

Simultaneous Multiple-Prompt Guided Generation Using Differentiable Optimal Transport

Yingtao Tian

Google Brain
Tokyo, Japan
alantian@google.com

Marco Cuturi

Google Brain (currently at Apple)
Paris, France
cuturi@apple.com

David Ha

Google Brain
Tokyo, Japan
hadavid@google.com

Abstract

Recent advances in deep learning, such as powerful generative models and joint text-image embeddings, have provided the computational creativity community with new tools, opening new perspectives for artistic pursuits. *Text-to-image synthesis* approaches that operate by generating images from text cues provide a case in point. These images are generated with a latent vector that is progressively refined to agree with text cues. To do so, patches are sampled within the generated image, and compared with the text prompts in the common text-image embedding space; The latent vector is then updated, using gradient descent, to reduce the mean (average) distance between these patches and text cues. While this approach provides artists with ample freedom to customize the overall appearance of images, through their choice in generative models, the reliance on a simple criterion (mean of distances) often causes mode collapse: The entire image is drawn to the *average* of all text cues, thereby losing their diversity. To address this issue, we propose using matching techniques found in the optimal transport (OT) literature, resulting in images that are able to reflect faithfully a wide diversity of prompts. We provide numerous illustrations showing that OT avoids some of the pit-

falls arising from estimating vectors with mean distances, and demonstrate the capacity of our proposed method to perform better in experiments, qualitatively and quantitatively.

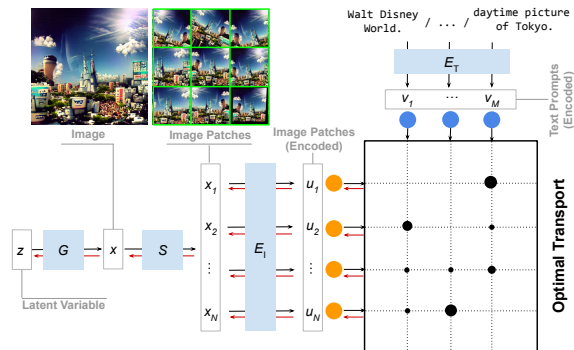
Introduction

The computational creativity community has been at the forefront of engaging with recent advances in deep learning, adopting early on generative models that are able to produce high-quality text and images. Such models offer varying degrees of realism and control to the artist, enabling the generation of results with artistic value. Recent advances have brought forward models that can produce images from natural language prompts, using pre-trained image generative models guided by text descriptions (Radford et al. 2021). The computational creativity community has seized this opportunity, has shared large bodies of code (Burton-King 2021; Murdock 2021a) and generated a large body of artwork, some of which has been curated online (Snell 2021; Murdock).

These tools are favoured by artists because they can shape generation in various ways: For instance, a relevant genera-



(a) Generated images from *two-prompts* using our method. (Left) “Walt Disney World.” and “a daytime picture of Tokyo.” (Right) “A painting of cat.” and “A painting of dog.”.



(b) The architecture of our work. Iteratively, the loss is computed forward (marked by \rightarrow) and the gradient is calculated backward (marked by \leftarrow) to update the latent variable z .

Figure 1: Our method illustrated with generated images and the architecture. In contrast, the existing method would fail with these two-prompts, producing images with less diverse features (left) or a painting with much different art style than single prompt (right). This is because the existing method of taking the mean cannot treat different parts of the image separately, and vector arithmetic in the latent space introduces uncontrollable changes in the semantics. Detailed analysis can be found in text. All figures in this paper are *generated* using pre-trained CLIP and VQGAN models, both publicly released under MIT license.

tive model can be used in that the style of pieces of art that can be produced can be efficiently guided by selecting a relevant generative model G . While this degree of freedom is useful, little has changed on how text prompts are handled in that pipeline:

Images can be generated using the following pipeline: The user supplies a generative model G and a text prompt t . An initial latent vector z is sampled randomly; the fit between its corresponding image $x = G(z)$ and the desired prompt t is quantified using their distance in a common CLIP embedding space; In order to minimize that distance, z is updated iteratively using gradient steps. Because both the CLIP embedding and G are differentiable, gradients for these distances are obtained using automatic differentiation.

In practice, a few more tricks are needed to produce convincing images. To accommodate the important artistic requirement that multiple concepts appear in images, several text prompts are allowed, but are pre-aggregated in embedding space to result in a composite prompt vector. Next, rather than consider the entire image against that composite prompt vector, several patches with random size, orientation and placement are sampled within the image x , and are then compared with the composite prompt, before these distances are averaged to form the overall loss.

These tricks rely therefore on aggregations: the mean of various prompt embeddings is used to define a single target prompt, and the various distances of all patches to that target are also reduced to their average. We argue, and we show later in the paper, that this reliance on averages can cause several issues, causing notably generated images to have parts that are uniformly closer to all prompts, thus defeating the original motivation of using multiple prompts to obtain artistic images with diverse objects. Another important drawback of averaging prompt embeddings is that it can potentially introduce uncontrollable changes in semantics, with a mean prompt embedding falling in a region of the embedding space with no corresponding meaning.

We propose to address this issue by treating the embedded patches of generated image and texts as vectors sampled from two probability distributions, and to use computational optimal transport (OT) (Peyré and Cuturi 2019) to find the best matching between them. As its name suggests, OT tries to find the minimal total effort required to “move” all patches towards texts, using the pairwise distance as the cost for measuring said effort. OT brings two advantages over simply taking the mean: (1) Since patches are randomly sampled, it encourages the intrinsic diversity *inside* a single generated image. (2) OT does not involve vector arithmetic in the latent space, sidestepping issues that may arise from the non-existing semantic of a mean prompt vector. Concretely, we use Sinkhorn’s Algorithm (Cuturi 2013; Séjourné et al. 2019) for the matching, in a way that is efficient and, most importantly, differentiable using OTT-JAX (Cuturi et al. 2022). Such differentiability is crucial to allow the computation of gradient all the way back to z .

Bringing all pieces together, our proposed use of OT enables the generation of images that are diverse and without the issue of unwanted extra semantics, as demonstrated empirically in the paper. Furthermore, since our proposed

method only changes how pairs of (patch, prompt) distances are recombined, it is orthogonal to other existing parts of the pipeline, and consists, implementation-wise, in a simple drop-in replacement of mean operations by optimised matchings (incidentally, taking means can be interpreted as the most naive approach conceivable to match pairs). We start this paper with a background section, needed to detail next our methodology, which is illustrated and validated in various experiments that showcase its performance, and explain why it is able to solve several issues arising from an over-reliance on mean distances and mean prompt embeddings.

Background

In this section, we review two pillars of our work, *prompt-guided image generation* and *differentiable optimal transport*. We argue in this paper that combining both is crucial to address issues we observe in existing generation methods.

Prompt Guided Image Generation

A notable trend in the field of computational creativity is to guide image generation using natural language as prompts. These *text-to-painting* synthesis tools allow artists to specify the content of a painting using prompts from natural languages. This text-driven generation has revolutionized the computational generation of artworks, as evidenced in online curated collections (Snell 2021; Murdock). These advances are made possible by combining two innovations from deep learning:

Powerful image generative models. Such models include recent generative adversarial networks (GANs) (Karras, Laine, and Aila 2019; Karras et al. 2020; Karras et al. 2021), variational autoencoders (van den Oord, Vinyals, and Kavukcuoglu 2017) and diffusion models (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020; Nichol and Dhariwal 2021; Dhariwal and Nichol 2021), that can produce images with high fidelity and diversity. Formally, this process can be denoted as $x = G(z)$ where the generative model $G : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{h \times w \times 3}$ converts a latent space variable $z \in \mathbb{R}^{d_z}$ to an RGB image of height h , weight w and 3 color channels. $x \in \mathbb{R}^{h \times w \times 3}$. z could be further manipulated to allow generating more suitable x (Li, Jin, and Zhu 2021), allowing artist to control the generation of artworks that fall in desired genres (Jin et al. 2017).

Joint modeling of images and natural language. This idea has been long in the making (Thomee et al. 2016; Li et al. 2017), but only recently given a convincing implementation thanks to progress in natural languages modeling (Raffel et al. 2019; Brown et al. 2020), and notably the ability to embed jointly images and text so well that the need for task-specific fine-tuning is eliminated, as shown in CLIP (Radford et al. 2021). CLIP provides two jointly-trained, differentiable encoders, $E_I : \mathbb{R}^{h \times w \times 3} \rightarrow \mathbb{R}^d$ and $E_T : \mathcal{T} \rightarrow \mathbb{R}^d$, for image and text respectively. We do not further elaborate the domain of text \mathcal{T} as it is not the focus of this work. Formally, given an image x and a text t , and a distance function $D : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ the encoded image $u = E_I(x)$ and the encoded text $v = E_T(t)$ are in a common comparable space $\mathcal{U} = \mathbb{R}^d$, and $D(u, v)$ measures the similarity between x and t . In the

case of CLIP that is trained with cosine distance, practically D could be chosen as cosine distance or geodesic distance, both effectively measuring the angle between the two vectors and being trivially differentiable. Ideally, text-driven image generation is now feasible by iteratively adjusting the latent space vector z , to minimize $D(u, v)$, the distance between the encoded image $x = G(z)$ and encoded user-specific prompt t . As G , E_I and D are differentiable, z could be updated using gradient Descent: $z \leftarrow z - \gamma \nabla_z F(z)$ where $\nabla_z F$ is the gradient of F defined as $F(z) = D(E_I(G(z)), E_T(t))$ and γ is a learning rate.

Using a distance from a single image to a single prompt is usually too restrictive. Therefore, and in practice, the distance is computed over pairs of multiple images and texts as follows: On the image side, n patches (a.k.a. cutouts. We use these two terms interchangeably), which we denote as $x_1, \dots, x_n = S(x)$ are randomly sampled from image x in the fashion of image data augmentation (Shorten and Khoshgoftaar 2019). We assume $x_i \in \mathbb{R}^{h \times w \times 3}$ still holds since we can trivially add a resizing step at the end of augmentation. This practice serves as a regularizer to ensure numerical stability and avoid fitting into regions of z where G has bad support. On the text side, m text prompts, denoted as t_1, \dots, t_m , are often considered, which allows artists to explore the possibilities of art by combining multiple texts as directions. Again, they are encoded accordingly, giving $u_1, \dots, u_n : u_i = E_I(x_i) \in \mathbb{R}^d$ and $v_1, \dots, v_m : v_j = E_T(t_j) \in \mathbb{R}^d$. These pairwise distances are then combined to form a loss, which is

$$F(z) = \text{Mean}_D(z) \stackrel{\text{def}}{=} \frac{1}{mn} \sum_{1 \leq i \leq n, 1 \leq j \leq m} D(u_i, v_j), \quad (1)$$

and thus the gradient $\nabla_z F$ reads

$$\nabla_z F = \left(\sum_{1 \leq i \leq n} \frac{\partial \text{Mean}_D}{\partial u_i} \frac{\partial u_i}{\partial x_i} \frac{\partial x_i}{\partial z} \right) \frac{\partial x}{\partial z} \quad (2)$$

where

$$\begin{aligned} \frac{\partial \text{Mean}_D}{\partial u_i} &= \frac{1}{nm} \sum_{1 \leq j \leq m} \frac{\partial D(u_i, v_j)}{\partial u_i} \\ \frac{\partial u_i}{\partial x_i} &= \nabla_x E_I(x_i), \quad \frac{\partial x}{\partial z} = \nabla_z G(z) \end{aligned} \quad (3)$$

and $\partial x_i / \partial x$ is defined as long as the random sampling is differentiable w.r.t. the input image x which is often the case of data augmentations. This framing of text-driven generation has been applied to different generators G , yielding a variety of artistic results: using unconditional GAN generation, like BigGAN (Murdock 2021a), VQGAN (Burton-King 2021) and SIREN (Murdock 2021b); conditional generation using GAN, such as StyleCLIP (Patashnik et al. 2021), that enables editing existing images. In addition to GANs, it can also be applied to Diffusion models (Crowson 2021; Kim and Ye 2021; Nichol et al. 2021).

Differentiable Optimal Transport

Optimal transport (OT), as its name suggests, can be understood as finding an efficient way to ‘move’ or ‘transport’, the

mass from a probability distribution to another distribution. We borrow notations from the survey book (Peyré and Cuturi 2019) and focus on one of the canonical OT formulations, one that was proposed in (Kantorovich 1942). A discrete measure with weights a on locations u_1, \dots, u_n would be denoted as $\alpha = \sum_{1 \leq i \leq n} a_i \delta_{u_i}$, where notation δ_u stands for a Dirac mass at location u . Similarly, for weights b on locations v_1, \dots, v_m we have $\beta = \sum_{1 \leq j \leq m} b_j \delta_{v_j}$. A possible way to map a discrete measure α onto β , given a cost matrix $\mathbf{C} \in \mathbb{R}_+^{n \times m}$, can be represented with a coupling matrix $\mathbf{P} \in \mathbb{R}_+^{n \times m}$, where the amount of mass transported from the i -th location in α to j -th location in β is stored as $\mathbf{P}_{i,j}$. The set of admissible couplings, \mathbf{U} , is defined through a and b as

$$\mathbf{U}(a, b) \stackrel{\text{def}}{=} \left\{ \mathbf{P} \in \mathbb{R}_+^{n \times m} : \sum_j \mathbf{P}_{i,j} = a, \quad \sum_i \mathbf{P}_{i,j} = b \right\},$$

These row- and column-sum constraints for \mathbf{P} indicate that the entire mass from α is indeed transported to β . Kantorovich’s problem of interest is

$$L(a, b, \mathbf{C}) \stackrel{\text{def}}{=} \min_{\mathbf{P} \in \mathbf{U}(a, b)} \langle \mathbf{C}, \mathbf{P} \rangle \stackrel{\text{def}}{=} \sum_{i,j} \mathbf{C}_{i,j} \mathbf{P}_{i,j},$$

which can be solved using linear programming, notably network flow solvers. The linear programming route, while well established, has a few drawbacks: it is slow, with an unstable solution. A possible workaround is to add an entropic regularization term, where the entropy of \mathbf{P} reads $\mathbf{H}(\mathbf{P}) \stackrel{\text{def}}{=} - \sum_{i,j} (\mathbf{P}_{i,j} (\log(\mathbf{P}_{i,j}) - 1))$. The regularized problem reads:

$$L^\epsilon(a, b, \mathbf{C}) \stackrel{\text{def}}{=} \min_{\mathbf{P} \in \mathbf{U}(a, b)} \langle \mathbf{C}, \mathbf{P} \rangle - \epsilon \mathbf{H}(\mathbf{P}).$$

This regularization has several practical virtues: the regularized problem can be solved efficiently with Sinkhorn’s Algorithm (Cuturi 2013; Séjourné et al. 2019), a fast iterative algorithm that only uses matrix-vector arithmetic. Another advantage, equally important in our setting, is that this approach, as implemented in OTT-JAX (Cuturi et al. 2022) results in fully differentiable quantities. Namely, assume that the cost matrix \mathbf{C} is provided in the form of a differentiable function resulting in entries $\mathbf{C}_{i,j} \stackrel{\text{def}}{=} \mathbf{C}(u_i, v_j)$. Then the gradient of L^ϵ w.r.t. u_i exists and is defined everywhere:

$$\forall i, 1 \leq i \leq n, \left\| \frac{\partial L^\epsilon}{\partial u_i} \right\| < \infty. \quad (4)$$

Not that the optimal solution \mathbf{P}^ϵ corresponding to L^ϵ can also be differentiated w.r.t. any of the relevant inputs, using the implicit function Theorem (Krantz and Parks 2002), as proposed in OTT-JAX (Cuturi et al. 2022), but this is not used in this paper because we rely on Danskin’s Theorem (Danskin 1966) (a.k.a Envelope Theorem) to differentiate L^ϵ w.r.t. \mathbf{C} .

Methodology

Our motivation comes from the concern arising from using an averaged loss Mean_D . By focusing on means, all sampled



(a) Our Proposed Method (OT)



(b) Baseline (Mean)

Figure 2: The generated image from two prompts: “Walt Disney World.” and “daytime picture of Tokyo.” Compare with the baseline, our methods generates images with better diversity (Disney-like architectures vs. city scene) while blending them well.

patches are encouraged to move uniformly to the mean of all prompts. This undermines the very motivation of introducing multiple prompts, which is to allow artists to obtain spatial diversity in the generated images, with various areas reflecting the diversity prompts. Furthermore, taking the mean in the embedding space introduces gradients in unwanted directions. Since the locations in the embedding space are associated with semantics, doing so may introduce uncontrollable, redundant semantics. To make things worse, the mean arithmetic effectively assumes an Euclidean space, which is inconsistent to the CLIP model that is trained with cosine distance in the embedding space.

To address these issues, it is possible to devise an arithmetic in non-Euclidean Space. However, finding a proper choice that works well with the rest of pipeline is not trivial and warrants a separate study. Instead we propose to eliminate the undesired simplifications brought by mean arithmetics, to replace Mean_D in Equation 1 with an optimal transport loss,

$$F = L^\epsilon(a, b, [D(u_i, v_j)]_{i,j}) \quad (5)$$

where $a_i = 1/n$ and $b_j = 1/m$, and the cost matrix \mathbf{C} is populated with pairwise distance D evaluations. Now, the gradient $\nabla_z F$ reads

$$\nabla_z F = \left(\sum_{1 \leq i \leq n} \frac{\partial L^\epsilon}{\partial u_i} \frac{\partial u_i}{\partial x_i} \frac{\partial x_i}{\partial x} \right) \frac{\partial x}{\partial z}. \quad (6)$$

Comparing with Equation 2, the only different term is $\frac{\partial L^\epsilon}{\partial u_i}$ which is also defined as in Equation 4. Along with other terms (see Equation 3), all terms are defined, and thus we know that $\nabla_z F$ is also well-defined and can be used in the iteratively updating of z :

$$z \leftarrow z - \gamma \nabla_z F$$

In doing so, the above mentioned issues are solved for the following reasons:

OT Treats different patches differently. As OT matches patches and text prompts, it naturally introduces a distinct treatment of patches according to their distances to text prompts. As the patches are randomly sampled, it encourages the intrinsic diversity *inside* a single generated image.

OT does not involve arithmetic in the latent space. OT relies on distances, but does not use averages in embedding spaces. Therefore it does not produce synthetic prompts in embeddings space that may not correspond to semantics. Furthermore, OT is agnostic to how distances are defined: any distance, other than cosine distance or geodesic distance, could be used to populate matrix \mathbf{C} .

Experiments

In this section, we highlight a few possibilities brought forward by using our methodology when handling multiple text prompts. Due to the creative nature of text-to-image synthesis, there is no standard measuring stick, such as classification accuracy, to provide a simple comparison between methods. Nevertheless, we consider a few tasks that can help us gain insight into the novelty, the properties and the behavior of our method. We consider:

Generated Image. Naturally the foremost task is to show the generated image x with multiple text prompts t_1, \dots, t_m . In this task, we focus on whether the generated image represents the text prompts in a way that is distinctive and subjectively recognized by human viewers.

Patches (Cutouts) from Generated Images. Our method improves the diversity of patches through increasing the correlation between the distribution of randomly sampled patches and multiple text prompts, as we identify as a source of issues from existing practices. In this task, we show the patches and organize them by text prompt. Formally, we show the n patches x_1, \dots, x_n sampled from x , and group x_i by $j^* = \arg \min_j D(u_i, v_j)$, the closest text prompt in



(a) Patches (cutouts) from our method (OT)



(b) Patches (cutouts) from baseline (Mean)

Prompt 0 : Walt Disney World.
36 out of 64 cutouts are closer to Prompt 0.



Prompt 1 : A daytime picture of Tokyo.
28 out of 64 cutouts are closer to Prompt 1.



(c) Prompts closer to each prompt, in our proposed method (OT)

Prompt 0 : Walt Disney World.
49 out of 64 cutouts are closer to Prompt 0.



Prompt 1 : A daytime picture of Tokyo.
15 out of 64 cutouts are closer to Prompt 1.



(d) Prompts closer to each prompt, in baseline (Mean)

Figure 3: The cutouts (patches) from generated images in Figure 2, for both our proposed method (OT) and the baseline. We show in (a) and (b) the sampled patches. Then in (c) and (d) we group these patches by the closer (measure by D) prompt they are to. Due to space constraints, we only show the number of each group and six patches that are mostly closet.

the embedding space.

Tangent of Patches (Cutouts) on Cost Plane. We identify the issue materialize in the way gradient information is pass from F back to patches, which is $\partial \text{Mean}_D / \partial u_i$ part in Equation 2, and propose to use L_C^ϵ such that the $\partial L_C^\epsilon / \partial u_i$ part in Equation 6, is better.

To quantitatively qualify such property, a few extra deliberations are needed. Concretely, we first define

$$\phi(u_i) : \mathbb{R}^d \rightarrow \mathbb{R}^m \stackrel{\text{def}}{=} [D(u_i, v_1), \dots, D(u_i, v_m)],$$

which is by definition a differentiable mapping from the aforementioned embedding space \mathbb{R}^d to \mathbb{R}^m , a m -d space of distances to prompts where the j -th element is the distance


to prompt j . As $\partial L_C^\epsilon / \partial u_i \in T_{u_i}$ (the tangent space of \mathbb{R}^d at u_i), the pushforward by ϕ at u_i is defined as $d\phi : T_{u_i} \mathbb{R}^d \rightarrow T_{\phi(u_i)} \mathbb{R}^m$ such that when applied to the gradient,

$$w_i = d\phi(\partial L_C^\epsilon / \partial u_i) \quad (7)$$

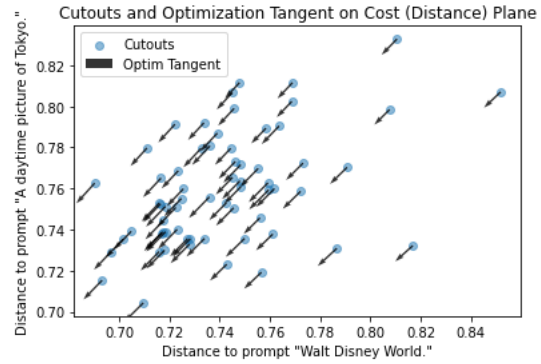
is in the tangent space of \mathbb{R}^m . Intuitively, w_i is a m -dimensional vector whose j -th element denotes the component of gradient that moves the i -th patch towards the j -th text prompt.

Comparing our Method with the Baseline for Two Prompts Setting

In this experiment, we focus on a scenario with $M = 2$ prompts, “Walt Disney World.” and “daytime picture of



(a) Our method (OT)



(b) Baseline (Mean)

Figure 4: Tangent after pushforward of the gradients on each patch (cutout) in the embedding space to the cost plane. Each blue dot is a patch (cutout), and intuitively, its coordinate shows the distance to one of two prompts, while its arrow shows the force of gradient that pushes it towards the prompts. On the left side, in our method the force of gradient pushes patches to prompts with different “mix ratio”, promoting the intrinsic diversity in the generated image from which patches are sampled. On the right side, in the baseline all patches are pushed for the same mix of prompts, thus leading to less diversity. Formally, the exact form and motivation for the tangent could be found mathematically in Equation 7 and its discussions.

Tokyo.” We compare two models, our proposed approach with Optimal Transport (Equation 5) and the baseline using Mean (Equation 1), with the purpose of investigating the behavior of these methods and the difference made by our approach. We keep all other configurations the same. Namely, we use a pre-trained VQGAN (Esser, Rombach, and Ommer 2021) on Imagenet dataset, $N = 64$ randomly sampled patch, and 1000 iterations of updating z . We organize the conducted tasks as explained before.

Generated Image and Patches (Cutouts) from it. In Figure 2 we show the generated image from both methods. Also in Figure 3 we show the patches (cutouts) sampled from the generated images at the end of all iterations.

We observe that OT helps generate images where patches (cutouts) are more balanced (36/28 vs 49/15). Furthermore, OT’s results are more diverse for two prompts. For OT, patches close to “Walt Disney World.” are more like close-ups and patches close to “A daytime picture of Tokyo.” are mostly zoomed-out. As patches are randomly done, it reflects the intrinsic property of generated images.

Tangent of Patches (Cutouts) on Cost Plane. We pushforward gradients on the patches’ embedding space to this cost plane, as explained in Equation 7, and show the results in Figure 4. We observe that our method using OT clearly shows that the positions in the cost plane reveal negative correlation, which means that different parts of the generated images are successfully encouraged to provide contribution to the similarities to different promoters. This is cross-verified by the “fan out” of tangents pushed forward to the cost plane, which shows the divergent gradients providing patch-specific directions in updating. In contrast, baseline methods are simply learning to be the mean of two prompts’ embeddings, as the tangents show the uniform gradient direction which does not distinguish between different prompts.

Our Method’s Behavior with Multiple Prompts

Having comparing our OT-based method with the baseline on the two prompts setting, we shift our focus to the scenario where our method is applied to multiple prompts. As this is we designed our method to expose fine differentiation among prompts, it becomes interesting to investigate such behavior when the number of prompts increases. In doing so, we consider totally $M = 6$ prompts, numbered from P0 to P5:

- * P0: Impressionism / Edgar Degas/ Landscape at Valery-sur-Somme
- * P1: Impressionism Laszlo Mednyanszky/ Landscape in the Alps (View from the Rax)
- * P2: Romanticism / J.M.W. Turner/ The Lake, Petworth, Sunset; Sample Study
- * P3: Romanticism / George Stubbs/ Hound Coursing a Stag
- * P4: Realism / Alexey Venetsianov/ In the Fields. Spring
- * P5: Realism / Alexey Venetsianov/ A Peasant Woman with Scythe and Rake

and as the prompts suggest, we use a pre-trained VQGAN on WikiArt dataset consisting mostly of paintings. The purpose is to both show that our method could be applied to generative models trained from different genre data, and also that the painting allows easier qualitative comparison of both objects and artistic styles. As the same setting mentioned above, $N = 64$ randomly sampled patch, and 1000 iterations of updating are used. We conduct tasks as explained before.

Generated Image. In Figure 5, we show in the first group the generated images corresponding to these prompts individually, and in the second group the generated images by combining prompts using our proposed method. We observe that our method is capable of composing the instructions from several prompts, in terms of styles and objects, into the same canvas.



(a) P0: Impressionism / Edgar Degas/ Landscape at Valery-sur-Somme



(b) P1: Impressionism Laszlo Mednyanszky/ Landscape in the Alps (View from the Rax)



(c) P2: Romanticism / J.M.W. Turner/ The Lake, Petworth, Sunset; Sample Study



(d) P3: Romanticism / George Stubbs/ Hound Coursing a Stag



(e) P4: Realism / Alexey Venetianov/ In the Fields. Spring



(f) P5: Realism / Alexey Venetianov/ A Peasant Woman with Scythe and Rake



(g) P0



(h) P0 + P1



(i) P0 + P1 + P2



(j) P0 + P1 + P2 + P3



(k) P0 + P1 + P2 + P3 + P4



(l) P0 + P1 + P2 + P3 + P4 + P5

Figure 5: The generated images from multiple (6) prompts, labeled P0 to P5. (a) - (f): The first group of 6 images are generated using each one prompt respectively, as a controlling group. (g) - (i): The second group of 6 images are the generated images with multiple (1 to 6) prompts respectively from our proposed method, each one of which using a combination of multiple problems specified in the caption.

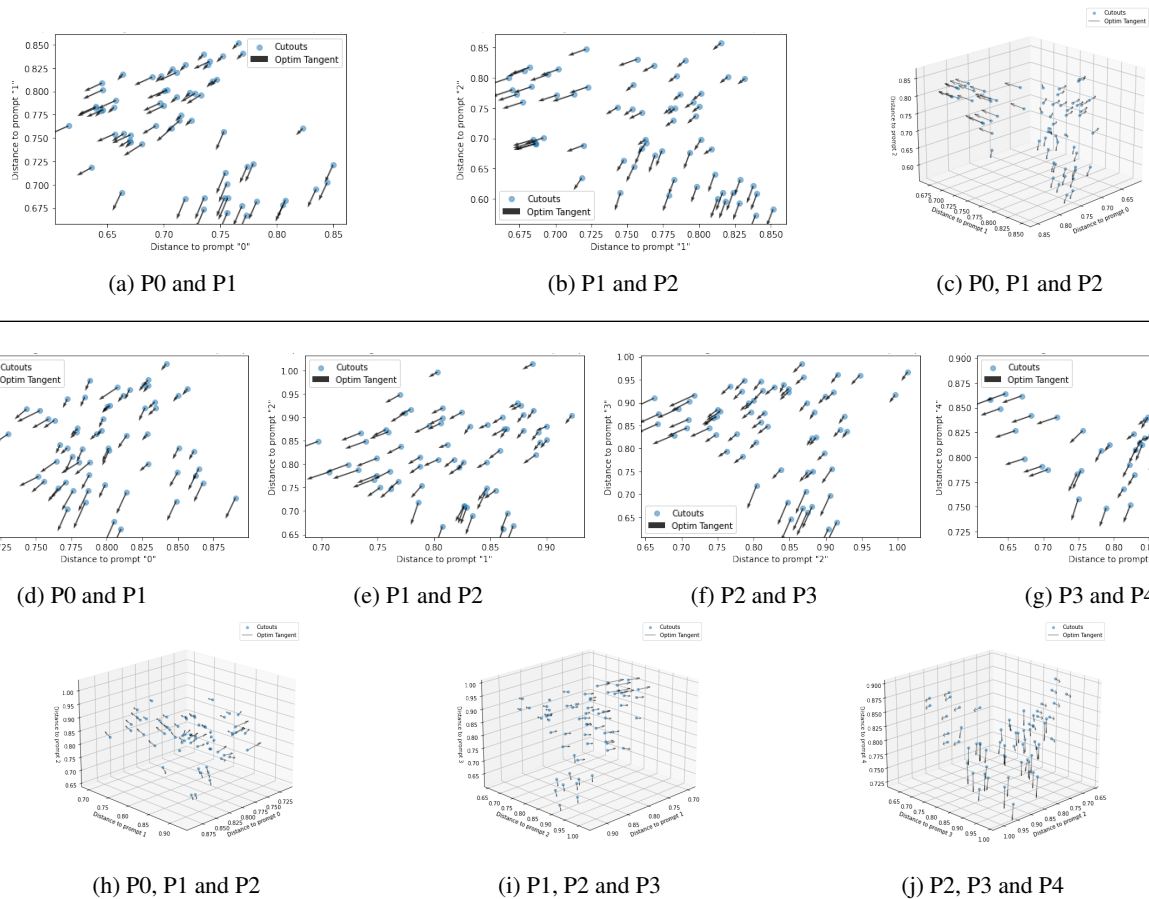


Figure 6: Tangent, representing the gradients on patches (cutouts) after they are pushed forward to Cost Plan. The first group is for the generation with 3 prompts and the second group is for the generation with 6 prompts, showing in 2D and 3D slices.

Tangent of Patches (Cutouts) on Cost Plane. In Figure 6, we show that the good behavior on tangent remains even for multiple prompts. This means that our method is capable of guiding generating images that are diverse in its contents w.r.t. multiple prompts.

Conclusion and Future Work

In this paper we discuss the problem in dealing with multiple text prompts in the setting of text-driven image generation for computational creativity setting. We then propose to address the issue using OT (Optimal Transport) between sampled patches in the generated image and multiple text prompts, and show its theoretical motivation and quantitative and qualitative empirical results highlighting the advantage brought by our proposed method.

One of the advantages in our method is that it is in theory orthogonal to other parts in the whole text driven image generation pipeline, as we show primarily that it works for VQGAN trained on several datasets. We envision that future work would investigate leveraging our proposed method to other drastically different forms of generative method, such as diffusion models. Another possible future direction may principally study the combination of optimal transport and

adaptive sampling where in our proposed work only random sampling is used for simplicity.

Author Contributions

Yingtiao: Ideated the problem and the method, conducted experiments, drafted the paper.

Marco Cuturi: Provided advises, helped designing the method / experiments, helped drafting the paper.

David Ha: Provided advises and gave feedback / suggestions for the whole work, helped drafting the paper.

Acknowledgement

The authors thank Yujin Tang, Jerry Li, Mathieu Blondel for their valuable comments and suggestions.

References

- [Brown et al. 2020] Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- [Burton-King 2021] Burton-King, S. 2021. Introduction to vqgan+clip. <https://bit.ly/3rcedh4>.

- [Crowson 2021] Crowson, K. 2021. Clip guided diffusion.
- [Cuturi et al. 2022] Cuturi, M.; Meng-Papaxanthos, L.; Tian, Y.; Bunne, C.; Davis, G.; and Teboul, O. 2022. Optimal transport tools (ott): A jax toolbox for all things wasserstein. *arXiv preprint arXiv:2201.12324*.
- [Cuturi 2013] Cuturi, M. 2013. Sinkhorn distances: Light-speed computation of optimal transport. *Advances in neural information processing systems* 26:2292–2300.
- [Danskin 1966] Danskin, J. M. 1966. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics* 14(4):641–664.
- [Dhariwal and Nichol 2021] Dhariwal, P., and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *arXiv preprint arXiv:2105.05233*.
- [Esser, Rombach, and Ommer 2021] Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12873–12883.
- [Ho, Jain, and Abbeel 2020] Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*.
- [Jin et al. 2017] Jin, Y.; Zhang, J.; Li, M.; Tian, Y.; Zhu, H.; and Fang, Z. 2017. Towards the automatic anime characters creation with generative adversarial networks. *arXiv preprint arXiv:1708.05509*.
- [Kantorovich 1942] Kantorovich, L. 1942. On the transfer of masses (in russian). In *Doklady Akademii Nauk*, 227–229.
- [Karras et al. 2020] Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8110–8119.
- [Karras et al. 2021] Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2021. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* 34.
- [Karras, Laine, and Aila 2019] Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4401–4410.
- [Kim and Ye 2021] Kim, G., and Ye, J. C. 2021. Diffusion-clip: Text-guided image manipulation using diffusion models. *arXiv preprint arXiv:2110.02711*.
- [Krantz and Parks 2002] Krantz, S. G., and Parks, H. R. 2002. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media.
- [Li et al. 2017] Li, A.; Jabri, A.; Joulin, A.; and van der Maaten, L. 2017. Learning visual n-grams from web data. In *Proceedings of the IEEE International Conference on Computer Vision*, 4183–4192.
- [Li, Jin, and Zhu 2021] Li, M.; Jin, Y.; and Zhu, H. 2021. Surrogate gradient field for latent space manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6529–6538.
- [Murdock] Murdock, R. @advadnoun. <https://twitter.com/advadnoun>.
- [Murdock 2021a] Murdock, R. 2021a. Big sleep: A simple command line tool for text to image generation, using openai’s clip and a biggan.
- [Murdock 2021b] Murdock, R. 2021b. Deep daze: A simple command line tool for text to image generation using openai’s clip and siren (implicit neural representation network).
- [Nichol and Dhariwal 2021] Nichol, A., and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*.
- [Nichol et al. 2021] Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- [Patashnik et al. 2021] Patashnik, O.; Wu, Z.; Shechtman, E.; Cohen-Or, D.; and Lischinski, D. 2021. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2085–2094.
- [Peyré and Cuturi 2019] Peyré, G., and Cuturi, M. 2019. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning* 11(5-6):355–607.
- [Radford et al. 2021] Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*.
- [Raffel et al. 2019] Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- [Séjourné et al. 2019] Séjourné, T.; Feydy, J.; Vialard, F.-X.; Trounev, A.; and Peyré, G. 2019. Sinkhorn divergences for unbalanced optimal transport. *arXiv preprint arXiv:1910.12958*.
- [Shorten and Khoshgoftaar 2019] Shorten, C., and Khoshgoftaar, T. M. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data* 6(1):1–48.
- [Snell 2021] Snell, C. 2021. Alien dreams: An emerging art scene. <https://ml.berkeley.edu/blog/posts/clip-art/>.
- [Song, Meng, and Ermon 2020] Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- [Thomee et al. 2016] Thomee, B.; Shamma, D. A.; Friedland, G.; Elizalde, B.; Ni, K.; Poland, D.; Borth, D.; and Li, L.-J. 2016. Yfcc100m: The new data in multimedia research. *Communications of the ACM* 59(2):64–73.
- [van den Oord, Vinyals, and Kavukcuoglu 2017] van den Oord, A.; Vinyals, O.; and Kavukcuoglu, K. 2017. Neural discrete representation learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6309–6318.