DoodleFormer: Creative Sketch Drawing with Transformers

Ankan Kumar Bhunia¹, Salman Khan^{1,2}, Hisham Cholakkal¹, Rao Muhammad Anwer^{1,3}, Fahad Shahbaz Khan^{1,4}, Jorma Laaksonen³, Michael Felsberg⁴

¹ Mohamed bin Zayed University of AI, UAE ² Australian National University, Australia ³ Aalto University, Finland ⁴ Linköping University, Sweden ankan.bhunia@mbzuai.ac.ae



Fig. 1: Examples of visual creative sketches generated using the proposed DoodleFormer. Here, we show creative sketches generated based on (a) the random input strokes, (b) text inputs and (c) incomplete sketch images provided by the user. In all three scenarios, the generated sketches are well aligned with the user provided inputs (*e.g.*, the creative sketches generated for the text inputs "walking forward" and "fly up high" in (b)). Similarly, the diversity in terms of appearance, posture and part size can be observed within the generated creative bird sketches in (a). Furthermore, DoodleFormer accurately completes the missing bird wings, legs and beak in the bottom right example in (c). Additional examples are available in Fig. 2, Fig. 4, Fig. 6 and the supplementary.

Abstract. Creative sketching or doodling is an expressive activity, where imaginative and previously unseen depictions of everyday visual objects are drawn. Creative sketch image generation is a challenging vision problem, where the task is to generate diverse, yet realistic creative sketches possessing the unseen composition of the visual-world objects. Here, we propose a novel coarse-to-fine two-stage framework, DoodleFormer, that decomposes the creative sketch generation problem into the creation of coarse sketch composition followed by the incorporation of fine-details in the sketch. We introduce graph-aware transformer encoders that effectively capture global dynamic as well as local static structural relations among different body parts. To ensure diversity of the generated creative sketches, we introduce a probabilistic coarse sketch decoder that explicitly models the variations of each sketch body part to be drawn. Experiments are performed on two creative sketch datasets: Creative Birds and Creative Creatures. Our qualitative, quantitative and human-based evaluations show that DoodleFormer outperforms the state-of-the-art on both datasets, yielding realistic and diverse creative sketches. On Creative Creatures, DoodleFormer achieves an absolute gain of 25 in Frèchet inception distance (FID) over state-of-the-art. We also demonstrate the effectiveness of DoodleFormer for related applications of text to creative sketch generation, sketch completion and house layout generation. Code is available at: https://github.com/ankanbhunia/doodleformer



Fig. 2: A visual comparison of creative sketch images generated by DoodlerGAN [10] (top row) and the proposed DoodleFormer (bottom row) for the same initial random input strokes. We show examples from both Creative Birds (a) and Creative Creatures (b) datasets. DoodlerGAN suffers from topological artefacts (*e.g.*, more than one head like region in the third bird sketch from the left), disconnected body parts (*e.g.*, the fifth sketch from the left in creatures). Further, the DoodlerGAN generated creative sketches have lesser diversity in terms of size, appearance and posture. The proposed DoodleFormer alleviates the issues of topological artefacts and disconnected body parts, generating creative sketches that are more realistic and diverse.

1 Introduction

Humans have an outstanding ability to easily communicate and express abstract ideas and emotions through sketch drawings. Generally, a sketch comprises several strokes, where each stroke can be considered as a group of points. In automatic sketch image generation, the objective is to generate recognizable sketches that are closely related to the real-world visual concepts. Here, the focus is to learn more canonical and mundane interpretations of everyday objects.

Different from the standard sketch generation problem discussed above, *creative* sketch generation [10] involves drawing more imaginative and previously unseen depictions of everyday visual concepts (see Fig. 1(a)). In this problem, creative sketches are generated according to externally provided random input strokes. Example of creative sketch generation includes doodling activity, where diverse, yet recognizable sketch images are generated through unseen composition of everyday visual concepts. Automatic generation of creative sketches can largely assist human creative process *e.g.*, inspiring further ideas by providing a possible interpretation of initial sketches by the user. However, such a creative task is more challenging compared to mimicking real-world scenes in to sketch images. This work investigates the problem of creative sketch generation.

Recently, Ge *et al.* [10] address the creative sketch image generation problem by proposing a part-based Generative Adversarial Network called DoodlerGAN. It utilizes a part-specific generator to produce each body part of the sketch. The generated body parts are then sequentially integrated with the externally provided random input, for obtaining final sketch image. Although DoodlerGAN utilizes a part-specific generator for creating each body part of the sketch, it does not comprise an explicit mechanism to ensure that each body part is placed appropriately with respect to the rest of the parts. This leads to topological artifacts and connectivity issues (see Fig. 2). Further, DoodlerGAN struggles to generate diverse sketch images, which is an especially desired property in creative sketch generation.

In this work, we argue that the aforementioned problems of topological artefacts, connectivity and diversity issues can be alleviated by imitating the natural *coarse-to-fine* creative sketch drawing process, where the artist first draws the holistic coarse structure

of the sketch and then fills the fine-details to generate the final sketch. By first drawing the holistic coarse structure of the sketch aids to appropriately decide the location and the size of each sketch body part to be drawn. To imitate such a coarse-to-fine creative sketch generation process, we look into a two-stage framework where the global as well as local structural relations among different body parts can be first captured at a coarse-level followed by obtaining the fine-level sketch. The coarse-to-fine framework is expected to further improve the diversity of the creative sketch images by explicitly modeling the variations in the location and size of each sketch body part to be drawn.

1.1 Contributions

We propose a novel two-stage encoder-decoder framework, DoodleFormer, for creative sketch generation. DoodleFormer decomposes the creative sketch generation problem into the construction of holistic coarse sketch composition followed by injecting fine-details to generate final sketch image. To generate realistic sketch images, we introduce graph-aware transformer (GAT) encoders that effectively encode the local structural relations between different sketch body parts by integrating a static adjacency-based graph into the dynamic self-attention block. We further introduce a probabilistic coarse sketch decoder that utilizes Gaussian mixture models (GMMs) to obtain diverse locations of each body part, thereby improving the diversity of output sketches (see Fig. 2).

We evaluate the proposed DoodleFormer by conducting extensive qualitative, quantitative and human-based evaluations on the recently introduced Creative Birds and Creative Creatures datasets. Our DoodleFormer performs favorably against DoodlerGAN on all three evaluations. For instance, DoodleFormer sketches were interpreted to be drawn by a human 86%, having better strokes integration 85% and being more creative 82%, over DoodlerGAN in terms of human-based evaluation. Further, DoodleFormer outperforms DoodlerGAN with absolute gains of 25 and 23 in terms of Frèchet inception distance (FID) on Creative Creatures and Creative Birds, respectively. In addition to sketch generation based on externally provided random initial strokes, we validate the effectiveness of DoodleFormer to generate creative sketches based on text inputs, incomplete sketch images provided by user as well as generating complete house layouts given coarse-level bubble diagrams. DoodleFormer achieves impressive performance for text to sketch generation, sketch completion (see Fig. 1 (b) and (c)) as well as house layout generation (see Fig. 8).

2 Related Work

The problem of sketch generation [12,16,20,5,37,7] has been studied extensively in literature. These methods generally aim to mimic the visual world by capturing its important aspects in the generated sketches. SketchRNN [12], utilizes sequence-to-sequence Variational Autoencoder (VAE) for conditional and unconditional generation of vector sketches. Cao *et al.* [5] propose a generative model that generates multi-class sketches. Moreover, alternative strategies such as differentiable rendering [37], attention-based architectures [27] and reinforcement learning [3,38,9] have been investigated for sketch generation. The work of [7] incorporates a convolutional encoder to capture the spatial 4 A. K. Bhunia *et al*.

layout of sketches, whereas [23,29,25] aim at completing the missing parts of sketches. The work of [21] targets recovering the masked parts of points in sketches. A few works [27,21] have also studied related tasks of sketch classification and retrieval.

Different from the aforementioned standard sketch generation task, creative sketch generation has been recently explored [10]. This task focuses on drawing more imaginative and previously unseen depictions of common visual concepts rather than generating canonical and mundane interpretations of visual objects. To this end, DoodlerGAN [10] introduces a part-based Generative Adversarial Network built on StyleGAN2 [17] to sequentially produce each body part of the creative sketch image. Here, the part-based GAN model needs to be trained separately for individual body parts (eye, head, beak, *etc.*) using part annotations. However, such a separate model for each body part results in a large computational overhead. During inference, these individual part-based GAN models are sequentially used to generate their respective body parts within the creative sketch. While generating recognizable creative sketches, DoodlerGANs struggles with topological artifacts, connectivity and diversity issues. In this work, we set out to overcome these issues to generate diverse, yet realistic creative sketches.

3 Our Approach

Motivation: To motivate our framework, we first distinguish two desirable properties to be considered when designing an approach for creative sketch generation.

Holistic Sketch Part Composition: As discussed earlier, DoodlerGAN employs a partspecific generator to produce each body part of the sketch. However, it does not utilize any explicit mechanism to ensure that the generated part is placed in an appropriate location relative to other parts, thereby suffering from topological artifacts and connectivity issues (see Fig. 2). Here, we argue that explicitly capturing the holistic arrangement of the sketch parts is desired to generate realistic sketch images that avoid topological artifacts and connectivity issues.

Fine-level Diverse Sketch Generation: Creative sketches exhibit a large diversity in appearance, posing a major challenge when generating diverse, yet realistic fine-detailed sketch images. Existing work of DoodlerGAN struggles to generate diverse sketch images since it typically ignores the noise input in the sketch generation process [26]. Although DoodlerGAN attempts to partially address this issue by introducing heuristics in the form of randomly translating the input partial sketch, the diversity of generated sketch images is still far from satisfactory (see Fig. 2). Instead, we argue that having an explicit probabilistic modeling within the framework is expected to further improve the diversity of the generated sketch images.

Overall Framework: The proposed two-stage DoodleFormer framework combines the two aforementioned desired properties by decomposing the creative sketch generation problem to first capture the holistic coarse sketch composition and then injecting fine-details to generate the final sketch. The overall architecture of the proposed two-stage DoodleFormer is shown in Fig. 3. DoodleFormer comprises two stages: *Part Locator* (PL-Net) and *Part Sketcher* (PS-Net). The first stage, Part Locator (PL-Net), learns to explicitly capture the holistic arrangement of the sketch parts conditioned on the externally provided random initial stroke points *C* represented in a vector form. PL-Net



Fig. 3: The proposed DoodleFormer comprises two stages: Part Locator (PL-Net) and Part Sketcher (PS-Net). (a) The first stage, PL-Net, takes the initial stroke points C as the conditional input and learns to return the bounding boxes corresponding to each body part (coarse structure of the sketch) to be drawn. PL-Net contains two graph-aware transformer (GAT) encoders (E_b, E_c) and a probablistic coarse sketch decoder utilizing GMM modelling for the coarse box prediction. Within the decoder, the bounding box parameters are predicted by the location-predictor (\mathcal{H}_{xy}) and size-predictor (\mathcal{H}_{wh}) modules. (b) The second stage, PS-Net, then takes the predicted box locations along with C as inputs and generates the final sketch image \bar{I}_{im} . Following the design of E_b and E_c , PS-Net also comprises GAT block-based encoders (\bar{E}_b, \bar{E}_c) . Further, PS-Net contains a convolutional encoder-decoder network $(\mathcal{R}_E, \mathcal{R}_D)$ and a mask regressor to generate rasterized high quality sketch image \bar{I}_{im} .

comprises graph-aware transformer (GAT) block-based encoders to capture structural relationship between different regions within a sketch. To the best of our knowledge, we are the first to introduce a GAT block-based transformer encoder for the problem of creative sketch image generation. Instead of directly predicting the box parameters as deterministic points from the transformer decoder, we further introduce probabilistic coarse sketch decoders that utilize GMM modelling for box prediction. This enables our DoodleFormer to achieve diverse, yet plausible coarse structure (bounding boxes) for sketch generation. The second stage, Part Sketcher (PS-Net), creates the final sketch image with appropriate line segments based on the coarse structure obtained from PL-Net. PS-Net also comprises GAT block-based encoders, as in PL-Net, along with a convolutional encoder-decoder network to generate the final rasterized sketch image.

Our carefully designed two-stage DoodleFormer architecture possesses both desired properties (holistic sketch part composition as well as fine-level diverse sketch generation) and creates diverse, yet realistic sketch images in a coarse-to-fine manner (see Fig. 4). Next, we describe in detail PL-Net (Sec. 3.1) and PS-Net (Sec. 3.2).

3.1 Part Locator Network (PL-Net)

As discussed above, PL-Net takes the initial stroke points C as the conditional input, and learns to return a coarse structure capturing the holistic part composition of the desired sketch. The *encoders* in PL-Net contain graph-aware transformer (GAT) blocks to encode the structural relationship between different parts (holistic sketch part composition), leading to realistic sketch image generation. The *decoder* in PL-Net utilizes GMM modeling for box prediction, enabling the generations of diverse sketch images.

6

Initial strokes		١	^	٦	`	`	6	١	7	ι	C	7	C	7	J	ح	$\overline{\ }$	7	`	٢
Single-stage	P	A	\$	9	×.	\$ \$	Ø	.8	E.	Z	Ş	B	Ł.	ß	2	¢		Ķ	c p	â
Two-stage	\$	\mathcal{D}	Ş	e G	4 Day	B	Ð	Ð	\$	·S.	Ę.	Ð	Â,	\$	*	Ø	ß	ŝ	₹¢	Ø
+GAT	Ę.	Þ	S)	Ş2	Ð	Ę	\$	Þ	S.	\$	Ð	Ð	Ś	S	Ę	S.	S	Ð	Z	Ď
+GMM (Ours)	R	Ø.	Þ	23	Ğ.	Se .	F.	Ð	5		S	R	X	A	Ŷ	A	ÇR	Ś	S.	R

Fig. 4: A visual comparison in terms of progressively integrating one contribution at the time, from top to bottom, for common initial strokes. Compared to the single-stage baseline (first row), the two-stage framework (without the GAT block and probabilistic modeling in the decoder) generates sketch in a coarse-to-fine manner. As a result, the two-stage framework (second row) produces a more complete sketch where each body part is placed at an appropriate location relative to other parts. The introduction of GAT block (third row) in the encoders of the two-stage framework improves the realism of the generated sketches by capturing the structural relationship between different parts (*e.g.*, the tenth image from the left, where there is a discontinuity between the beak and the head of the bird). Further, the introduction of probabilistic modelling in the decoder of the two-stage framework (last row), improves the diversity (*e.g.*, appearance, size, orientation and posture) of generated sketch images. Our final two-stage framework (last row) produces realistic and diverse sketch images.

Graph-aware Transformer Block-based Encoder PL-Net consists of two graphaware transformer (GAT) block-based encoders E_b and E_c , which are used to obtain contextualized representation of the coarse (holistic) structure \mathcal{B} and the conditional input \mathcal{C} , respectively.

To encode the identity t of each body part present in a sketch, we define $v_t \in$ \mathbb{R}^d as a learned part embedding. We concatenate v_t with a feature representation obtained from $oldsymbol{b}_t \in \mathcal{B}$ (box location and size information (x_t, y_t, w_t, h_t) of each body part). This concatenated feature is then used as an input to the encoder E_b . The conditional input strokes C are passed through a linear layer before being input to the encoder E_c . We add special cls tokens [8] at the beginning of input sequences to the encoders (E_b and E_c). The output of this token is considered as the contextualized representation of the whole sequence. Further, we use fixed positional encodings to the input of each attention layer to retain information regarding the sequence order. Next, we introduce our GAT block used in both en-



Fig. 5: Our proposed graph-aware transformer (GAT) block (c) replaces the standard selfattention (b) in the conventional transformer encoder layer (a). Our GAT block injects the graph structure into self-attention by learning to reweight the attention matrix based on the pairwise relations between the graph nodes. In this way, the proposed GAT block combines the local connectivity patterns from the learned adjacency graph with the dynamic attention from the self-attention block.

coders (E_b, E_c) to encode the holistic structural composition of (sketch) body parts. **Graph-aware Transformer (GAT) Block:** The structure of our GAT block is shown in Fig. 5(c). Each GAT block consists of a graph-aware multi-headed self-attention (MHSA) module followed by a feed-forward network [18]. Given the queries Q, keys K, and values V the standard self-attention module [32] computes the attention according to the following equation (also shown in Fig. 5(b)),

$$\boldsymbol{\alpha} = \operatorname{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^{T}}{\sqrt{d}}\right). \tag{1}$$

While the standard self-attention module is effective towards learning highly contextualized feature representation, it does not explicitly emphasize on the local structural relation. However, creative sketches are structured inputs with definite connectivity patterns between sketch parts. To model this structure, we propose to encode an adjacency based graph implemented with spectral graph convolution [19]. Our proposed GAT block combines the definite connectivity patterns from the learned adjacency graph with the dynamic attention from self-attention block. Let us consider a graph where each node *i* is associated with a structure-aware representation n_i and corresponding neighbour set $\mathcal{N}_r(i)$. To represent the neighbor set $\mathcal{N}_r(i)$ for each node *i*, we define an adjacency matrix A where each entry represents whether two nodes *i* and *j* are adjacent. The edge weight e_{ij} between two adjacent nodes *i* and *j* is given by,

$$e_{ij} = \boldsymbol{W}_b^T \text{ReLU}\left(\boldsymbol{W}_a\left[\boldsymbol{n}_i, \boldsymbol{n}_j\right]\right) \ \forall j \in \mathcal{N}_r(i), \tag{2}$$

where W_a and W_b are learned parameters and $[\cdot, \cdot]$ is a concatenation operator. We set $e_{ij} = 0 \forall j \notin \mathcal{N}_r(i)$. For each GAT block l, the spectral graph convolution operation is,

$$\boldsymbol{n}_{i}^{(l+1)} = \operatorname{ReLU}\left(\boldsymbol{n}_{i}^{(l)} + \sum_{j \in \mathcal{N}_{r}(i)} e_{ij} \boldsymbol{W}_{c} \boldsymbol{n}_{j}^{(l)}\right),$$
(3)

where W_c is a learned matrix. Our main intuition is that the adjacency matrix representing the neighbourhood graph structure is static which is computed over the connected components in the graph and predetermined for each input, it is also symmetric and generally sparse. In contrast, attention learned from the self-attention layer is dynamic, can be dense and also non-symmetric. We propose to combine these two complementary representations through the following equation where we calculate the attention weight α_{ij} for nodes $j \in \mathcal{N}_r(i)$ as follows,

$$\alpha_{ij} = \frac{e_{ij} \exp(\varphi_{ij})}{\sum_{j \in \mathcal{N}_r(i)} e_{ij} \exp(\varphi_{ij})}, \text{ s.t. } \varphi_{ij} \in \frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d}},$$
(4)

where φ_{ij} is an element of the standard attention matrix.

The special token (cls) output from E_c is then utilized as an input to a Prior-Net for approximating the conditional prior latent distribution. Similarly, the cls token outputs of both E_b and E_c are provided as input to a Recog-Net for approximating the variational latent distribution. Both the Prior-Net and the Recog-Net are parameterized by multi-layer perceptrons (MLPs) to approximate prior and variational latent normal distributions. During training, we sample the latent variable z from the variational distribution and provide it as input to the probabilistic coarse sketch decoder. 8 A. K. Bhunia *et al*.

Probabilistic Coarse Sketch Decoder The probabilistic coarse sketch decoder within our PL-Net utilizes probabilistic modelling to generate diverse coarse structure. The decoder comprises two modules: a location-predictor \mathcal{H}_{xy} and a size-predictor \mathcal{H}_{wh} . Here, the location-predictor \mathcal{H}_{xy} estimates the center coordinates (x_t, y_t) of bounding boxes around body parts, while the size-predictor \mathcal{H}_{wh} predicts their width and height (w_t, h_t) . Both these modules consist of multi-headed self- and encoder-decoder attention mechanisms [32]. The encoder-decoder attention obtains the key and value vectors from the output of the encoder E_c . This allows every position in the decoder to attend to all positions in the conditional input sequence. The part embedding v_t from the encoder is used as a query positional encoding to each attention layer of the decoder. Over multiple consecutive decoding layers, the decoder modules produce respective output features $f_t^{xy} \in \mathbb{R}^d$ and $f_t^{wh} \in \mathbb{R}^d$ that lead to the distribution parameters of bounding boxes being associated with each body part, representing the coarse structure of the final sketch to be generated.

To enhance the diversity of generated sketch images, we model the box predictions from each decoder module by Gaussian Mixture Models (GMMs) [4,11]. Different from the conventional box prediction [6,39] that directly maps the decoder output features as deterministic box parameters, our GMM-based box prediction is modeled with M normal distributions $\mathcal{N}(\cdot)$ where each distribution is parameterized by θ_k and a mixture weight π_k ,

$$p(\boldsymbol{b}_t | \mathcal{C}, \boldsymbol{z}) = \sum_{k=1}^{M} \pi_{k,t} \mathcal{N}(\boldsymbol{b}_t; \theta_{k,t}), \text{ for } \sum_{k=1}^{M} \pi_{k,t} = 1.$$
(5)

The GMM parameters can be obtained by minimizing the negative log-likehood for all P body parts in a sketch,

$$\mathcal{L}_{b} = -\frac{1}{P} \sum_{t=1}^{P} \log \left(\sum_{k=1}^{M} \pi_{k,t} \mathcal{N}(\boldsymbol{b}_{t}; \boldsymbol{\theta}_{k,t}) \right).$$
(6)

Here, we simplify the quadvariate distribution of GMMs in Eq. 6 by decomposing it into two bivariate distributions as $p(\mathbf{b}_t | \mathcal{C}, \mathbf{z}) = p(x_t, y_t | \mathcal{C}, \mathbf{z})p(w_t, h_t | x_t, y_t, \mathcal{C}, \mathbf{z})$. The parameters of these bivariate GMMs are obtained by employing linear layers and appropriate normalization on the outputs f_t^{xy} , f_t^{wh} of \mathcal{H}_{xy} and \mathcal{H}_{wh} , respectively. In addition to GMM parameters, these linear layers also estimate the presence of a body part using an indicator variable, which is trained with a binary cross entropy loss \mathcal{L}_c .

PL-Net Loss function (\mathcal{L}_{PL}): The overall loss function \mathcal{L}_{PL} to train the PL-Net is the weighted sum of the reconstruction loss \mathcal{L}_{rec} , and the KL divergence loss \mathcal{L}_{KL} ,

$$\mathcal{L}_{PL} = \mathcal{L}_{rec} + \lambda_{KL} \mathcal{L}_{KL}.$$
(7)

Here, the reconstruction loss term is $\mathcal{L}_{rec} = \mathcal{L}_b + \mathcal{L}_c$. The KL divergence loss term \mathcal{L}_{KL} regularizes the variational distribution [28] from the Recog-Net to be closer to the prior distribution from the conditional Prior-Net, whereas λ_{KL} is a scalar loss weight.

Our carefully designed PL-Net architecture, presented above, provides a coarse structure of the sketch that is used to generate a diverse, yet realistic final sketch image in the second stage (PS-Net) of the proposed two-stage DoodleFormer framework.

Next, we present the PS-Net that takes the coarse structure of the sketch along with initial partial sketch C as inputs and generates the final sketch image.

3.2 Part Sketcher Network (PS-Net)

Our PS-Net comprises two graph-aware transformer (GAT) block-based encoders \bar{E}_b and \bar{E}_c , following the design of encoders E_b and E_c in the PL-Net. Here, the encoder \bar{E}_b produces a contextualized feature representation of bounding box b_t associated with each body part. Similarly, the encoder \bar{E}_c outputs a contextualized feature representation of initial stroke points C. Both these contextualized feature representations from \bar{E}_b and \bar{E}_c are then concatenated and passed through a linear layer to obtain u_t .

The initial stroke points C is converted to its raster form I_C and passed through a convolutional encoder \mathcal{R}_E that outputs a spatial representation $g = \mathcal{R}_E(I_C)$. Consequently, g and $\{u_t\}_{t=1}^P$ are provided as input to a convolutional decoder \mathcal{R}_D for generating the final sketch image \bar{I}_{im} ,

$$\bar{\boldsymbol{I}}_{im} = \mathcal{R}_D\left(\boldsymbol{g}, \{\boldsymbol{u}_t\}_{t=1}^P\right). \tag{8}$$

The decoder network \mathcal{R}_D utilizes the ResNet [13] architecture as a backbone. To introduce diversity in the generated images, a zero-mean unit-variance multivariate random noise is added with g before passing it to the decoder network. For fine-grained shape prediction, we utilize a mask regressor [30,31] having up-sampling convolutions, followed by sigmoid transformation to generate an auxiliary mask for each bounding box. The predicted masks are resized to the sizes of corresponding bounding boxes, which are then used to compute the instance-specific and structure-aware affine transformation parameters in the normalization layer of the decoder \mathcal{R}_D .

The training of PS-Net follows the standard GAN formulation where the PS-Net generator \mathcal{G} is followed by additional discriminator networks \mathcal{D}_{im} , \mathcal{D}_{part} , and \mathcal{D}_{app} to obtain image-level (\mathcal{L}_{im}), part-level (\mathcal{L}_{part}), and appearance (\mathcal{L}_{app}) adversarial losses [14,30], respectively. The loss function is then given by,

$$\mathcal{L}_{PS} = \mathcal{L}_{im} + \lambda_p \mathcal{L}_{part} + \lambda_a \mathcal{L}_{app},\tag{9}$$

where λ_p and λ_a are the loss weight hyper-parameters.

The introduction of the GAT block in the PL-Net and PS-Net encoders contributes towards the generation of realistic sketch images, whereas the effective utilization of probabilistic modelling in the PL-Net decoder leads to improved diversity. In summary, our two-stage DoodleFormer generates diverse, yet realistic sketch images (see Fig. 4).

4 Experiments

Datasets: We perform extensive experiments on the recently introduced Creative Birds and Creative Creatures datasets [10]. The Creative Birds has 8067 sketches of birds, whereas the Creative Creatures contains 9097 sketches of various creatures. In both datasets, all sketches come with part annotations. Both datasets also contain free-form natural language phrase as a text description for each sketch.

Table 1: Comparison of DoodleFormer with DoodlerGAN [10], StyleGAN2 [17] and SketchRNN [12] in terms of Frèchet inception distance (FID), generation diversity (GD), characteristic score (CS) and semantic diversity score (SDS). Our DoodleFormer performs favorably against existing methods on both datasets.

Methods	Cr FID(↓)	eative Bird GD(†)	ls CS(↑)	$\begin{tabular}{ c c c c } \hline Creative Creatures \\ FID(\downarrow) & GD(\uparrow) & CS(\uparrow) & SDS(\uparrow) \end{tabular} \end{tabular}$					
Training Data	-	19.40	0.45	-	18.06	0.60	1.91		
SketchRNN [12]	82.17	17.29	0.18	54.12	16.11	0.48	1.34		
StyleGAN2 [17]	130.93	14.45	0.12	56.81	13.96	0.37	1.17		
DoodlerGAN [10]	39.95	16.33	0.69	43.94	14.57	0.55	1.45		
DoodleFormer (Ours)	16.45	18.33	0.55	18.71	16.89	0.56	1.78		

Implementation Details: As discussed, both PL-Net and PS-Net utilize graph-aware transformer (GAT) block-based encoders. In each GAT-block based encoder, we define an adjacency matrix A based on the connectivity patterns of the adjacency graph. Every pair of overlapping bounding boxes on a coarse structure is connected in the adjacency graph. Similarly, for initial strokes, the corresponding adjacency graph connects adjacent points on each single stroke. Each of these encoders consist of L=6 graph-aware transformer blocks. Here, each block comprises multi-headed attention having 8 heads. In the probabilistic coarse sketch decoder, the location-predictor and size-predictor utilize 3 self- and encoder-decoder attention layers. Further, we set the embedding size d=512. We augment the vector sketch images by applying small affine transformations and these vector sketches are converted to raster images of size 128×128 . Our DoodleFormer is trained as follows. In the first stage, for training PL-Net, we initially obtain the bounding boxes for all body parts in a sketch using the part annotations. The bounding boxes are normalized to values between 0 and 1. In the second stage, we train PS-Net using the raster sketch images and their corresponding ground-truth boxes. In both stages, the initial stroke points are provided in a vector form as a conditional input. In all experiments, we use a batch size of 32. The learning rate is set to $1e^{-4}$ and the loss weights λ_{KL} , λ_p , λ_a are set to 1, 10 and 10.

4.1 Quantitative and Qualitative Comparisons

We first present a comparison (Tab. 1) of our DoodleFormer with state-of-the-art approaches [12,33,10] on both Creative Birds and Creative Creatures. For a fair comparison, we evaluate all methods using two widely used metrics, namely Frèchet inception distance (FID) [15] and generation diversity (GD) [5], as in DoodlerGAN [10]. Here, the FID and GD scores are computed using an Inception model trained on the Quick-Draw3.8M dataset [34], that embeds the images onto a feature space [10]. Tab. 1 shows that DoodleFormer outperforms existing methods in terms of both FID and GD scores, on the two datasets. The higher GD score indicates the ability of DoodleFormer to generate diverse sketch images, whereas the lower FID score indicates the superior quality of its generated creative sketches.

Furthermore, similar to DoodlerGAN [10], we use two additional metrics: characteristic score (CS) and semantic diversity score (SDS). The CS metric evaluates how



Fig. 6: Applications of DoodleFormer. (a) Creative sketch generation based on random input strokes. (b) Creative sketch completion: Here, DoodleFormer accurately completes missing parts (*e.g.*, beak, head and body of bird is well connected in second row and column in (b)), compared to DoodlerGAN (DG). (c) Text to creative sketch generation: We compare DoodleFormer with AttnGAN [35] (AG), StackGAN [36] (SG). DoodleFormer produces sketches that are well aligned with user provided input texts. Best viewed zoomed in.

often a generated sketch is classified to be a bird (for Creative Birds) or creature (for Creative Creatures) by the Inception model trained on the QuickDraw3.8M dataset. The SDS metric measures the diversity of the sketches in terms of the different creature categories they represent. While the CS score can give us a basic understanding of the generation quality, it does not necessarily reflects the creative abilities of a model. For instance, if a model generates only canonical and mundane sketches of birds, then the generated sketches would more likely to be correctly classified by the trained Inception model. In that case, the CS score will still be high. In contrast, the SDS score is more reliable in measuring the diversity of the generated sketch images. Tab. 1 shows that DoodleFormer performs favourably against existing methods, in terms of SDS score, on both datasets. Fig. 6(a) shows a visual comparison of DoodleFormer with Doodler-GAN for creative sketch generation¹.

4.2 User Study

Here, we present our user study to evaluate the human plausibility of creative sketches generated by our DoodleFormer. Specifically, we show 100 participants pairs of sketches – one generated by DoodleFormer and the other by a competing approach. For each pair of images, similar to DoodlerGAN [10], each participant is provided with 5 questions which are shown in the legend of Fig. 7 (a-e). DoodleFormer performs favorably against DoodlerGAN for all five questions on both datasets. For instance, Doodle-Former sketches were interpreted to be drawn by a human 86%, having better initial strokes integration 85% and being more creative 82%, over DoodlerGAN on Creative

11

Birds dataset. Further, for all the five questions, the DoodleFormer generated sketch images were found to be comparable with the human drawn sketches in Creative Datasets¹.

4.3 Ablation Study

We perform multiple ablation studies to validate the impact of proposed contributions in our framework. Tab. 2 shows the impact of two-stage framework, GAT blocks and GMM-based modeling on Creative Birds. Our single-stage baseline (referred as baseline*) is a standard transformer-based encoder-decoder architecture, where initial strokes are given as input to the transformer encoder. The decoder sequentially generates all body parts which are then integrated to obtain final output sketch. The generated sketches using baseline* are unrealistic and suffer from body parts misplace-



Fig. 7: User study results on Creative Birds (left) and Creative Creatures (right) based on the five questions (a-e) mentioned in the legend. Higher values indicate DoodleFormer is preferred more often over the compared approaches (Doodler-GAN and human drawn Creative datasets).

ment. The introduction of two-stage framework leads to an absolute gain of 26.3 in terms of FID score, highlighting the importance of a coarse-to-fine framework for realistic creative sketch generation. Our two-stage framework baseline neither uses the GAT block in encoders nor employs the GMM-based modeling in decoder. Instead of GMM-based modeling, we use a deterministic L1 loss in first stage (PL-Net) of two-stage baseline. While this two-stage baseline improves realism of generated sketches, it stills suffers from topological artifacts. The introduction of GAT block in encoders of two-stage baseline improves realism of generated sketches by capturing the structural relationship between different parts. Although GAT blocks improve FID score by a margin of 3.25, the generated sketches still lack diversity as indicated by only a marginal change in GD score. The introduction of GMM-based modelling in decoder improves diversity (*e.g.*, appearance, size, orientation and posture) of generated sketches that leads to an absolute gain of 1.17 in GD (see also Fig. 4). Our final DoodleFormer (two-stage baseline + GAT blocks + GMM) achieves absolute gains of 30.0 and 3.7 in terms of FID score over baseline* and two-stage baseline, respectively.

We also evaluate the design choices of our GAT blocks (see Tab. 3). First, we replace GAT blocks in encoders with simple GCN layers. The standard transformer-based encoders outperforms this GCN-based baseline by a margin of 7.31. Further, we adapt the Mesh Graphormer [22] by using their Graphormer in encoders of our framework. Mesh Graphormer stacks transformer encoder layer and GCN block together in series. In our experiments, we observe this design based on loosely connected components performs slightly worse than standard transformer-based baseline. In contrast, an integrated design like ours performs comparatively better.

¹ Additional details and results are provided in supplementary material.

Table 2: Impact of our two-stage framework, GAT blocks and GMM-based probabilistic modelling on Creative Birds.

Table 3: Comparison of alternative design										
choices	for	the	proposed	GAT	blocks	on				
Creative	Bire	ds.								

Design Choices	Methods	$\text{FID}(\downarrow)$	$\text{GD}(\uparrow)$	
Single-stage	baseline*	46.45	16.87	
Two-stage	baseline baseline + GAT baseline + GAT + GMM	20.14 16.89 16.45	17.05 17.16 18.33	-

 Methods
 FID(↓)
 GD(↑)

 GCN layers
 27.45
 17.23

 Transformer layers
 20.14
 17.05

 Mesh Graphormer [22]
 20.34
 16.78

 GAT layers (ours)
 16.45
 18.33

Table 4: **House Layout Generation:** We compare our approach with the existing methods in terms of FID and Compatibility scores (obtained by the graph edit distance). The dataset samples are split into five groups based on the room counts (1-3, 4-6, 7-9, 10-12, and 13+).

Mathada			FID (J	.)	Compatibility (\downarrow)					
Methods	1-3	4-6	7-9	10-12	13+	1-3	4-6	7-9	10-12	13+
Ashual et al. [12]	64.0	92.2	87.6	122.8	149.9	0.2	2.7	6.2	19.2	36.0
Johnson et al. [17]	69.8	86.9	80.1	117.5	123.2	0.2	2.6	5.2	17.5	29.3
House-GAN [10]	13.6	9.4	14.4	11.6	20.1	0.1	1.1	2.9	3.9	10.8
Ours	9.6	10.1	11.2	9.7	18.2	0.1	1.0	2.1	2.4	8.3

4.4 Related Applications

We also analyze DoodleFormer on three related tasks: user provided text to creative sketch generation, creative sketch completion and house layout generation.

Text to Creative Sketch Generation: Here, the text description is given as conditional input to encoder E_c in PL-Net, yielding a coarse structure of desired sketch which is fed to PS-Net to generate final sketch. We remove \mathcal{R}_E from the PS-Net, and the *cls* token output from the encoder \bar{E}_c is directly passed as input to \mathcal{R}_D . We use 80%-20% train-test split. We compare DoodleFormer with two popular text-to-image methods: StackGAN [36], AttnGAN [35] on Creative Birds and Creative Creatures. DoodleFormer performs favorably against these methods in terms of FID and GD scores on both datasets. On Creative Birds, StackGAN, AttnGAN and DoodleFormer achieve respective FID scores of [53.1, 45.2, **18.5**], and GD scores of [16.7, 16.5, **17.3**]. Fig. 6 (c) shows a qualitative comparison¹.

Creative Sketch Completion: Given an incomplete sketch as input, DoodleFormer attempts to creatively complete the rest of the sketch. First, PL-Net obtains the bounding boxes for missing parts. Then, PS-Net generates an image containing the required missing parts which is then integrated with the incomplete sketch input to obtain the final output. On both Creative Birds and Creative Creatures, DoodleFormer achieves favorable results compared to DoodlerGAN in terms of FID and GD scores. On Creative Birds, DoodlerGAN and DoodleFormer achieve respective FID scores of [44.2 and **18.3**] and GD scores of [15.1 and **17.8**]. Fig. 6 (b) shows the qualitative comparison¹.

House Layout Generation: Finally, we use our proposed PL-Net architecture for the house-plan generation [24] task. The goal is to take a bubble diagram as an input, and generate a diverse set of realistic and compatible house layouts. A bubble diagram is represented by a graph where each node contains information about rooms and edges

13



Fig. 8: Qualitative results of House Layout Generation. Given the input bubble diagram, We compare the house layout sample generated using our method with the House-GAN [24]. Our method produces house layouts that are well aligned with input bubble diagram texts.

indicate their spatial adjacency. The output house layout is represented as axis-aligned bounding boxes. The Encoder E_c takes the room type information as input and Probabilistic Decoder subsequently outputs boundary boxes for each room. To transform the obtained boundary box layout to a floor plan layout, we employ a floor-plan postprocessing strategy. In this process, we first extract boundary lines of the generated boundary boxes. Next, we merge the adjacent line segments together and further align them to obtain a closed polygon.

We perform the house-plan generation experiments on LIFULL HOME's dataset [1]. For fair comparison, we follow the same setting used by House-GAN [24]. We divide the samples into five groups based on the number of rooms: 1-3, 4-6, 7-9, 10-12, and 13+. To test the generalization ability in each group, we train a model while excluding samples in the same group. At test time, we randomly pick a bubble input diagram from each group and generate 10 samples. Similar to House-GAN [24], we quantitatively measure the performance of our method in terms of FID and compatibility scores. The compatibility score is the graph editing distance [2] between the input bubble diagram and the bubble diagram constructed from the output layout. Tab. 4 shows that our method outperforms existing house-plan generation methods both in terms of FID and compatibility scores. Fig. 8 shows the qualitative comparison of our house layout generation approach with House-GAN [24].

5 Conclusion

We proposed a novel coarse-to-fine two-stage approach, DoodleFormer, for creative sketch generation. We introduce graph-aware transformer encoders that effectively capture global dynamic as well as local static structural relations among different body parts. To ensure diversity of generated creative sketches, we introduce a probabilistic coarse sketch decoder that explicitly models variations of each sketch body part to be drawn. We show the effectiveness of DoodleFormer on two datasets by performing extensive qualitative, quantitative and human-based evaluations. In addition, we demonstrate promising results on related applications such as text to creative sketch generation, sketch completion and house layout generation.

References

- Lifull home's dataset. https://www.nii.ac.jp/dsc/idr/lifull, accessed: 2010-09-30 14
- Abu-Aisheh, Z., Raveaux, R., Ramel, J.Y., Martineau, P.: An exact graph edit distance algorithm for solving pattern recognition problems. In: 4th International Conference on Pattern Recognition Applications and Methods 2015 (2015) 14
- Balasubramanian, S., Balasubramanian, V.N., et al.: Teaching gans to sketch in vector format. arXiv preprint arXiv:1904.03620 (2019) 3
- 4. Bishop, C.M.: Mixture density networks. Aston University (1994) 8
- 5. Cao, N., Yan, X., Shi, Y., Chen, C.: Ai-sketcher: a deep generative model for producing high-quality sketches. In: AAAI (2019) 3, 10
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020) 8
- Chen, Y., Tu, S., Yi, Y., Xu, L.: Sketch-pix2seq: a model to generate sketches of multiple categories. arXiv preprint arXiv:1709.04121 (2017) 3
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL (2019) 6
- Ganin, Y., Kulkarni, T., Babuschkin, I., Eslami, S.A., Vinyals, O.: Synthesizing programs for images using reinforced adversarial learning. In: ICML (2018) 3
- Ge, S., Goswami, V., Zitnick, C.L., Parikh, D.: Creative sketch generation. In: ICLR (2021) 2, 4, 9, 10, 11, 13
- 11. Graves, A.: Generating sequences with recurrent neural networks. arXiv preprint arXiv:1308.0850 (2013) 8
- 12. Ha, D., Eck, D.: A neural representation of sketch drawings. In: ICLR (2018) 3, 10, 13
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) 9
- He, S., Liao, W., Yang, M.Y., Yang, Y., Song, Y.Z., Rosenhahn, B., Xiang, T.: Context-aware layout to image generation with enhanced object appearance. In: CVPR (2021) 9
- 15. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) 10
- Hinton, G.E., Nair, V.: Inferring motor programs from images of handwritten digits. In: NeurIPS (2006) 3
- 17. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of StyleGAN. In: CVPR (2020) 4, 10, 13
- Khan, S., Naseer, M., Hayat, M., Zamir, S.W., Khan, F.S., Shah, M.: Transformers in vision: A survey. arXiv preprint arXiv:2101.01169 (2021) 7
- Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017) 7
- Li, Y., Song, Y.Z., Hospedales, T.M., Gong, S.: Free-hand sketch synthesis with deformable stroke models. In: IJCV (2017) 3
- Lin, H., Fu, Y., Xue, X., Jiang, Y.G.: Sketch-bert: Learning sketch bidirectional encoder representation from transformers by self-supervised learning of sketch gestalt. In: CVPR (2020) 4
- 22. Lin, K., Wang, L., Liu, Z.: Mesh graphormer. In: ICCV (2021) 12, 13
- Liu, F., Deng, X., Lai, Y.K., Liu, Y.J., Ma, C., Wang, H.: Sketchgan: Joint sketch completion and recognition with gan. In: CVPR (2019) 4
- Nauata, N., Chang, K.H., Cheng, C.Y., Mori, G., Furukawa, Y.: House-gan: Relational generative adversarial networks for graph-constrained house layout generation. In: European Conference on Computer Vision. pp. 162–177. Springer (2020) 13, 14

- 16 A. K. Bhunia *et al*.
- Qi, Y., Su, G., Chowdhury, P.N., Li, M., Song, Y.Z.: Sketchlattice: Latticed representation for sketch manipulation. In: ICCV (2021) 4
- 26. Ramasinghe, S., Farazi, M., Khan, S., Barnes, N., Gould, S.: Rethinking conditional gan training: An approach using geometrically structured latent manifolds. In: NeurIPS (2021) 4
- Ribeiro, L.S.F., Bui, T., Collomosse, J., Ponti, M.: Sketchformer: Transformer-based representation for sketched structure. In: CVPR (2020) 3, 4
- Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: NeurIPS (2015) 8
- 29. Su, G., Qi, Y., Pang, K., Yang, J., Song, Y.Z.: Sketchhealer: A graph-to-sequence network for recreating partial human sketches. In: BMVC (2020) 4
- 30. Sun, W., Wu, T.: Image synthesis from reconfigurable layout and style. In: ICCV (2019) 9
- Sun, W., Wu, T.: Learning layout and style reconfigurable gans for controllable image synthesis. PAMI (2021) 9
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017) 7, 8
- Viazovetskyi, Y., Ivashkin, V., Kashin, E.: Stylegan2 distillation for feed-forward image manipulation. In: ECCV (2020) 10
- Xu, P., Hospedales, T.M., Yin, Q., Song, Y.Z., Xiang, T., Wang, L.: Deep learning for freehand sketch: A survey and a toolbox. arXiv preprint arXiv:2001.02600 (2020) 10
- Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X.: Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In: CVPR (2018) 11, 13
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV (2017) 11, 13
- Zheng, N., Jiang, Y., Huang, D.: Strokenet: A neural painting environment. In: ICLR (2018)
 3
- Zhou, T., Fang, C., Wang, Z., Yang, J., Kim, B., Chen, Z., Brandt, J., Terzopoulos, D.: Learning to doodle with stroke demonstrations and deep q-networks. In: BMVC (2018) 3
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: ICLR (2021) 8