
Generative Adversarial Transformers

Drew A. Hudson^{§1} C. Lawrence Zitnick²

Abstract

We introduce the GANsformer, a novel and efficient type of transformer, and explore it for the task of visual generative modeling. The network employs a bipartite structure that enables long-range interactions across the image, while maintaining computation of linear efficiency, that can readily scale to high-resolution synthesis. It iteratively propagates information from a set of latent variables to the evolving visual features and vice versa, to support the refinement of each in light of the other and encourage the emergence of compositional representations of objects and scenes. In contrast to the classic transformer architecture, it utilizes multiplicative integration that allows flexible region-based modulation, and can thus be seen as a generalization of the successful StyleGAN network. We demonstrate the model’s strength and robustness through a careful evaluation over a range of datasets, from simulated multi-object environments to rich real-world indoor and outdoor scenes, showing it achieves state-of-the-art results in terms of image quality and diversity, while enjoying fast learning and better data-efficiency. Further qualitative and quantitative experiments offer us an insight into the model’s inner workings, revealing improved interpretability and stronger disentanglement, and illustrating the benefits and efficacy of our approach. An implementation of the model is available at <https://github.com/dorarad/gansformer>.

1. Introduction

The cognitive science literature speaks of two reciprocal mechanisms that underlie human perception: the *bottom-up* processing, proceeding from the retina up to the cortex, as local elements and salient stimuli hierarchically group together to form the whole (Gibson, 2002), and the *top-down* processing, where surrounding context, selective attention

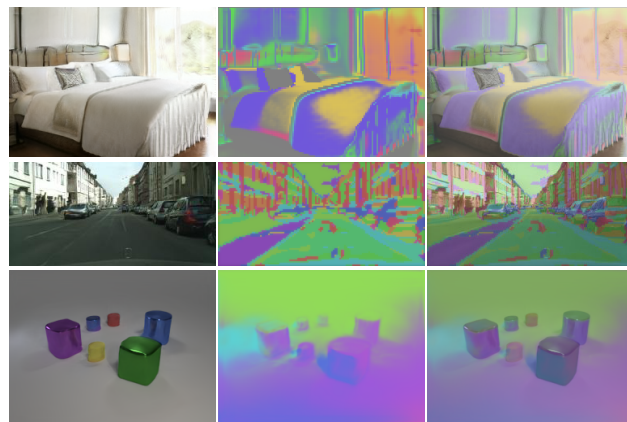


Figure 1. Sample images generated by the GANsformer, along with a visualization of the model attention maps.

and prior knowledge inform the interpretation of the particular (Gregory, 1970). While their respective roles and dynamics are being actively studied, researchers agree that it is the interplay between these two complementary processes that enables the formation of our rich internal representations, allowing us to perceive the world in its fullest and create vivid imageries in our mind’s eye (Intaitė et al., 2013).

Nevertheless, the very mainstay and foundation of computer vision over the last decade – the Convolutional Neural Network, surprisingly does not reflect this bidirectional nature that so characterizes the human visual system, and rather displays a one-way feed-forward progression from raw sensory signals to higher representations. Their local receptive field and rigid computation reduce their ability to model long-range dependencies or develop holistic understanding of global shapes and structures that goes beyond the brittle reliance on texture (Geirhos et al., 2019), and in the generative domain especially, they are linked to considerable optimization and stability issues (Zhang et al., 2019) due to their fundamental difficulty in coordinating between fine details across a generated scene. These concerns, along with the inevitable comparison to cognitive visual processes, beg the question of whether convolution alone provides a complete solution, or some key ingredients are still missing.

[§]I wish to thank Christopher D. Manning for the fruitful discussions and constructive feedback in developing the bipartite transformer, especially when explored for language, as well as for the kind financial support that allowed this work to happen.

¹Computer Science Department, Stanford University, CA, USA
²Facebook AI Research, CA, USA. Correspondence to: Drew A. Hudson <dorarad@cs.stanford.edu>.

Meanwhile, the NLP community has witnessed a major revolution with the advent of the Transformer network (Vaswani et al., 2017), a highly-adaptive architecture centered around relational attention and dynamic interaction. In response, several attempts have been made to integrate the transformer into computer vision models, but so far they have met only limited success due to scalability limitations stemming from their quadratic mode of operation.

Motivated to address these shortcomings and unlock the full potential of this promising network for the field of computer vision, we introduce the Generative Adversarial Transformer, or GANsformer for short, a simple yet effective generalization of the vanilla transformer, explored here for the task of visual synthesis. The model features a bipartite construction for computing soft attention, that iteratively aggregates and disseminates information between the generated image features and a compact set of latent variables to enable bidirectional interaction between these dual representations. This proposed design achieves a favorable balance, being capable of flexibly modeling global phenomena and long-range interactions on the one hand, while featuring an efficient setup that still scales linearly with the input size on the other. As such, the GANsformer can sidestep the computational costs and applicability constraints incurred by prior work, due to their dense and potentially excessive pairwise connectivity of the standard transformer (Zhang et al., 2019; Brock et al., 2019), and successfully advance generative modeling of images and scenes.

We study the model’s quantitative and qualitative behavior through a series of experiments, where it achieves state-of-the-art performance for a wide selection of datasets, of both simulated as well as real-world kinds, obtaining particularly impressive gains in generating highly-compositional multi-object scenes. As indicated by our analysis, the GANsformer requires less training steps and fewer samples than competing approaches to successfully synthesize images of high quality and diversity. Further evaluation provides robust evidence for the network’s enhanced transparency and compositionality, while ablation studies empirically validate the value and effectiveness of our approach. We then present visualizations of the model’s produced attention maps to shed more light upon its internal representations and synthesis process. All in all, as we will see through the rest of the paper, by bringing the renowned GANs and Transformer architectures together under one roof, we can integrate their complementary strengths, to create a strong, compositional and efficient network for visual generative modeling.

2. Related Work

Generative Adversarial Networks (GANs), originally introduced in 2014 (Goodfellow et al., 2014), have made remarkable progress over the past few years, with signif-

icant advances in training stability and dramatic improvements in image quality and diversity, turning them to be nowadays a leading paradigm in visual synthesis (Radford et al., 2016; Brock et al., 2019; Karras et al., 2019). In turn, GANs have been widely adopted for a rich variety of tasks, including image-to-image translation (Isola et al., 2017), super-resolution (Ledig et al., 2017), style transfer (Choi et al., 2018), and representation learning (Donahue et al., 2016), to name a few. But while automatically produced images for faces, single objects or natural scenery have reached astonishing fidelity, becoming nearly indistinguishable from real samples, the unconditional synthesis of more structured or compositional scenes is still lagging behind, suffering from inferior coherence, reduced geometric consistency and, at times, lack of global coordination (Johnson et al., 2018; Casanova et al., 2020; Zhang et al., 2019). As of now, faithful generation of structured scenes is thus yet to be reached.

Concurrently, the last couple of years saw impressive progress in the field of NLP, driven by the innovative architecture called Transformer (Vaswani et al., 2017), which has attained substantial gains within the language domain and consequently sparked considerable interest across the deep learning community (Vaswani et al., 2017; Devlin et al., 2019). In response, several attempts have been made to incorporate self-attention constructions into vision models, most commonly for image recognition, but also in segmentation (Fu et al., 2019), detection (Carion et al., 2020), and synthesis (Zhang et al., 2019). From structural perspective, they can be roughly divided into two streams: those that apply local attention operations, failing to capture global interactions (Cordonnier et al., 2020; Parmar et al., 2019; Hu et al., 2019; Zhao et al., 2020; Parmar et al., 2018), and others that borrow the original transformer structure as-is and perform attention globally, across the entire image, resulting in prohibitive computation due to its quadratic complexity, which fundamentally hinders its applicability to low-resolution layers only (Zhang et al., 2019; Brock et al., 2019; Bello et al., 2019; Wang et al., 2018; Esser et al., 2020; Dosovitskiy et al., 2020; Jiang et al., 2021). Few other works proposed sparse, discrete or approximated variations of self-attention, either within the adversarial or autoregressive contexts, but they still fall short of reducing memory footprint and computation costs to a sufficient degree (Shen et al., 2020; Ho et al., 2019; Huang et al., 2019; Esser et al., 2020; Child et al., 2019).

Compared to these prior works, the GANsformer stands out as it manages to avoid the high costs ensued by *self attention*, employing instead *bipartite attention* between the image features and a small collection of latent variables. Its design fits naturally with the generative goal of transforming source latents into an image, facilitating long-range interaction without sacrificing computational efficiency. Rather,




Figure 2. We introduce the GANsformer network, that leverages a bipartite structure to allow long-range interactions, while evading the quadratic complexity standard transformers suffer from. We present two novel attention operations over the bipartite graph: *simplex* and *duplex*, the former permits communication in one direction, in the generative context – from the latents to the image features, while the latter enables both top-down and bottom up connections between these two variable groups.

the network maintains a scalable linear efficiency across all layers, realizing the transformer full potential. In doing so, we seek to take a step forward in tackling the challenging task of compositional scene generation. Intuitively, and as is later corroborated by our findings, holding multiple latent variables that interact through attention with the evolving generated image, may serve as a structural prior that promotes the formation of compact and compositional scene representations, as the different latents may specialize for certain objects or regions of interest. Indeed, as demonstrated in section 4, the Generative Adversarial Transformer achieves state-of-the-art performance in synthesizing both controlled and real-world indoor and outdoor scenes, while showing indications for semantic compositional disposition along the way.

In designing our model, we drew inspiration from multiple lines of research on generative modeling, compositionality and scene understanding, including techniques for scene decomposition, object discovery and representation learning. Several approaches, such as (Burgess et al., 2019; Greff et al., 2019; Eslami et al., 2016; Engelcke et al., 2020), perform iterative variational inference to encode scenes into multiple slots, but are mostly applied in the contexts of synthetic and oftentimes fairly rudimentary 2D settings. Works such as Capsule networks (Sabour et al., 2017) leverage ideas from psychology about Gestalt principles (Smith, 1988; Hamlyn, 2017), perceptual grouping (Brooks, 2015) or analysis-by-synthesis (Biederman, 1987), and like us, introduce ways to piece together visual elements to discover compound entities or, in the cases of Set Transformers (Lee et al., 2019) and A^2 -Nets (Chen et al., 2018), group local information into global aggregators, which proves useful for a broad spectrum of tasks, spanning unsupervised segmentation (Greff et al., 2017; Locatello et al., 2020), clustering (Lee et al., 2019), image recognition (Chen et al., 2018), NLP (Ravula et al., 2020) and viewpoint generalization (Kosiorrek et al., 2019). However, our work stands out incor-

porating new ways to integrate information between nodes, as well as novel forms of attention (*Simplex* and *Duplex*) that iteratively update and refine the assignments between image features and latents, and is the first to explore these techniques in the context of high-resolution generative modeling.

Most related to our work are certain GAN models for conditional and unconditional visual synthesis: A few methods (van Steenkiste et al., 2020; Gu et al., 2020; Nguyen-Phuoc et al., 2020; Ehrhardt et al., 2020) utilize multiple replicas of a generator to produce a set of image layers, that are then combined through alpha-composition. As a result, these models make quite strong assumptions about the independence between the components depicted in each layer. In contrast, our model generates one unified image through a cooperative process, coordinating between the different latents through the use of soft attention. Other works, such as SPADE (Park et al., 2019; Zhu et al., 2020), employ region-based feature modulation for the task of layout-to-image translation, but, contrary to us, use fixed segmentation maps or static class embeddings to control the visual features. Of particular relevance is the prominent StyleGAN model (Karras et al., 2019; 2020), which utilizes a single global style vector to consistently modulate the features of each layer. The GANsformer generalizes this design, as multiple style vectors impact different regions in the image concurrently, allowing for a spatially finer control over the generation process. Finally, while StyleGAN broadcasts information in one direction from the *global* latent to the *local* image features, our model propagates information both from latents to features and vice versa, enabling top-down and bottom-up reasoning to occur simultaneously¹.

¹Note however that our model certainly does not claim to serve as a biologically-accurate reflection of cognitive top-down processing. Rather, this analogy played as a conceptual source of inspiration that aided us through the idea development.

3. The Generative Adversarial Transformer

The Generative Adversarial Transformer is a type of Generative Adversarial Network, which involves a *generator* network (G) that maps a sample from the *latent* space to the output space (e.g. an image), and a *discriminator* network (D) which seeks to discern between real and fake samples (Goodfellow et al., 2014). The two networks compete with each other through a minimax game until reaching an equilibrium. Typically, each of these networks consists of multiple layers of convolution, but in the GANsformer case, we instead construct them using a novel architecture, called *Bipartite Transformer*, formally defined below.

The section is structured as follows: we first present a formulation of the Bipartite Transformer, a general-purpose generalization of the Transformer (section 3.1). Then, we provide an overview of how the transformer is incorporated into the generator and discriminator framework (section 3.2). We conclude by discussing the merits and distinctive properties of the GANsformer, that set it apart from the traditional GAN and transformer networks (section 3.3).

3.1. The Bipartite Transformer

The standard *transformer* network consists of alternating multi-head self-attention and feed-forward layers. We refer to each pair of self-attention and feed-forward layers as a transformer layer, such that a Transformer is considered to be a stack composed of several such layers. The Self-Attention operator considers all pairwise relations among the input elements, so to update each single element by attending to all the others. The Bipartite Transformer generalizes this formulation, featuring instead a bipartite graph between two groups of variables (in the GAN case, latents and image features). In the following, we consider two forms of attention that could be computed over the bipartite graph – Simplex attention, and Duplex attention, depending on the direction in which information propagates² – either in one way only, or both in top-down and bottom-up ways. While for clarity purposes we present the technique here in its one-head version, in practice we make use of a multi-head variant, in accordance with (Vaswani et al., 2017).

3.1.1. SIMPLEX ATTENTION

We begin by introducing the *simplex attention*, which distributes information in a single direction over the Bipartite Transformer. Formally, let $X^{n \times d}$ denote an input set of n vectors of dimension d (where, for the image case, $n = W \times H$), and $Y^{m \times d}$ denote a set of m aggregator variables (the latents, in the case of the generator). We can then compute attention over the derived bipartite graph between

²In computer networks, *simplex* refers to single direction communication, while *duplex* refers to communication in both ways.

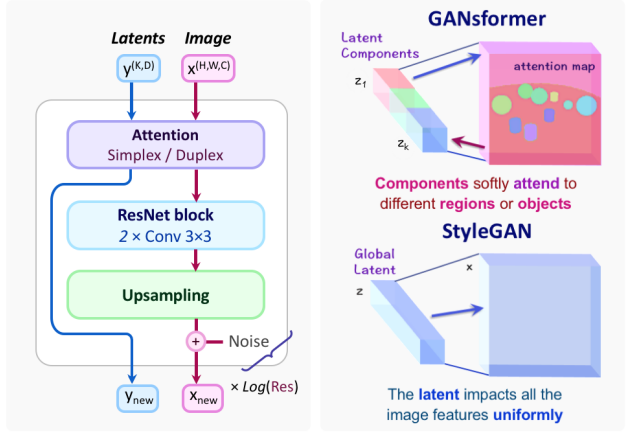


Figure 3. **Model Structure.** *Left:* The GANsformer layer, composed of a bipartite attention to propagate information from the latents to the evolving features grid, followed by a standard convolution and upsampling. These layers are stacked multiple times starting from a 4×4 grid, until a final high-resolution image is produced. *Right:* The latents and image features attend to each other to capture the scene compositionality. The GANsformer’s compositional latent space contrasts with the StyleGAN monolithic one, where a single latent modulates uniformly the whole scene.

these two groups of elements. Specifically, we then define:

$$Attention(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} \right) V$$

$$a_b(X, Y) = Attention(q(X), k(Y), v(Y))$$

Where a stands for *Attention*, and $q(\cdot), k(\cdot), v(\cdot)$ are functions that respectively map elements into queries, keys, and values, all maintaining the same dimensionality d . We also provide the query and key mappings with respective positional encoding inputs, to reflect the distinct position of each element in the set (e.g. in the image) (further details on the specifics of the positional encoding scheme in section 3.2).

We can then combine the attended information with the input elements X , but whereas the standard transformer implements an additive update rule of the form: $u^a(X, Y) = LayerNorm(X + a_b(X, Y))$ (where $Y = X$ in the standard self-attention case) we instead use the retrieved information to control both the *scale* as well as the *bias* of the elements in X , in line with the practices promoted by the StyleGAN model (Karras et al., 2019). As our experiments indicate, such multiplicative integration enables significant gains in the model performance. Formally:

$$u^s(X, Y) = \gamma(a_b(X, Y)) \odot \omega(X) + \beta(a_b(X, Y))$$

Where $\gamma(\cdot), \beta(\cdot)$ are mappings that compute multiplicative and additive styles (gain and bias), maintaining a dimension of d , and $\omega(X) = \frac{X - \mu(X)}{\sigma(X)}$ normalizes each element, with respect to the other features. This update rule fuses together




Figure 4. Samples of images generated by the GANsformer for the CLEVR, Bedroom and Cityscapes datasets, and a visualization of the produced attention maps. The different colors correspond to the latents that attend to each region.

the normalization and information propagation from Y to X , by essentially letting Y control X statistical tendencies of X , which for instance can be useful in the case of visual synthesis for generating particular objects or entities.

3.1.2. DUPLEX ATTENTION

We can go further and consider the variables Y to poses a key-value structure of their own (Miller et al., 2016): $Y = (K^{n \times d}, V^{n \times d})$, where the values store the *content* of the Y variables (e.g. the randomly sampled latents for the case of GANs) while the keys track the centroids of the attention-based assignments from X to Y , which can be computed as: $K = a_b(Y, X)$. Consequently, we can define a new update rule, that is later empirically shown to work more effectively than the simplex attention:

$$u^d(X, Y) = \gamma(a(X, K, V)) \odot \omega(X) + \beta(a(X, K, V))$$

This update compounds two attention operations on top of each other, where we first compute soft attention assignments between X and Y , and then refine the assignments by considering their centroids, analogously to the k-means algorithm (Lloyd, 1982; Locatello et al., 2020).

Finally, to support bidirectional interaction between the elements, we can chain two reciprocal simplex attentions from X to Y and from Y to X , obtaining the **duplex attention**, which alternates computing $Y := u^a(Y, X)$ and $X := u^d(X, Y)$, such that each representation is refined in light of its interaction with the other, integrating together bottom-up and top-down interactions.

3.1.3. OVERALL ARCHITECTURE STRUCTURE

Vision-Specific adaptations. In the standard Transformer used for NLP, each self-attention layer is followed by a Feed-Forward FC layer that processes each element independently

(which can be deemed a 1×1 convolution). Since our case pertains to images, we use instead a kernel size of $k = 3$ after each application of attention. We also apply a Leaky ReLU nonlinearity after each convolution (Maas et al., 2013) and then upsample or downsmample the features X , as part of the generator or discriminator respectively, following e.g. StyleGAN2 (Karras et al., 2020). To account for the features location within the image, we use a sinusoidal positional encoding along the horizontal and vertical dimensions for the visual features X (Vaswani et al., 2017), and a trained embedding for the set of latent variables Y .

Overall, the bipartite transformer is thus composed of a stack that alternates attention (simplex or duplex) and convolution layers, starting from a 4×4 grid up to the desirable resolution. Conceptually, this structure fosters an interesting communication flow: rather than densely modeling interactions among all the pairs of pixels in the images, it supports adaptive long-range interaction between far away pixels in a moderated manner, passing through through a compact and global bottleneck that selectively gathers information from the entire input and distribute it to relevant regions. Intuitively, this form can be viewed as analogous to the top-down notions discussed in section 1, as information is propagated in the two directions, both from the local pixel to the global high-level representation and vice versa.

We note that both the simplex and the duplex attention operations enjoy a bilinear efficiency of $\mathcal{O}(mn)$ thanks to the network’s bipartite structure that considers all pairs of corresponding elements from X and Y . Since, as we see below, we maintain Y to be of a fairly small size, choosing m in the range of 8–32, this compares favorably to the potentially prohibitive $\mathcal{O}(n^2)$ complexity of the self-attention, that impedes its applicability to high-resolution images.




Figure 5. Sample images and attention maps. Samples of images generated by the GANsformer for the CLEVR, LSUN-Bedroom and Cityscapes datasets, and a visualization of the produced attention maps. The different colors correspond to the latent variables that attend to each region. For the CLEVR dataset we should multiple attention maps in different layers of the model, revealing how the latent variables roles change over the different layers – while they correspond to different objects as the layout of the scene is being formed in early layers, they behave similarly to a surface normal in the final layers of the generator.

3.2. The Generator and Discriminator Networks

We use the celebrated StyleGAN model as a starting point for our GAN design. Commonly, a generator network consists of a multi-layer CNN that receives a randomly sampled vector z and transforms it into an image. The StyleGAN approach departs from this design and, instead, introduces a feed-forward *mapping network* that outputs an intermediate vector w , which in turn interacts directly with each convolution through the *synthesis network*, globally controlling the feature maps statistics of every layer.

Effectively, this approach attains *layer-wise decomposition* of visual properties, allowing the model to control particular global aspects of the image such as pose, lighting conditions or color schemes, in a coherent manner over the entire image. But while StyleGAN successfully disentangles global properties, it is potentially limited in its ability to perform *spatial decomposition*, as it provides no means to control the style of a localized regions within the generated image.

Luckily, the bipartite transformer offers a solution to meet

this goal. Instead of controlling the style of all features globally, we use instead our new attention layer to perform adaptive and local region-wise modulation. We split the latent vector z into k components, $z = [z_1, \dots, z_k]$ and, as in StyleGAN, pass each of them through a shared mapping network, obtaining a corresponding set of intermediate latent variables $Y = [y_1, \dots, y_k]$. Then, during synthesis, after each CNN layer in the generator, we let the feature map X and latents Y to play the roles of the two element groups, mediate their interaction through our new attention layer (either simplex or duplex). This setting thus allows for a flexible and dynamic style modulation at the region level. Since soft attention tends to group elements based on their proximity and content similarity (Vaswani et al., 2017), we see how the transformer architecture naturally fits into the generative task and proves useful in the visual domain, allowing the network to exercise finer control in modulating semantic regions. As we see in section 4, this capability turns to be especially useful in modeling compositional scenes.

For the discriminator, we similarly apply attention after ev-

Table 1. Comparison between the GANsformer and competing methods for image synthesis. We evaluate the models along commonly used metrics such as FID, Inception, and Precision & Recall scores. FID is considered to be the most well-received as a reliable indication of images fidelity and diversity. We compute each metric 10 times over 50k samples with different random seeds and report their average.

Model	CLEVR				LSUN-Bedroom			
	FID ↓	IS ↑	Precision ↑	Recall ↑	FID ↓	IS ↑	Precision ↑	Recall ↑
GAN	25.02	2.17	21.77	16.76	12.16	2.66	52.17	13.63
k-GAN	28.29	2.21	22.93	18.43	69.90	2.41	28.71	3.45
SAGAN	26.04	2.17	30.09	15.16	14.06	2.70	54.82	7.26
StyleGAN2	16.05	2.15	28.41	23.22	11.53	2.79	51.69	19.42
VQGAN	32.60	2.03	46.55	63.33	59.63	1.93	55.24	28.00
GANsformer _s	10.26	2.46	38.47	37.76	8.56	2.69	55.52	22.89
GANsformer _d	9.17	2.36	47.55	66.63	6.51	2.67	57.41	29.71

Model	FFHQ				Cityscapes			
	FID ↓	IS ↑	Precision ↑	Recall ↑	FID ↓	IS ↑	Precision ↑	Recall ↑
GAN	13.18	4.30	67.15	17.64	11.57	1.63	61.09	15.30
k-GAN	61.14	4.00	50.51	0.49	51.08	1.66	18.80	1.73
SAGAN	16.21	4.26	64.84	12.26	12.81	1.68	43.48	7.97
StyleGAN2	9.24	4.33	68.61	25.45	8.35	1.70	59.35	27.82
VQGAN	63.12	2.23	67.01	29.67	173.80	2.82	30.74	43.00
GANsformer _s	8.12	4.46	68.94	10.14	14.23	1.67	64.12	2.03
GANsformer _d	7.42	4.41	68.77	5.76	5.76	1.69	48.06	33.65

ery convolution, in this case using trained embeddings to initialize the aggregator variables Y , which may intuitively represent background knowledge the model learns about the scenes. At the last layer, we concatenate these variables to the final feature map to make a prediction about the identity of the image source. We note that this construction holds some resemblance to the PatchGAN discriminator introduced by (Isola et al., 2017), but whereas PatchGAN pools features according to a fixed predetermined scheme, the GANsformer can gather the information in a more adaptive and selective manner. Overall, using this structure endows the discriminator with the capacity to likewise model long-range dependencies, which can aid the discriminator in its assessment of the image fidelity, allowing it to acquire a more holistic understanding of the visual modality.

In terms of the loss function, optimization and training configuration, we adopt the settings and techniques used in the StyleGAN and StyleGAN2 models (Karras et al., 2019; 2020), including in particular style mixing, stochastic variation, exponential moving average for weights, and a non-saturating logistic loss with lazy R1 regularization.

3.3. Summary

To recapitulate the discussion above, the GANsformer successfully unifies the GANs and Transformer for the task of scene generation. Compared to the traditional GANs and transformers, it introduces multiple key innovations:

- **Bipartite Structure** that balances between expressiveness and efficiency, modeling long-range dependencies while maintaining linear computational costs.
- **Compositional Latent Space** with multiple variables that coordinate through attention to produce the image cooperatively, in a manner that matches the inherent compositionality of natural scenes.

- **Bidirectional Interaction** between the latents and the visual features, which allows the refinement and interpretation of each in light of the other.
- **Multiplicative Integration** rule to affect features’ style more flexibly, akin to StyleGAN but in contrast to the transformer architecture.

As we see in the following section, the combination of these design choices yields a strong architecture that demonstrates high efficiency, improved latent space disentanglement, and enhanced transparency of its generation process.

4. Experiments

We investigate the GANsformer through a suite of experiments to study its quantitative performance and qualitative behavior. As we will see below, the GANsformer achieves state-of-the-art results, successfully producing high-quality images for a varied assortment of datasets: FFHQ for human faces (Karras et al., 2019), CLEVR for multi-object scenes (Johnson et al., 2017), and the LSUN-Bedroom (Yu et al., 2015) and Cityscapes (Cordts et al., 2016) datasets for challenging indoor and outdoor scenes. The use of these datasets and their reproduced images are only for the purpose of scientific communication. Further analysis we then conduct in section 4.1, 4.2 and 4.3 provides evidence for several favorable properties the GANsformer possesses, including better data-efficiency, enhanced transparency, and stronger disentanglement compared to prior approaches. Section 4.4 then quantitatively assesses the network semantic coverage of the natural image distribution for the CLEVR dataset, while ablation studies in the supplementary empirically validate the relative importance of each of the model’s design choices. Taken altogether, our evaluation offers solid evidence for the GANsformer effectiveness and efficacy in modeling compositional images and scenes.

We compare our network with multiple related approaches including both baselines as well as leading models for image synthesis: (1) A baseline GAN (Goodfellow et al., 2014): a standard GAN that follows the typical convolutional architecture. (2) StyleGAN2 (Karras et al., 2020), where a single global latent interacts with the evolving image by modulating its style in each layer. (3) SAGAN (Zhang et al., 2019), a model that performs self-attention across all pixel pairs in the low-resolution layers of the generator and discriminator. (4) k-GAN (van Steenkiste et al., 2020) that produces k separated images, which are then blended through alpha-composition. (5) VQGAN (Esser et al., 2020) that has been proposed recently and utilizes transformers for discrete autoregressive auto-encoding.

To evaluate all models under comparable conditions of training scheme, model size, and optimization details, we implement them all within the codebase introduced by the StyleGAN authors. The only exception to that is the recent VQGAN model for which we use the official implementation. All models have been trained with images of 256×256 resolution and for the same number of training steps, roughly spanning a week on 2 NVIDIA V100 GPUs per model (or equivalently 3-4 days using 4 GPUs). For the GANsformer, we select k – the number of latent variables, from the range of 8–32. Note that increasing the value of k does not translate to increased overall latent dimension, and we rather kept the overall latent equal across models. See supplementary material A for further implementation details, hyperparameter settings and training configuration.

As shown in table 1, our model matches or outperforms prior work, achieving substantial gains in terms of FID score which correlates with image quality and diversity (Heusel et al., 2017), as well as other commonly used metrics such as Inception score (IS) and Precision and Recall (P&R).³ As could be expected, we obtain the least gains for the FFHQ human faces dataset, where naturally there is relatively lower diversity in images layout. On the flip side, most notable are the significant improvements in the performance for the CLEVR case, where our approach successfully lowers FID scores from 16.05 to 9.16, as well as the Bedroom dataset where the GANsformer nearly halves the FID score from 11.32 to 6.5, being trained for equal number of steps. These findings suggest that the GANsformer is particularly adept in modeling scenes of high compositionality (CLEVR) or layout diversity (Bedroom). Comparing

³Note that while the StyleGAN paper (Karras et al., 2020) reports lower FID scores in the FFHQ and Bedroom cases, they obtain them by training their model for 5-7 times longer than our experiments (StyleGAN models are trained for up to 17.5 million steps, producing 70M samples and demanding over 90 GPU-days). To comply with a reasonable compute budget, in our evaluation we equally reduced the training duration for all models, maintaining the same number of steps.

between the Simplex and Duplex Attention versions further reveals the strong benefits of integrating the reciprocal bottom-up and top-down processes together.

4.1. Data and Learning Efficiency

We examine the learning curves of our and competing models (figure 6, middle) and inspect samples of generated image at different stages of the training (figure 11 in the supplementary). These results both reveal that our model learns significantly faster than competing approaches, in the case of CLEVR producing high-quality images in approximately 3-times less training steps than the second-best approach. To explore the GANsformer learning aptitude further, we have performed experiments where we reduced the size of the dataset that each model (and specifically, its discriminator) is exposed to during the training (figure 6, rightmost) to varied degrees. These results similarly validate the model superior data-efficiency, especially when as little as 1k images are given to the model.

4.2. Transparency & Compositionality

To gain more insight into the model’s internal representation and its underlying generative process, we visualize the attention distributions produced by the GANsformer as it synthesizes new images. Recall that at each layer of the generator, it casts attention between the k latent variables and the evolving spatial features of the generated image.

From the samples in figures 4 and 5, we can see that particular latent variables tend to attend to coherent regions within the image in terms of content similarity and proximity. Figure 5 shows further visualizations of the attention computed by the model in various layers, showing how it behaves distinctively in different stages of the synthesis process. These visualizations imply that the latents carry a semantic sense, capturing objects, visual entities or constituent components of the synthesized scene. These findings can thereby attest for an enhanced compositionality that our model acquires through its multi-latent structure. Whereas models such as StyleGAN use a single monolithic latent vector to account for the whole scene and modulate features only at the global scale, our design lets the GANsformer exercise finer control in impacting features at the object granularity, while leveraging the use of attention to make its internal representations more explicit and transparent.

To quantify the compositionality level exhibited by the model, we use a pre-trained segmentor (Wu et al., 2019) to produce semantic segmentations for a sample set of generated scenes, and use them to measure the correlation between the attention cast by latent variables and various semantic classes. In figure 7 in the supplementary, we present the classes that on average have shown the highest correlation with respect to latent variables in the model, indicating

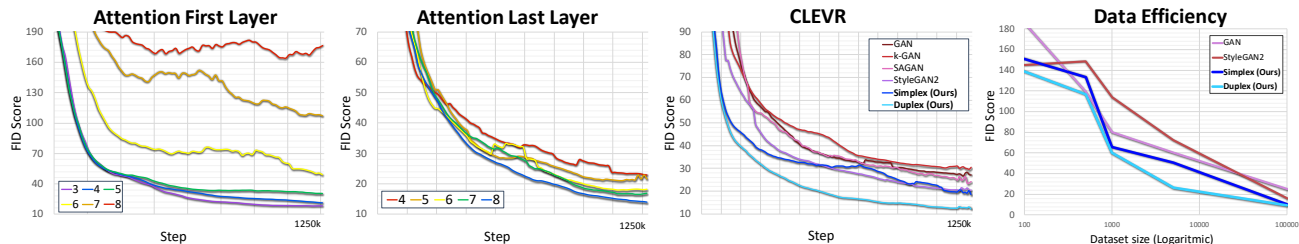


Figure 6. (1-2) Performance as a function of initial and final layers the attention is applied to. The more layers attention is applied to, the better the model’s performance get and the faster it learns, verifying the effectiveness of the GANsformer approach. (3) Learning curve for the GANsformer and competing approaches. (4): Data-efficiency for CLEVR.

that the model coherently attend to semantic concepts such as windows, pillows, sidewalks and cars, as well as coherent background regions like carpets, ceiling, and walls.

4.3. Disentanglement

We consider the DCI metrics commonly used in the disentanglement literature (Eastwood & Williams, 2018), to provide more evidence for the beneficial impact our architecture has on the model’s internal representations. These metrics assess the Disentanglement, Completeness and Informativeness of a given representation, essentially evaluating the degree to which there is 1-to-1 correspondence between latent factors and global image attributes. To obtain the attributes, we consider the area size of each semantic class (bed, carpet, pillows) obtained through a pre-trained segmentor and use them as the output response features for measuring the latent space disentanglement, computed over 1k images. We follow the protocol proposed by (Wu et al., 2020) and present the results in table 3. This analysis confirms that the GANsformer latent representations enjoy higher disentanglement when compared to the baseline StyleGAN approach.

4.4. Image Diversity

One of the advantages of compositional representations is that they can support combinatorial generalization – a key foundation of human intelligence (Battaglia et al., 2018). Inspired by this observation, we perform an experiment to measure that property in the context of visual synthesis of multi-object scenes. We use an object detector on the generated CLEVR scenes to extract the objects and properties within each sample. We then compute Chi-Square statistics on the sample set to determine the degree to which each model manages to covers the natural uniform distribution of CLEVR images. Table 2 summarizes the results, where we can see that our model obtains better scores across almost all the semantic properties of the image distribution. These metrics complement the common FID and IS scores as they emphasize structure over texture, focusing on object existence, arrangement and local properties, and thereby substantiating further the model compositionality.

Table 2. **Chi-Square statistics** of the generated scenes distribution for the CLEVR dataset, based on 1k samples. Samples were processed by a pre-trained object detector to identify objects and semantic attributes, to compute the properties’ distribution over the generated scenes.

	GAN	StyleGAN	GANsformer _s	GANsformer _d
Object Area	0.038	0.035	0.045	0.068
Object Number	2.378	1.622	2.142	2.825
Co-occurrence	13.532	9.177	9.506	13.020
Shape	1.334	0.643	1.856	2.815
Size	0.256	0.066	0.393	0.427
Material	0.108	0.322	1.573	2.887
Color	1.011	1.402	1.519	3.189
Class	6.435	4.571	5.315	16.742

Table 3. **Disentanglement metrics (DCI)**, which assess the Disentanglement, Completeness and Informativeness of latent representations, computed over 1k CLEVR images. The GANsformer achieves the strongest results compared to competing approaches.

	GAN	StyleGAN	GANsformer _s	GANsformer _d
Disentanglement	0.126	0.208	0.556	0.768
Modularity	0.631	0.703	0.891	0.952
Completeness	0.071	0.124	0.195	0.270
Informativeness	0.583	0.685	0.899	0.972
Informativeness’	0.434	0.332	0.848	0.963

5. Conclusion

We have introduced the GANsformer, a novel and efficient bipartite transformer that combines top-down and bottom-up interactions, and explored it for the task of generative modeling, achieving strong quantitative and qualitative results that attest for the model robustness and efficacy. The GANsformer fits within the general philosophy that aims to incorporate stronger inductive biases into neural networks to encourage desirable properties such as transparency, data-efficiency and compositionality – properties which are at the core of human intelligence, and serve as the basis for our capacity to reason, plan, learn, and imagine. While our work focuses on visual synthesis, we note that the bipartite transformer is a general-purpose model, and expect it may be found useful for other tasks in both vision and language. Overall, we hope that our work will help taking us a little closer in our collective search to bridge the gap between the intelligence of humans and machines.

References

- Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Irwan Bello, Barret Zoph, Quoc Le, Ashish Vaswani, and Jonathon Shlens. Attention augmented convolutional networks. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 3285–3294. IEEE, 2019. doi: 10.1109/ICCV.2019.00338.
- Irving Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94 (2):115, 1987.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Joseph L Brooks. Traditional and new principles of perceptual grouping. 2015.
- Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. MONet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pp. 213–229. Springer, 2020.
- Arantxa Casanova, Michal Drozdal, and Adriana Romero-Soriano. Generating unseen complex scenes: are we there yet? *arXiv preprint arXiv:2012.04027*, 2020.
- Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. A2-nets: Double attention networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 350–359, 2018.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Yunjey Choi, Min-Je Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 8789–8797. IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00916.
- Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the relationship between self-attention and convolutional layers. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 3213–3223. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.350.
- 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 Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.
- Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018.
- Sébastien Ehrhardt, Oliver Groth, Aron Monszpart, Martin Engelcke, Ingmar Posner, Niloy J. Mitra, and Andrea Vedaldi. RELATE: physically plausible multi-object scene synthesis using structured latent spaces. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in*

- Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Martin Engelcke, Adam R. Kosior, Oivi Parker Jones, and Ingmar Posner. GENESIS: generative scene inference and sampling with object-centric latent representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Koray Kavukcuoglu, and Geoffrey E. Hinton. Attend, infer, repeat: Fast scene understanding with generative models. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3225–3233, 2016.
- Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. *arXiv preprint arXiv:2012.09841*, 2020.
- Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pp. 3146–3154. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00326.
- Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- James J Gibson. A theory of direct visual perception. *Vision and Mind: selected readings in the philosophy of perception*, pp. 77–90, 2002.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.
- Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. Neural expectation maximization. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6691–6701, 2017.
- Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2424–2433. PMLR, 2019.
- Richard Langton Gregory. The intelligent eye. 1970.
- Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code GAN prior. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 3009–3018. IEEE, 2020. doi: 10.1109/CVPR42600.2020.00308.
- David Walter Hamlyn. *The psychology of perception: A philosophical examination of Gestalt theory and derivative theories of perception*, volume 13. Routledge, 2017.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6626–6637, 2017.
- Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.
- Han Hu, Zheng Zhang, Zhenda Xie, and Stephen Lin. Local relation networks for image recognition. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 3463–3472. IEEE, 2019. doi: 10.1109/ICCV.2019.00356.
- Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 603–612. IEEE, 2019. doi: 10.1109/ICCV.2019.00069.
- Monika Intaitė, Valdas Noreika, Alvydas Šoliūnas, and Christine M Falter. Interaction of bottom-up and top-down processes in the perception of ambiguous figures. *Vision Research*, 89:24–31, 2013.

- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 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*, pp. 5967–5976. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.632.
- Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two transformers can make one strong gan. *arXiv preprint arXiv:2102.07074*, 2021.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 1988–1997. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.215.
- Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1219–1228, 2018.
- Tero Karras, Samuli Laine, and Timo Aila. 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*, pp. 4401–4410. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00453.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 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*, pp. 8107–8116. IEEE, 2020. doi: 10.1109/CVPR42600.2020.00813.
- Adam R. Kosiorek, Sara Sabour, Yee Whye Teh, and Geoffrey E. Hinton. Stacked capsule autoencoders. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 15486–15496, 2019.
- Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 105–114. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.19.
- Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R. Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3744–3753. PMLR, 2019.
- Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *arXiv preprint arXiv:2006.15055*, 2020.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, pp. 3. Citeseer, 2013.
- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1400–1409, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1147.
- Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yong-Liang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. *arXiv preprint arXiv:2002.08988*, 2020.
- Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pp. 2337–2346. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00244.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4052–4061. PMLR, 2018.

- Niki Parmar, Prajit Ramachandran, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 68–80, 2019.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- Anirudh Ravula, Chris Alberti, Joshua Ainslie, Li Yang, Philip Minh Pham, Qifan Wang, Santiago Ontanon, Sumit Kumar Sanghai, Vaclav Cvicek, and Zach Fisher. Etc: Encoding long and structured inputs in transformers. 2020.
- Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic routing between capsules. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 3856–3866, 2017.
- Zhuoran Shen, Irwan Bello, Raviteja Vemulapalli, Xuhui Jia, and Ching-Hui Chen. Global self-attention networks for image recognition. *arXiv preprint arXiv:2010.03019*, 2020.
- Barry Smith. Foundations of gestalt theory. 1988.
- Sjoerd van Steenkiste, Karol Kurach, Jürgen Schmidhuber, and Sylvain Gelly. Investigating object compositionality in generative adversarial networks. *Neural Networks*, 130: 309–325, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008, 2017.
- Xiaolong Wang, Ross B. Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 7794–7803. IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00813.
- Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- Zongze Wu, Dani Lischinski, and Eli Shechtman. StyleSpace analysis: Disentangled controls for StyleGAN image generation. *arXiv preprint arXiv:2011.12799*, 2020.
- Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- Han Zhang, Ian J. Goodfellow, Dimitris N. Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7354–7363. PMLR, 2019.
- Hengshuang Zhao, Jiaya Jia, and Vladlen Koltun. Exploring self-attention for image recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 10073–10082. IEEE, 2020. doi: 10.1109/CVPR42600.2020.01009.
- Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. SEAN: image synthesis with semantic region-adaptive normalization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 5103–5112. IEEE, 2020. doi: 10.1109/CVPR42600.2020.00515.