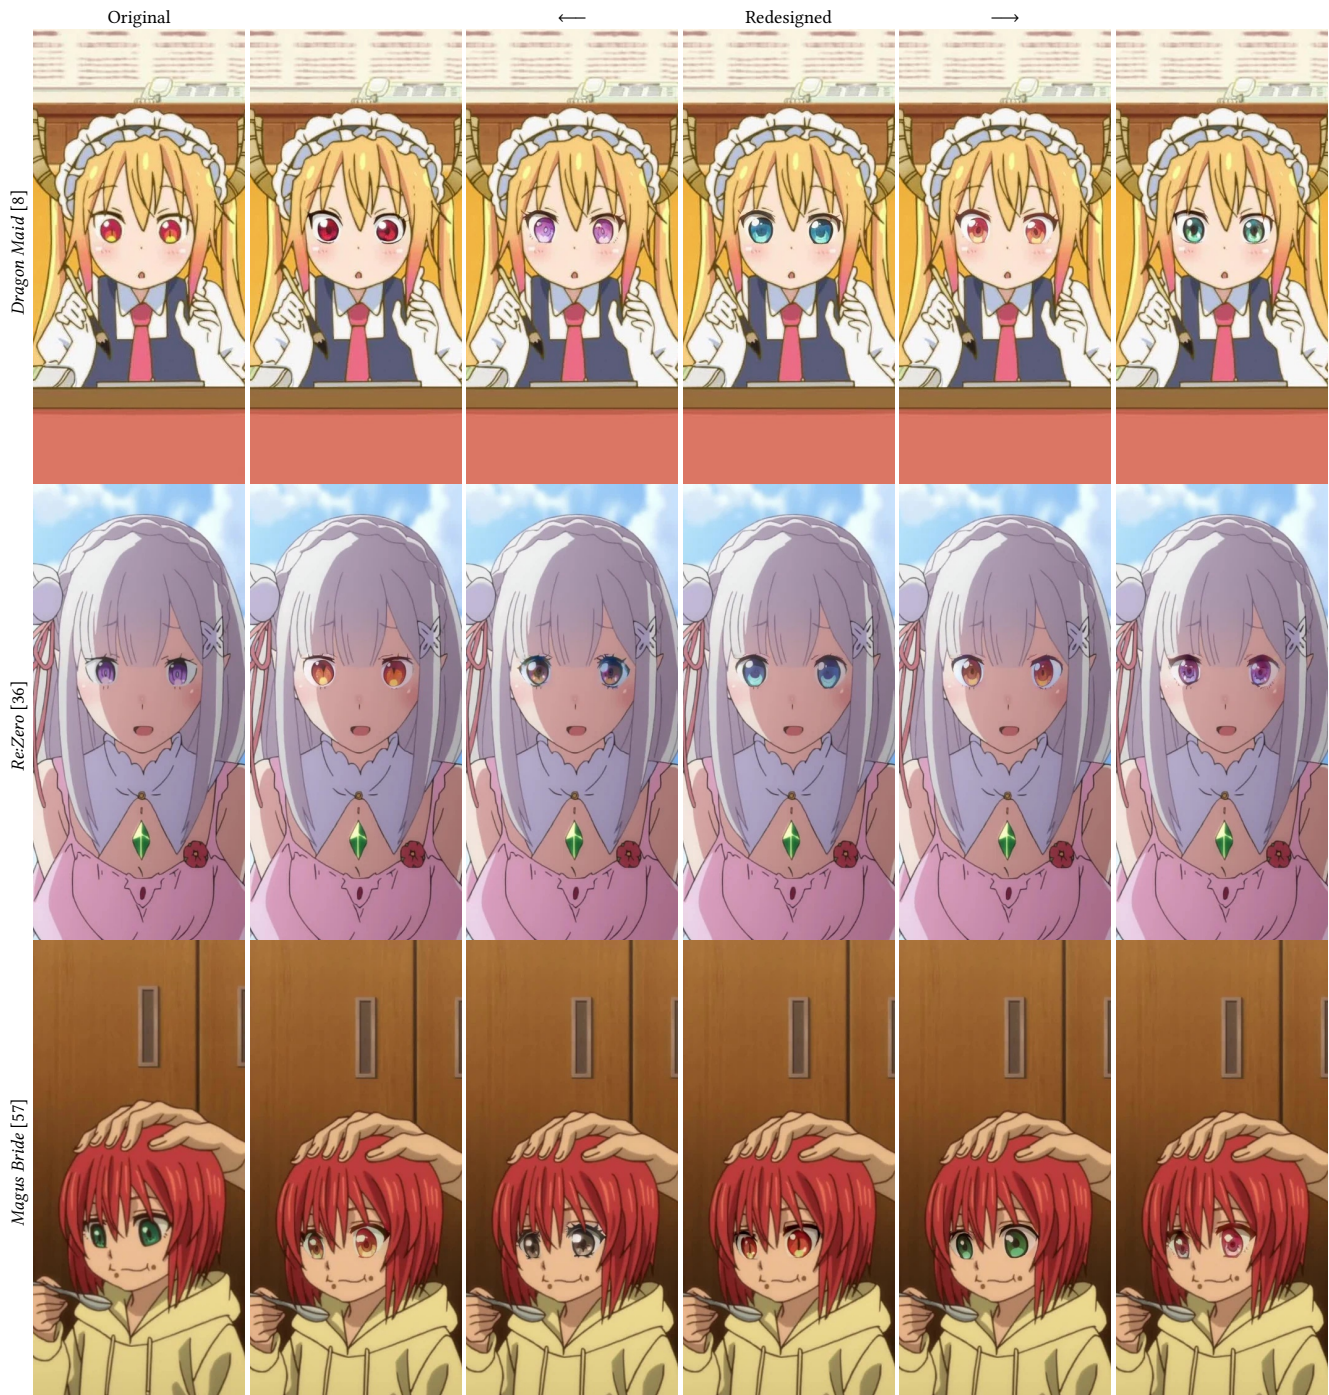
## 7 Acknowledgments

We would like to especially thank Jarret Martin for teaching us about the animation pipeline needs, his continuous feedback throughout the course of our research and help running the industry survey. We would also like to thank Tonari Animation and OtakuVS for offering us the production data and other elements from their works Otachan and Second Self for illustration of the research, which we did whenever possible.

All copyrighted content belongs to its original owners. Characters, music, and all other forms of creative content are owned by their original creators, and we make no claim to them. This content was made in compliance with Fair Use law for the purpose of illustration of non-profit scientific research.

## References

- [1] Kenta Akita, Yuki Morimoto, and Reiji Tsuruno. 2020. Deep-Eyes: Fully Automatic Anime Character Colorization with Painting of Details on Empty Pupils. In *Eurographics 2020 - Short Papers*, Alexander Wilkie and Francesco Banterle (Eds.). The Eurographics Association. <https://doi.org/10.2312/egs.20201023>
- [2] Anonymous, Danbooru community, and Gwern Branwen. 2022. Danbooru2021: A Large-Scale Crowdsourced and Tagged Anime Illustration Dataset. <https://www.gwern.net/Danbooru2021>. <https://www.gwern.net/Danbooru2021> Accessed: DATE.
- [3] Olivier Augereau, Motoi Iwata, and Koichi Kise. 2018. A Survey of Comics Research in Computer Science. *J. Imaging* 4, 7 (2018), 87. <https://doi.org/10.3390/jimaging4070087>
- [4] Stephen Balaban. 2015. Deep learning and face recognition: the state of the art. *Biometric and surveillance technology for human and activity identification XII* 9457 (2015), 68–75.
- [5] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. 2016. Learning local feature descriptors with triplets and shallow convolutional neural networks.. In *Bmvc*, Vol. 1. 3.
- [6] Boichi, Shueisha, and Dr.STONE Production Committee. 2019. Dr.Stone. <https://dr-stone.jp/>. [Online; accessed 4-August-2023].
- [7] Yuanzheng Ci, Xinzhu Ma, Zhihui Wang, Haojie Li, and Zhongxuan Luo. 2018. User-Guided Deep Anime Line Art Colorization with Conditional Adversarial Networks. In *Proceedings of the 26th ACM International Conference on Multimedia (Seoul, Republic of Korea) (MM '18)*. Association for Computing Machinery, New York, NY, USA, 1536–1544. <https://doi.org/10.1145/3240508.3240661>
- [8] Coolkyousinnjya and Dragon Life Improvement Committee. 2017. Miss Kobayashi’s Dragon Maid. <https://maidragon.jp/1st/> [Online; accessed 4-August-2023].
- [9] Darling in the Franxx Production Committee. 2018. Darling in the Franxx. <https://darli-fra.jp/staffcast/>. [Online; accessed 4-August-2023].



**Figure 18: Redesign Results.** A side-effect of how our proposed networks are trained is that our method is also capable of applying entirely different designs to characters. While not the focus of our work, it demonstrates the versatility of our method and the robustness of the learned model, even when applied to high-resolution imagery.



**Figure 19: Redesign Results.** Further examples of character redesign which demonstrate that our method is robust enough to handle characters in a variety of challenging scenarios. This includes dealing with oblique camera angles, irregular lighting conditions, head tilts and occlusions caused by hair.



**Figure 20: Temporal Coherence.** Example of redesigning the eyes of a character from *Re:Zero* [36] in a particularly lengthy shot. Output is temporally consistent.

- [10] H. A. David. 1988. *The Method of Paired Comparisons, 2nd ed.* Oxford University Press.
- [11] Eden of the East Production Committee. 2009. Eden of the East. <https://juiz.jp/>. [Online; accessed 4-August-2023].
- [12] Furansujin Connection. 2016. Les étapes de fabrication. <https://web.archive.org/web/20220401075414/http://www.furansujinconnection.com/les-etapes-de-fabrication/>. [Online; accessed 2022-March-21].
- [13] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2016. Image Style Transfer Using Convolutional Neural Networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. IEEE Computer Society, 2414–2423. <https://doi.org/10.1109/CVPR.2016.265>
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [15] Alexander Hermans, Lucas Beyer, and Bastian Leibe. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017).
- [16] Xun Huang and Serge J. Belongie. 2017. Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*. IEEE Computer Society, 1510–1519. <https://doi.org/10.1109/ICCV.2017.167>
- [17] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. 2018. Multimodal Unsupervised Image-to-image Translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [18] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*. IEEE Computer Society, 5967–5976. <https://doi.org/10.1109/CVPR.2017.632>
- [19] Yanghua Jin, Jiakai Zhang, Minjun Li, Yingtao Tian, Huachun Zhu, and Zhihao Fang. 2017. Towards the Automatic Anime Characters Creation with Generative Adversarial Networks. CoRR abs/1708.05509 (2017). <http://arxiv.org/abs/1708.05509>
- [20] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-Free Generative Adversarial Networks. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.), 852–863. <https://proceedings.neurips.cc/paper/2021/hash/076ccd93ad68be51f23707988e934906-Abstract.html>
- [21] Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 4401–4410. <https://doi.org/10.1109/CVPR.2019.00453>
- [22] Tero Karras, Samuli Laine, and Timo Aila. 2021. A Style-Based Generator Architecture for Generative Adversarial Networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 12 (2021), 4217–4228. <https://doi.org/10.1109/TPAMI.2020.2970919>
- [23] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 8107–8116. <https://doi.org/10.1109/CVPR42600.2020.00813>
- [24] Shoji Kawamori, SATELIGHT, and Project AQUARION Evol. 2011. Aquarion Evol. [http://aqevol.com/staff\\_cast/index.php](http://aqevol.com/staff_cast/index.php). [Online; accessed 4-August-2023].
- [25] Key and Kyoto Animation. 2007. Clannad. <https://www.tbs.co.jp/clannad/clannad1/02staffcast/staffcast.html>. [Online; accessed 4-August-2023].
- [26] Hyunsu Kim, Ho Young Jho, Eunhyeok Park, and Sungjoo Yoo. 2019. Tag2Pix: Line Art Colorization Using Text Tag With SECat and Changing Loss. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. IEEE, 9055–9064. <https://doi.org/10.1109/ICCV.2019.00915>
- [27] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwang Hee Lee. 2020. U-GAT-IT: Unsupervised Generative Attentional Networks with Adaptive Layer-Instance Normalization for Image-to-Image Translation. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BJLZ5ySKPH>
- [28] LahInTheFutureland. 2023. Mysteriousv4 SDXL 1.0. <https://civitai.com/models/118441/lah-mysterious-or-sdxl>. [Online; accessed 19-October-2023].
- [29] Gayoung Lee, Dohyun Kim, Youngjoon Yoo, Dongyoon Han, Jung-Woo Ha, and Jaehyuk Chang. 2019. Unpaired Sketch-to-Line Translation via Synthesis of Sketches. In *SIGGRAPH Asia 2019 Technical Briefs* (Brisbane, QLD, Australia) (SA ’19). Association for Computing Machinery, New York, NY, USA, 45–48. <https://doi.org/10.1145/3355088.3365163>
- [30] Bing Li, Yuanlue Zhu, Yitong Wang, Chia-Wen Lin, Bernard Ghanem, and Linlin Shen. 2021. AniGAN: Style-Guided Generative Adversarial Networks for Unsupervised Anime Face Generation. *IEEE Transactions on Multimedia* (2021), 1–1. <https://doi.org/10.1109/TMM.2021.3113786>
- [31] Chengze Li, Xueting Liu, and Tien-Tsin Wong. 2017. Deep Extraction of Manga Structural Lines. *ACM Trans. Graph.* 36, 4, Article 117 (7 2017), 12 pages. <https://doi.org/10.1145/3072959.3073675>
- [32] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. 2019. Few-shot unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10551–10560.
- [33] Akinobu Maejima, Hiroyuki Kubo, Seitaro Shinagawa, Takuya Funatomi, Tatsuo Yotsukura, Satoshi Nakamura, and Yasuhiro Mukaigawa. 2021. *Anime Colorization Using Few-Shot Learning*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3478512.3488604>
- [34] Hiromichi Masuda, Tadashi Sudo, Kazuo Rikukawa, Yuji Mori, Naofumi Ito, Yasuo Kameyama, and Megumi Nonouchi. 2019. *Anime Industry Report*. Technical Report. The Association of Japanese Animations.
- [35] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. 2018. Which training methods for GANs do actually converge?. In *International conference on machine learning*. PMLR, 3481–3490.
- [36] Tappei Nagatsuki, KADOKAWA, and Re: Life in a Different World Starting from Zero 1 Production Committee. 2016. Re: Life in a Different World Starting from Zero. <http://re-zero-anime.jp/tv/>. [Online; accessed 4-August-2023].
- [37] Oron Nir, Gal Rapoport, and Ariel Shamir. 2022. CAST: Character labeling in Animation using Self-supervision by Tracking. *Computer Graphics Forum* (2022). <https://doi.org/10.1111/cgf.14464>
- [38] Ori Nizan and Ayellet Tal. 2020. Breaking the Cycle - Colleagues Are All You Need. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 7857–7866. <https://doi.org/10.1109/CVPR42600.2020.00788>
- [39] Eiichiro Oda and Toei Animation. 1997. One Piece. <https://one-piece.com/anime/>. [Online; accessed 4-August-2023].
- [40] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. 2021. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10743–10752.
- [41] Monyreak Oum and Team RWBY Project. 2022. RWBY: Ice Queendom. <https://anime.team-rwby-project.jp/staff/>. [Online; accessed 4-August-2023].
- [42] Patrick Pérez, Michel Gangnet, and Andrew Blake. 2003. Poisson image editing. In *ACM SIGGRAPH 2003 Papers*. 313–318.
- [43] Erik Reinhard, Michael Adhikmin, Bruce Gooch, and Peter Shirley. 2001. Color transfer between images. *IEEE Computer graphics and applications* 21, 5 (2001), 34–41.
- [44] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1 (Montreal, Canada) (NIPS’15)*. MIT Press, Cambridge, MA, USA, 91–99.
- [45] Kuniaki Saito, Kate Saenko, and Ming-Yu Liu. 2020. COCO-FUNIT: Few-Shot Unsupervised Image Translation with a Content Conditioned Style Encoder. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 12348)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 382–398. [https://doi.org/10.1007/978-3-030-58580-8\\_23](https://doi.org/10.1007/978-3-030-58580-8_23)
- [46] Makoto Sanada, Kadokawa, and Angels of Death Production Committee. 2018. Angels of Death. [http://www.jcstaff.co.jp/sakuhin/nenpyo/2018/06\\_satsuriku/satsuriku.htm](http://www.jcstaff.co.jp/sakuhin/nenpyo/2018/06_satsuriku/satsuriku.htm). [Online; accessed 4-August-2023].
- [47] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 815–823.
- [48] Yugo Shimizu, Ryosuke Furuta, Delong Ouyang, Yukinobu Taniguchi, Ryota Hinami, and Shonosuke Ishiwatari. 2021. Painting Style-Aware Manga Colorization Based On Generative Adversarial Networks. In *2021 IEEE International Conference on Image Processing, ICIP 2021, Anchorage, AK, USA, September 19-22, 2021*. IEEE, 1739–1743. <https://doi.org/10.1109/ICIP42928.2021.9506254>
- [49] Sidney Siegel and N. John Castellan. 1988. *Nonparametric Statistics for the Behavioral Sciences*. McGrall-Hill International.
- [50] Felipe Coelho Silva, Paulo André Lima de Castro, Hélio Ricardo Júnior, and Ernesto Cordeiro Marujo. 2019. Mangan: Assisting Colorization Of Manga Characters Concept Art Using Conditional GAN. In *2019 IEEE International Conference on Image Processing, ICIP 2019, Taipei, Taiwan, September 22-25, 2019*. IEEE, 3257–3261. <https://doi.org/10.1109/ICIP.2019.8803667>
- [51] Edgar Simo-Serra, Satoshi Iizuka, and Hiroshi Ishikawa. 2018. Mastering Sketching: Adversarial Augmentation for Structured Prediction. *ACM Trans. Graph.* 37, 1, Article 11 (1 2018), 13 pages. <https://doi.org/10.1145/3132703>
- [52] Edgar Simo-Serra, Satoshi Iizuka, and Hiroshi Ishikawa. 2018. Real-Time Data-Driven Interactive Rough Sketch Inking. *ACM Trans. Graph.* 37, 4, Article 98 (7 2018), 14 pages. <https://doi.org/10.1145/3197517.3201370>
- [53] Edgar Simo-Serra, Satoshi Iizuka, Kazuma Sasaki, and Hiroshi Ishikawa. 2016. Learning to Simplify: Fully Convolutional Networks for Rough Sketch Cleanup. *ACM Trans. Graph.* 35, 4, Article 121 (7 2016), 11 pages. <https://doi.org/10.1145/>

2897824.2925972

- [54] Hao Su, Jianwei Niu, Xuefeng Liu, Qingfeng Li, Jiahe Cui, and Ji Wan. 2021. MangaGAN: Unpaired Photo-to-Manga Translation Based on The Methodology of Manga Drawing. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*. AAAI Press, 2611–2619. <https://ojs.aaai.org/index.php/AAAI/article/view/16364>
- [55] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [56] Xinrui Wang and Jinze Yu. 2020. Learning to Cartoonize Using White-Box Cartoon Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [57] Kore Yamazaki, Mag Garden, and The Ancient Magus' Bride Production Committee. 2017. The Ancient Magus Bride. <https://mahoyome.jp/staffcast/>. [Online; accessed 4-August-2023].
- [58] Alfred L. Yarbus. 1967. *Eye Movements and Vision*. Plenum Press, New York. Translated from Russian.
- [59] Hajime Yatate and Sunrise Bandai Visual. 2007. Idolmaster: Xenoglossia. <https://www.sunrise-inc.co.jp/idolmaster/>. [Online; accessed 4-August-2023].
- [60] Sheng You, Ning You, and Minxue Pan. 2019. PI-REC: Progressive Image Reconstruction Network With Edge and Color Domain. *CoRR abs/1903.10146* (2019). arXiv:1903.10146 <http://arxiv.org/abs/1903.10146>
- [61] Lvmin Zhang, Yi Ji, and Chunping Liu. 2020. DanbooRegion: An Illustration Region Dataset. In *European Conference on Computer Vision (ECCV)*.
- [62] Lvmin Zhang, Chengze Li, Edgar Simo-Serra, Yi Ji, Tien-Tsin Wong, and Chunping Liu. 2021. User-Guided Line Art Flat Filling With Split Filling Mechanism. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*. Computer Vision Foundation / IEEE, 9889–9898. [https://openaccess.thecvf.com/content/CVPR2021/html/Zhang\\_User-Guided\\_Line\\_Art\\_Flat\\_Filling\\_With\\_Split\\_Filling\\_Mechanism\\_CVPR\\_2021\\_paper.html](https://openaccess.thecvf.com/content/CVPR2021/html/Zhang_User-Guided_Line_Art_Flat_Filling_With_Split_Filling_Mechanism_CVPR_2021_paper.html)
- [63] Lvmin Zhang, Chengze Li, Tien-Tsin Wong, Yi Ji, and Chunping Liu. 2018. Two-Stage Sketch Colorization. *ACM Trans. Graph.* 37, 6, Article 261 (12 2018), 14 pages. <https://doi.org/10.1145/3272127.3275090>
- [64] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*. IEEE Computer Society, 2242–2251. <https://doi.org/10.1109/ICCV.2017.244>