

Prompt Augmentation for Self-supervised Text-guided Image Manipulation

Rumeysa Bodur¹, Binod Bhattarai², Tae-Kyun Kim^{1,3}

¹Imperial College London, UK, ²University of Aberdeen, UK, ³KAIST, South Korea

r.bodur18@imperial.ac.uk, binod.bhattarai@abdn.ac.uk, kimtaekyun@kaist.ac.kr

Abstract

Text-guided image editing finds applications in various creative and practical fields. While recent studies in image generation have advanced the field, they often struggle with the dual challenges of coherent image transformation and context preservation. In response, our work introduces prompt augmentation, a method amplifying a single input prompt into several target prompts, strengthening textual context and enabling localised image editing. Specifically, we use the augmented prompts to delineate the intended manipulation area. We propose a Contrastive Loss tailored to driving effective image editing by displacing edited areas and drawing preserved regions closer. Acknowledging the continuous nature of image manipulations, we further refine our approach by incorporating the similarity concept, creating a Soft Contrastive Loss. The new losses are incorporated to the diffusion model, demonstrating improved or competitive image editing results on public datasets and generated images over state-of-the-art approaches.

1. Introduction

The field of text-guided image generation has made significant progress, particularly with the advent of diffusion models [20, 31, 32, 37], ushering in a new era of content creation. This progress extends to text-guided image manipulation, which offers a wide range of applications, spanning from artistic expression to the enhancement of image interpretability. Text-guided image manipulation involves altering an input image based on a user-provided textual prompt, such as changing the appearance or shape of objects, modifying the background, and adding, removing or replacing features. Some sample manipulations performed by our method can be seen in Figure 1. In this context, two intertwined challenges are present: the task of transforming image content in accordance with a provided textual description, and the need to preserve the salient aspects of the original visual information that remain contextually relevant. While recent methodologies [8, 16, 27] have excelled in the domain of image manipulation, they often encounter limitations in addressing both of these challenges



Figure 1. **Text-guided image manipulation.** Illustrative examples generated by our method (bottom row) with localised manipulations based on given text prompts and input images (top row).

concurrently.

Many pioneering text-guided image editing methods rely on text-guided image generation models [11, 16, 18, 27, 30, 41], which tend to be ill-equipped to address the nuanced challenges of image manipulation. Additionally, methods tailored specifically for image editing often require domain-specific training [2, 19]. These methods require training for each instance or translation domain, limiting their scalability across diverse tasks. Despite excelling within their trained domains, they often struggle with the broader generalisation needed for accommodating various editing scenarios. Although there exist generic methods, such as Instruct-Pix2Pix [8], they often struggle to achieve both the faithful transformation of image content and the preservation of contextual details. Some existing methods incorporate masks at inference time [1, 11, 31] or during training [40]. Although these methods show a good performance at content preservation, Dall-E 2 [31] and Imagen Editor [40] fall short as they treat it as an inpainting task without taking into consideration the masked content and DiffEdit [11] shows weaknesses as it relies on a pre-trained image generation model that is not specifically trained for manipulation purposes and is sensitive to the mask detected at inference time.

Augmentation techniques have proven useful for enhancing image translation [5, 6, 9], but their application in label space remains underexplored [3, 4, 25, 26]. Hence, to overcome the aforementioned challenges, this paper intro-

duces the concept of prompt augmentation, which extends a single input prompt into a multitude of target prompts. This augmentation provides the model with a broader and more nuanced understanding of the image editing task at hand. Taking inspiration from DiffEdit [11], which uses user given source and target prompts to generate masks at inference time, we leverage augmented prompts to automatically derive an attention mask that serves as a guide for image editing at training time. This mask plays a pivotal role in distinguishing which regions of the input image should undergo alterations and which should remain unaltered, contributing to the task’s overall precision in localising the desired edit.

In addition to facilitating attention map computation, augmented prompts empower us to establish relations between their respective manipulations. Specifically, areas falling into the obtained masked for each pair should exhibit dissimilarity, reflecting the diverse edits dictated by different prompts, while the unmasked regions should maintain similarity, preserving aspects unaffected by the prompt. With this aim we propose to incorporate a Contrastive Loss (CL) that serves this dual purpose, compelling the model to simultaneously displace edited areas while drawing closer the preserved regions. As manipulation is not a discreet task and the required amount of modification is influenced by the target prompt’s relation to the source image, we propose to soften our contrastive loss by considering this relation. Building upon this foundation, we introduce a novel loss function, the Soft Contrastive Loss (Soft CL) to incorporate the concept of similarity between textual prompts. This approach results in a more dynamic and nuanced interaction between textual prompts and the image, contributing to a higher level of performance.

- We introduce prompt augmentation, expanding input prompts into multiple targets, in order to enhance contextual understanding for image editing, and compute a dynamic attention mask to guide editing localisation.
- We introduce CL to encourage effective editing, pulling preserved regions closer while pushing edited areas to align with the target prompt. Our novel Soft CL incorporates similarity for dynamic prompt interaction, improving adaptability and performance.
- We assess the performance of our method through a comparative analysis with state-of-the-art approaches, achieving competitive results without relying on masks during inference or utilising a paired dataset.

2. Related Work

Text-guided Image Manipulation. Text-guided image manipulation using diffusion models is an evolving field that harnesses diffusion models’ progressive and attention-based architecture to facilitate fine-grained image adjustments. This research trajectory aims to enhance con-

trollability and flexibility while maintaining the fidelity of the input image. DiffusionCLIP [19] leverages CLIP-based losses to guide the diffusion process where they fine-tune a pretrained diffusion model for a specific domain. Imagic [18], Text2Live [2] and Dreambooth [34] involve fine-tuning the entire model for each image, primarily generating variations for objects. Plug-and-Play [39] explores the injection of spatial features and self-attention maps to maintain the overall structural integrity of the image. Prompt-to-Prompt (P2P) [16] dispenses with fine-tuning, instead retaining image structure by assigning cross-attention maps from the original image to the edited one based on corresponding text tokens. Another study [15] uses self-guidance by constraining attention maps or intermediate activations to control the sampling process. Similarly, MasaCtrl [10] proposes to exploit innate features to use mutual attention during inference. InstructPix2Pix [8] utilises P2P to generate target images by accommodating human-like instructions for image editing. SDEdit [27] adopts a two-step approach, where they first introduce noise into the input image and then employ the SDE prior for denoising, ultimately enhancing realism while aligning with user guidance. BlendedDiffusion [1] follows a similar approach but introduces manually created masks at inference time. Similarly, DiffEdit [11] also uses masks at inference but automatically predicts them by using difference of latent noises obtained by input and target prompts provided by the user. Imagen Editor [40] incorporates masks during training by using an object detector for randomly masking out objects to propose inpainting masks during training. In contrast, our method integrates masks during training, specifically tailored to the target manipulation. We determine these masks through the generation of multiple augmentations of the source prompt enabling us to employ a strategy akin to [11] to obtain an attention mask.

Prompt generation and augmentation for image generation. InstructPix-to-Pix [8] proposes to generate instructions for image editing by using a GPT model that they train on manually constructed pairs of captions and instruction. Specialist Diffusion [26] augments the prompts to define the same image with multiple captions that convey the *same* meaning in order to improve the generalisation of the image generation network. They retrieve similar captions, replace words with synonyms or reflect the image augmentation, e.g. horizontal flip, in the text prompt. Another work [24] proposes to generate random sentences including source and target domain in order to calculate a mean difference that will serve as a direction while editing. iEdit [7] generates target prompts by changing words in the input caption in order to retrieve pseudo-target images and guide the model. In [25], the authors propose to augment the prompts with cultural descriptions in order to reduce the culture bias in generative models. Our

approach uniquely integrates prompt augmentation during training, automatically generating multiple augmentations of the source prompt that define possible target prompts. We leverage these augmented prompts to determine masks tailored specifically to the target manipulation, and to enable self-supervised learning, which enhance the model’s ability to learn editing tasks.

Contrastive Loss for image generation. Being a powerful self-supervised learning method, contrastive learning ensures the consistency of image representations across various augmentations by comparing positive pairs against negative ones. While it has delivered remarkable outcomes in numerous fields, its application in image generation, and more specifically, image manipulation, remains relatively uncharted. In Cntr-GAN[43] the discriminator is trained with a loss to push different image representations apart, while drawing augmentations of the same image closer. The first study to explore contrastive learning in the domain of image manipulation [29] proposes PatchNCE loss to minimise the distance between the feature representations of patches from a source image and corresponding patches from generated images while treating randomly sampled patches from other locations as negative samples. XMC-GAN [42] proposes to enforce text–image resemblance with a contrastive discriminator for text-to-image generation.

Soft Contrastive Loss. Most studies utilising contrastive learning rely on hard assignment of samples [14, 29, 42, 43] while only a few explore soft assignment. In a study [21] tested on image retrieving, the authors propose a log ratio loss to enforce the label and feature representation distances to be proportional. Another study [38] follows a similar approach for visual localisation, in which they aim to have the ordering of euclidean distances between features respect the ordering of geometric proximity measure between the corresponding images. Given the continuous nature of textual prompts and image manipulation, in our work, we leverage the similarity between augmented prompts to guide the soft assignment of image representations.

3. Methodology

In this section, we present a comprehensive description of our proposed approach. We start by explaining the prompt augmentation process, followed by how these augmented prompts contribute to obtaining an attention mask and training the model using a contrastive loss. We then delve into the finer refinements of our approach, aiming to accommodate the continuous nature of textual prompts through the integration of Soft-CL and soft prompt augmentation. For a visual overview of our methodology, refer to Figure 2.

3.1. Prompt Augmentation

We implement prompt augmentation as a critical component of our approach, with the primary goal of generat-

ing multiple target prompts to facilitate *self-supervision* for localised image editing. We employ a publicly available large-scale dataset, i.e LAION-5B [35], that contains image-caption pairs. However, given the inherent noise in captions from web-derived datasets, manipulating these captions is challenging. To address this, we employ an image captioning model, i.e. BLIP [22], to generate cleaner descriptions associated with the source images. This step is imperative as it enhances the clarity of captions, providing a structured foundation essential for subsequent modifications. For instance, the original caption for the LAION image shown in Figure 2 is ‘2004 Jaguar XJ8 Daphne AL’, which is not descriptive of the image and cannot be easily manipulated to obtain target prompts. The caption generated with BLIP, ‘the front-end view of a black car’, is cleaner, more descriptive and enables easier manipulation.

With the cleaner captions in hand, we proceed to augment the single input prompt by creating a range of target prompts. This process begins by masking a random noun or adjective within the input prompt. Leveraging the capabilities of a masked language model, specifically BERT [13], we generate potential replacement words for the masked term. Additionally, we enrich our candidate word pool by incorporating semantically related words such as synonyms, antonyms and co-hyponyms, gathered through the NLTK¹ library. From this pool, we randomly select a set of words to generate multiple variations of the input prompt. For instance, for the aforementioned sample in Figure 2, some generated target prompts are: ‘the front-end view of a blue car’, ‘the front-end view of a red car’, ‘the front-end view of a white car’. Please refer to the supplementary material for more samples. These augmented prompts not only diversify the translation capabilities of our network but also lay the groundwork for localised image manipulation by facilitating our self-supervised framework.

3.2. On-The-Fly Mask Generation

Diffusion models conditioned on text generate different noise estimates for each prompt. Drawing inspiration from the approach used in DiffEdit during inference, we exploit this characteristic to reveal the intended areas of manipulation within images during training. In our work, we use Latent Diffusion Model (LDM) [32], which operates in the latent space to mitigate computational complexities associated with diffusion models. The loss for conditional LDM is given in Eqn. 1, where ϵ_θ represents the noise estimation network, which is a UNet [33]. τ_θ is a domain-specific encoder projecting the conditioning target prompt to an intermediate representation, and the step t is uniformly sampled from $\{1, 2, \dots, T\}$. The network parameters θ are optimised to predict the noise $\epsilon_1 \sim \mathcal{N}(0, 1)$ that is used for corrupting the encoded version of the input image.

¹<https://www.nltk.org/>

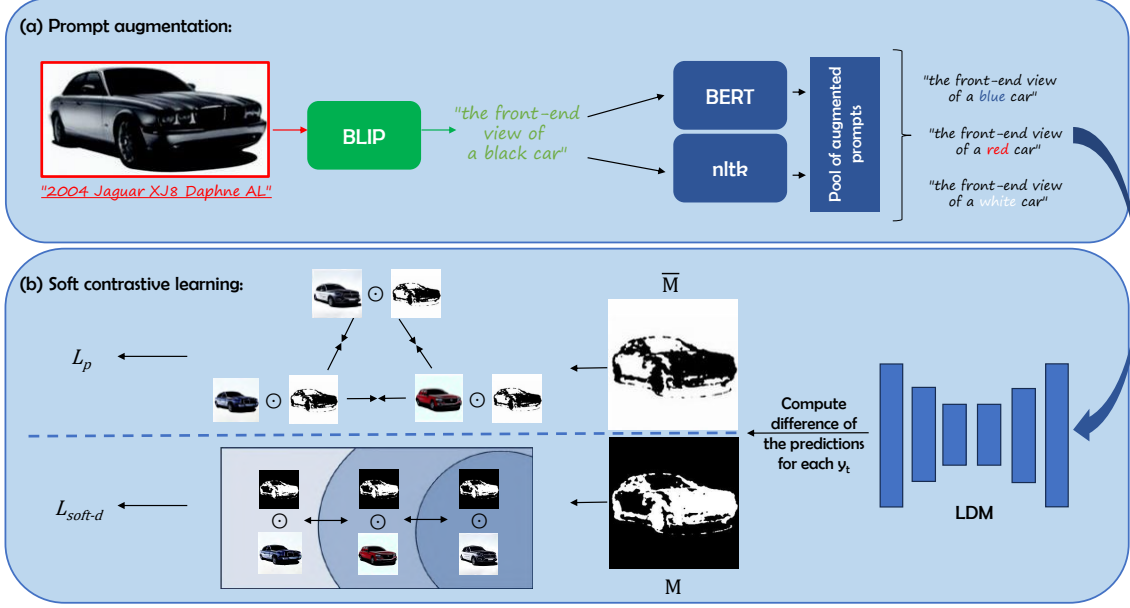


Figure 2. **Overview of our method.** (a) Prompt Augmentation: In order to augment the prompts to facilitate localised image editing we start by refining textual descriptions for source images using the BLIP captioning model [22], resulting in cleaner captions suitable for further processing. Subsequently, we augment this input prompt by generating a range of target prompts using masked language modeling and exploiting word relations. (b) Soft Contrastive Loss (Soft-CL): The augmented prompts are instrumental in computing an attention mask based on the differences between the generated images. This attention mask is used to bring the inverse masked areas of the generated images closer while pushing away masked areas considering their similarity to the prompts.

$$\mathcal{L}_{LDM} = \mathbb{E}_{\mathcal{E}(x), y, \epsilon \sim \mathcal{N}(0,1), t} [|\epsilon - \epsilon_{\theta}(z_t, t, \tau_{\theta}(y))|^2] \quad (1)$$

By passing the augmented prompts through the LDM, the model outputs a set of estimated noises. We use the differences between these estimations to delineate the areas of image manipulation. Given the corrupted version of an encoded input image z_t and N augmented target prompts y_1, y_2, \dots, y_N , the noise estimation network, ϵ_{θ} , outputs N estimated noises during each iteration. We calculate the average of the differences of these estimation as shown in Eqn. 2 where N_p represents the number of unique pairs, and it is given by $N_p = \frac{N \cdot (N-1)}{2}$.

$$\Delta\epsilon = \frac{1}{N_p} \sum_{i,j} |\epsilon_{\theta}(z_t, t, \tau_{\theta}(y_i)) - \epsilon_{\theta}(z_t, t, \tau_{\theta}(y_j))| \quad (2)$$

The outcome, $\Delta\epsilon$, provides valuable insights into the varying degrees of noise introduced by different prompts, signifying areas in the image that are most susceptible to modifications. Subsequently, these dissimilarity values are thresholded to establish a binary mask that distinctly highlights regions of alteration within the image. An example can be seen in Figure 2, where the mask obtained for the car image outlines only areas that would be manipulated in

case the colour of the car would change. In our particular implementation, we employ a threshold empirically set at 0.4 of the range of absolute differences in noise estimates. The resulting binary mask, M , serves as a crucial component for defining the areas targeted for image manipulation within the generated images.

3.3. Contrastive Loss (CL)

Using the mask M derived from augmented prompts, our goal is to guide the model to encourage the masked areas of the latent images, z_1, z_2, \dots, z_N , generated with augmented prompts, y_1, y_2, \dots, y_N , to be in alignment with its respective prompt, while preserving other relevant regions. To achieve this, we introduce a CL that promotes dissimilarity in the masked areas, while drawing the inverse masked areas \bar{M} of the images closer to each other and the source image. The *dissimilarity* within the masked regions is expressed in the first component of the CL as follows:

$$L_d = \frac{1}{N_p} \sum_{i,j} (1 - |(M \odot z_i) - (M \odot z_j)|_2) \quad (3)$$

where \odot represents element-wise multiplication. This term promotes diversity within the areas of the images that should undergo edits, ensuring that the edited content aligns closely with the target prompts. The second component of

CL is equally significant. It operates in parallel, aiming to minimise differences within unmasked regions of the images. These regions are critical for *preserving* the parts of the image that should remain unaltered during the editing process. This part of the loss can be defined as:

$$L_p = \frac{1}{N_p} \sum_{i,j} |(\bar{M} \odot z_i) - (\bar{M} \odot z_j)|_2 \quad (4)$$

The combined CL, unifying both the dissimilarity and preservation components, can be represented as follows:

$$L_{CL} = L_p + \beta \cdot L_d \quad (5)$$

The overall objective to fine-tune the LDM is:

$$L_{obj} = L_{LDM} + \alpha \cdot L_{CL} \quad (6)$$

3.4. Soft Contrastive Loss (Soft-CL)

In image editing, textual prompts present a spectrum of demands that may vary from subtle modifications to profound transformations. For instance, let’s consider the difference between transforming ‘a girl playing in a park’ into ‘a girl playing in a garden’ or ‘a girl playing in a sandbox’. The change from a park to a garden, for example, requires fewer modifications compared to the more extensive transformation into a sandbox. To cater to this diversity in editing requirements, we introduce a refinement to our method in the form of Soft Contrastive Loss (Soft-CL). Soft-CL represents an enhancement over the conventional Contrastive Loss, infusing the loss function with the concept of similarity. To accommodate the varying editing demands, we adapt our loss function by revising Eqn. 3 as follows:

$$L_d^{soft} = \frac{1}{N_p} \sum_{i,j} \left(1 - |(M_i \odot z_i) - (M_j \odot z_j)|_2 \cdot \gamma(y_i, y_j)\right) \quad (7)$$

where $\gamma()$ refers to a similarity measure between the prompts, for which we have used the cosine distance between the CLIP embeddings of the prompts.

In accordance with the Soft-CL we have also enhanced our prompt augmentation strategy to mine prompts in a softer manner rather than selecting them randomly from the pool we construct. For soft prompt augmentation (Soft-PA), we choose the prompts based on differing similarity levels. As we have done in Soft-CL we specifically use the distance between their CLIP embeddings. This process results in a diverse set of prompts, further enhancing the Soft-CL, and consequently the model’s ability to handle a wide range of textual inputs to produce contextually relevant image edits.

4. Experiments

In this section, we analyse our method’s performance through experiments, demonstrating its performance with qualitative and quantitative results, comparing to state-of-the-art methods. We also conduct an ablation study, a hyperparameter study, and discuss limitations.

4.1. Experimental Setup

Datasets. For training, we used Laion-5B images [35] with aesthetics scores greater than 7. For evaluation, we combined images generated using Stable Diffusion v1.4 [32] with real images from COCO [23] and ImageNet [12], totaling 135 images.

Baselines. Our method is compared against several state-of-the-art baselines, namely SDEdit [27], DiffEdit [11], DALL-E 2 [31] and InstructPix2Pix [8]. It should be noted that DALL-E 2 and DiffEdit use manually created or automatically generated masks at inference time, and Instruct-Pix2Pix is trained on a paired dataset. Hence, these methods are not directly comparable to ours.

Evaluation Metrics. To assess the fidelity of our image translation, we utilise CLIPScore [17], which measures the faithfulness of the image to the target prompt. Additionally, SSIM (Structural Similarity Index) is employed to gauge the faithfulness of the translated image to the input image. We also compute the FID [36] in order to evaluate the quality of the generated samples.

4.2. Qualitative Results

In Figure 3, we present a comparison of the results generated by our method against SDEdit [27], DALL-E 2 [31], DiffEdit [11] and InstructPix2Pix [8] using both generated and real images along with heatmaps showing the pixel differences between input and generated images. We observe that SDEdit [27] often encounters challenges in preserving background details, e.g. “*a victorian lamp*”. Furthermore, the editing performance of SDEdit may fall short, e.g. “*a strawberry cake*”. In some scenarios, it struggles with both aspects, as seen in “*a green and white check pillow*”. These limitations can be attributed to the inherent trade-off within SDEdit, balancing between editing and preservation. Additionally, SDEdit is not specifically trained for editing purposes, which contributes to its challenges in maintaining both background integrity and effective editing.

As an inpainting method requiring manual user-provided masks for the intended area of manipulation, DALL-E 2 [31] falls short in faithfully reproducing the characteristics of the input image, despite its effective preservation of the background. Notably, in instances such as “*a strawberry cake*”, the shape of the cake undergoes significant alterations, compromising the fidelity to the original. Similarly, the orientation of the flower in “*a yellow rose*” is not adequately preserved. These observed shortcomings underscore the challenges inherent in maintaining fidelity to input details in the inpainting-based DALL-E 2 framework. DiffEdit [11] employs masks derived from input and output prompts during inference. However, DiffEdit’s performance proves to be sensitive to the precision of the obtained masks and, consequently, the prompts given by the user. While it yields commendable results when the mask




Figure 3. Qualitative comparison of our method against SDEdit [27], DALL-E 2 [31], DiffEdit [11] and InstructPixtoPix [8] using both generated and real images.

Method	CLIPScore (%) \uparrow	FID \downarrow	SSIM- \bar{M} (%) \uparrow	CLIP-R-Precision (%) \uparrow	Human Study
SDEdit [27]	76.76	174	63.07	75.91 \pm 1.5	76.3% - 23.7%
DiffEdit [11]	72.77	85	85.14	67.60 \pm 1.7	73.84% - 26.16%
DALL-E 2 [31]	78.47	151	96.74	81.83 \pm 1.4	68.3% - 31.7%
InstructPix2Pix [8]	77.66	123	81.71	77.45 \pm 1.6	62.15% - 37.85%
Our Method	78.19	133	70.39	80.69 \pm 1.45	-

Table 1. Quantitative comparison of baselines and our method. (\uparrow higher is better, \downarrow lower is better).

is accurately detected, e.g. “*a yellow rose*”, it faces challenges in other scenarios, such as “*a strawberry cake*”, where the change predominantly occurs in the plates due to inferred mask inaccuracies. It should be noted that DiffEdit, designed primarily for inference time, can be seamlessly integrated into our method. InstructPix2Pix [8] undergoes training with fairly aligned target images generated by P2P [16]. Despite its ability to preserve the background and while achieving successful translations in many cases, it carries a notable caveat associated with the use of generated images. The outcomes may exhibit non-realistic qualities, e.g. “*a bear*”, or manipulations can extend beyond the intended region, impacting the entire image, as evident in the check patterned background of “*a green and white check pillow*”. This limitation stems from weaknesses inherited from P2P. On the other hand, our method demonstrates proficient editing capabilities, preserving the background with minimal undesired changes.

4.3. Quantitative Results

In Table 1, we present a quantitative comparison of our method with state-of-the-art approaches. The CLIPScore metric indicates the translation coherency regarding the target prompt, while SSIM- \bar{M} scores offer a perspective on background preservation, providing a comprehensive assessment of diverse methodological approaches. SDEdit, serving as our baseline, achieves a CLIPScore of 76.76 and an SSIM- \bar{M} score of 63.07, aligning with qualitative observations of challenges in preserving background details. DiffEdit, relying on automatically generated masks, obtains a CLIPScore of 72.77 and an SSIM- \bar{M} score of 85.14, showing that despite proficiency in background preservation, it falls short in prompt-based translation. DALL-E 2, as an inpainting method with manual masks, demonstrated a CLIPScore of 78.47 and an outstanding SSIM- \bar{M} score of 96.74, showcasing remarkable background preservation due to its inpainting nature. InstructPix2Pix, trained with fairly aligned target images, achieved a CLIPScore of 77.66 and an SSIM- \bar{M} score of 81.71. While competitive, the SSIM- \bar{M} score indicates some compromise in background preservation, aligning with qualitative analysis. Our method, with a CLIPScore of 78.19 and an SSIM- \bar{M} score of 70.39, demonstrates effective editing with a moderate level of background preservation despite not being trained on a paired dataset or using masks at inference.

Examining the FID values in Table 1 provides valuable in-

sights into the perceptual quality and realism of generated images across different methods. SDEdit reveals a higher FID, indicating a noticeable gap in quality when compared to source images. DiffEdit, relying on automatically generated masks, presents a lower FID, suggesting a closer resemblance to source images. This lower score may be attributed to its effective use of masks for coherent image generation. DALL-E 2 exhibits a moderate FID, which could be attributed to the challenges of seamlessly blending manipulated regions with the rest of the image, impacting the overall perceptual quality. InstructPix2Pix showcases a competitive FID, maintaining a balance between translation and background preservation. Our method maintains a moderate FID, indicating a balance between effective editing and maintaining a satisfactory level of image quality.

We further evaluated our method using CLIP-R-Precision score [28] (R=1)(Tab. 1), known for better alignment with human preferences compared to CLIPScore. We observe that it aligns more closely with our qualitative results as our method is outperformed only by DALL-E 2. We argue that, CLIP-Precision-Score still falls short in capturing the identified shortcomings of DALL-E, highlighted in our qualitative findings.

Following established protocols [2, 18], we conduct a user study on Microworkers, where participants were presented pairs of images and selected the more successful manipulation based on faithfulness to the prompt and input image. In Tab. 1, blue indicates preference for our method and red for the counterpart baseline. The results align more closely with our qualitative observations than the automatic metrics. DALL-E, despite achieving the highest ClipScore, is preferred less than our method and InstructPix2Pix. This is attributed to its low fidelity to masked content, emphasising nuances that automatic metrics might overlook. Similarly, IP2P’s weaknesses, not directly measurable automatically, contribute to its lag behind our method. Further details of this study can be found in the supplementary.

4.4. Ablation Study

Table 2 provides insights from our ablation study, investigating the influence of key components in our proposed method. The absence of CL, Soft-CL, Soft-PA is indicated by \times in the respective columns. Remarkably, the introduction of contrastive loss alone leads to a noticeable enhancement in all metrics, indicating improved translation coherency and content preservation compared to the baseline,



Figure 4. Qualitative comparison of ablation study results.

CL	Soft-CL	Soft-PA	CLIPScore (%) \uparrow	FID \downarrow	SSIM- \bar{M} (%) \uparrow
\times	\times	\times	76.76	174	63.07
\checkmark	\times	\times	77.94	159	68.37
\times	\checkmark	\times	77.53	140	67.21
\times	\checkmark	\checkmark	78.19	133	70.39

Table 2. Ablation study results for the proposed method components. CL denotes contrastive loss, Soft-CL is soft contrastive loss, and Soft-PA is soft prompt augmentation.

which is equivalent to stable diffusion-based SDEdit. While the introduction of soft contrastive loss also shows a rise in all metrics, it becomes superior to CL only with the introduction of soft prompt augmentation, which refines the method’s ability to balance translation coherence and background preservation. In summary, our ablation study underscores the pivotal role of contrastive loss in enhancing translation quality. The incorporation of soft contrastive loss, coupled with soft prompt augmentation, contributes to substantial improvements in overall performance.

In Figure 4, the qualitative ablation study unveils the nuanced impact of each component in refining the translation. The samples, generating different variations of a cake: “an orange cake”, “a mandarin cake”, “a chocolate cake” and “a marble cake” illustrate the model’s behavior in the absence of each component. The baseline model produces results with some difference between “orange” and “mandarin,” despite their similar semantics. Additionally, it struggles to synthesise a “chocolate cake” and “a marble cake”. The introduction of hard-assigned CL improves generation outputs but does not establish a clear relation between the prompts. With the incorporation of Soft-CL the model becomes adept at discerning fine-grained differences. In this context, “an orange cake” and “a mandarin cake” demonstrate the model’s ability to generate semantically related descriptions, reflecting heightened sensitivity to shared characteristics between these citrus fruits. Soft-PA amplifies these effects by introducing diversity in the input

Weight (α)	CLIPScore (%) \uparrow	FID \downarrow	SSIM- \bar{M} (%) \uparrow
0.25	76.11	148	70.75
0.5	76.54	146	70.3
1.0	78.20	133	70.39
2.0	72.04	121	73.22

Table 3. Impact of different weight values for α in Eqn. 6 on model performance metrics.

prompts. In the case of “a chocolate cake” Soft-PA encourages the model to explore a broader spectrum of possibilities, resulting in varied and contextually rich outputs while showing a clear similarity between ‘orange’ and ‘mandarin’ results. The qualitative examples underscore the role of Soft-CL in enhancing the model’s capacity to capture semantic similarities, fostering more coherent and contextually relevant image manipulations. The interplay between Soft-CL and Soft-PA contributes to the generation of diverse and nuanced outputs, showcasing the effectiveness of these components in refining the editing capabilities of our method.

4.5. Hyper-parameter Study

In Table 3, the impact of the weight hyperparameter (α in Eqn. 6) on our model’s performance metrics is shown. Increasing α correlates positively with CLIPScore, indicating improved translation coherency. However, a sharp decrease in CLIPScore is observed with a high value of $\alpha = 2.0$. Concurrently, the FID metric decreases with higher α values, reflecting enhanced image quality. The SSIM- \bar{M} score, emphasising fidelity to input images, peaks at $\alpha = 2.0$. Our analysis suggests that $\alpha = 1.0$ strikes a favorable equilibrium, yielding strong performance across CLIPScore, FID, and SSIM- \bar{M} , making it a judicious choice for our method.

5. Conclusion

We introduced a novel method using prompt augmentation for dynamic mask generation and self-supervised learning in image manipulation. Our soft-contrastive loss achieves effective translations in delineated areas while preserving the rest with minimal undesired changes. Addressing challenges in existing methods, our approach provides a promising solution for localised image manipulations. Experimental evaluations, including comparisons with state-of-the-art techniques and ablation studies, were conducted both qualitatively and quantitatively. Our proposed image editing method stands as a significant contribution, paving the way for prompt augmentation and striking a nuanced balance between successful translations and background preservation.

Acknowledgements. R. Bodur is funded by the Turkish Ministry of National Education. This work was in part sponsored by NST grant (CRC 21011, MSIT), KOCCA grant (R2022020028, MCST), IITP grant (RS-2023-00228996, MSIT). BB is supported by NCS University of Aberdeen Startup grant.

References

- [1] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18208–18218, 2022. [1](#), [2](#)
- [2] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 707–723. Springer, 2022. [1](#), [2](#), [7](#)
- [3] Binod Bhattarai and Tae-Kyun Kim. Inducing optimal attribute representations for conditional gans. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [1](#)
- [4] Binod Bhattarai, Rumeysa Bodur, and Tae-Kyun Kim. Auglabel: Exploiting word representations to augment labels for face attribute classification. *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 2308–2312, 2020. [1](#)
- [5] Rumeysa Bodur, Binod Bhattarai, and Tae-Kyun Kim. 3d dense geometry-guided facial expression synthesis by adversarial learning. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*, 2021. [1](#)
- [6] Rumeysa Bodur, Binod Bhattarai, and Tae-Kyun Kim. Joint training of hierarchical gans and semantic segmentation for expression translation. *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 2308–2312, 2023. [1](#)
- [7] Rumeysa Bodur, Erhan Gundogdu, Binod Bhattarai, Tae-Kyun Kim, Michael Donoser, and Loris Bazzani. iedit: Localised text-guided image editing with weak supervision. *arXiv preprint arXiv:2305.05947*, 2023. [2](#)
- [8] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [2](#), [5](#), [6](#), [7](#)
- [9] Jie Cao, Luanxuan Hou, Ming-Hsuan Yang, Ran He, and Zhenan Sun. Remix: Towards image-to-image translation with limited data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#)
- [10] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22560–22570, 2023. [2](#)
- [11] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. DiffEdit: Diffusion-based Semantic Image Editing with Mask Guidance. In *International Conference on Learning Representations (ICLR)*, 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [3](#)
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. [5](#)
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Association for Computational Linguistics*, 2019. [3](#)
- [14] Yueqi Duan, Lei Chen, Jiwen Lu, and Jie Zhou. Deep embedding learning with discriminative sampling policy. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4959–4968, 2019. [3](#)
- [15] Dave Epstein, Allan Jabri, Ben Poole, Alexei Efros, and Aleksander Holynski. Diffusion self-guidance for controllable image generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 16222–16239. Curran Associates, Inc., 2023. [2](#)
- [16] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. In *International Conference on Learning Representations (ICLR)*, 2022. [1](#), [2](#), [7](#)
- [17] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *EMNLP*, 2021. [5](#)
- [18] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [2](#), [7](#)
- [19] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2426–2435, 2022. [1](#), [2](#)
- [20] Jinseok Kim and Tae-Kyun Kim. Arbitrary-scale image generation and upsampling using latent diffusion model and implicit neural decoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [1](#)
- [21] Sungyeon Kim, Minkyoo Seo, Ivan Laptev, Minsu Cho, and Suha Kwak. Deep metric learning beyond binary supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2288–2297, 2019. [3](#)
- [22] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning (ICML)*, 2022. [3](#), [4](#), [1](#)
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2014. [5](#)
- [24] Yupei Lin, Sen Zhang, Xiaojun Yang, Xiao Wang, and Yukai Shi. Regeneration learning of diffusion models with rich prompts for zero-shot image translation. *arXiv preprint arXiv:2305.04651*, 2023. [2](#)
- [25] Zhixuan Liu, Youeun Shin, Beverley-Claire Okogwu, Youngsik Yun, Lia Coleman, Peter Schaldenbrand, Jihie

- Kim, and Jean Oh. Towards equitable representation in text-to-image synthesis models with the cross-cultural understanding benchmark (ccub) dataset. *arXiv preprint arXiv:2301.12073*, 2023. 1, 2
- [26] Haoming Lu, Hazarapet Tunanyan, Kai Wang, Shant Navasardyan, Zhangyang Wang, and Humphrey Shi. Specialist diffusion: Plug-and-play sample-efficient fine-tuning of text-to-image diffusion models to learn any unseen style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14267–14276, 2023. 1, 2
- [27] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2022. 1, 2, 5, 6, 7
- [28] Dong Huk Park, Samaneh Azadi, Xihui Liu, Trevor Darrell, and Anna Rohrbach. Benchmark for compositional text-to-image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*. Curran, 2021. 7
- [29] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for conditional image synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 3
- [30] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference (SIGGRAPH)*. ACM, 2023. 1
- [31] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1, 5, 6, 7, 2, 3
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3, 5
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, Cham, 2015. 3
- [34] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arxiv:2208.12242*, 2022. 2
- [35] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W. Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R. Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 3, 5
- [36] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, 2020. Version 0.3.0. 5
- [37] Jaskirat Singh, Stephen Gould, and Liang Zheng. High-fidelity guided image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [38] Janine Thoma, Danda Pani Paudel, and Luc V Gool. Soft contrastive learning for visual localization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 11119–11130, 2020. 3
- [39] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. *arXiv preprint arXiv:2211.12572*, 2022. 2
- [40] Su Wang, Chitwan Saharia, Ceslee Montgomery, Jordi Pont-Tuset, Shai Noy, Stefano Pellegrini, Yasumasa Onoe, Sarah Laszlo, David J. Fleet, Radu Soricut, Jason Baldridge, Mohammad Norouzi, Peter Anderson, and William Chan. Imagen editor and editbench: Advancing and evaluating text-guided image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18359–18369, 2023. 1, 2
- [41] Chen Henry Wu and Fernando De la Torre. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7378–7387, 2023. 1
- [42] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-to-image generation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [43] Zhengli Zhao, Zizhao Zhang, Ting Chen, Sameer Singh, and Han Zhang. Image augmentations for gan training. *arXiv preprint arXiv:2006.02595*, 2020. 3