THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Universal Sketch Perceptual Grouping

Citation for published version:

Li, K, Pang, K, Song, J, Song, Y-Z, Xiang, T, Hospedales, T & Zhang, H 2018, Universal Sketch Perceptual Grouping. in *European Conference on Computer Vision (ECCV 2018)*. Lecture Notes in Computer Science, vol. 11212, Image Processing, Computer Vision, Pattern Recognition, and Graphics, vol. 11212, Springer, Cham, Munich, Germany, pp. 593-609, European Conference on Computer Vision 2018, Munich, Germany, 8/09/18. https://doi.org/10.1007/978-3-030-01237-3_36

Digital Object Identifier (DOI):

[10.1007/978-3-030-01237-3_36](https://doi.org/10.1007/978-3-030-01237-3_36)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

European Conference on Computer Vision (ECCV 2018)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Universal Sketch Perceptual Grouping

Ke Li^{1,2}, Kaiyue Pang², Jifei Song², Yi-Zhe Song², Tao Xiang², Timothy M. Hospedales³, and Honggang Zhang¹

¹Beijing University of Posts and Telecommunications

²SketchX, Queen Mary University of London, ³The University of Edinburgh

{like1990,zhhg}@bupt.edu.cn, {t.hospedales}@ed.ac.uk

{kaiyue.pang,j.song,yizhe.song,t.xiang}@qmul.ac.uk

Fig. 1: Examples of the SPG dataset. Stroke groups are colour coded.

Abstract. In this work we aim to develop a universal sketch grouper. That is, a grouper that can be applied to sketches of any category in any domain to group constituent strokes/segments into semantically meaningful object parts. The first obstacle to this goal is the lack of large-scale datasets with grouping annotation. To overcome this, we contribute the largest sketch perceptual grouping (SPG) dataset to date, consisting of 20,000 unique sketches evenly distributed over 25 object categories. Furthermore, we propose a novel deep universal perceptual grouping model. The model is learned with both generative and discriminative losses. The generative losses improve the generalisation ability of the model to unseen object categories and datasets. The discriminative losses include a local grouping loss and a novel global grouping loss to enforce global grouping consistency. We show that the proposed model significantly outperforms the state-of-the-art groupers. Further, we show that our grouper is useful for a number of sketch analysis tasks including sketch synthesis and fine-grained sketch-based image retrieval (FG-SBIR).

Keywords: Sketch Perceptual grouping, Universal grouper, Deep grouping model, Dataset.

1 Introduction

Humans effortlessly detect objects and object parts out of a cluttered background. The Gestalt school of psychologists [1, 2] argued that this ability to perceptually group visual cues/patterns into objects is built upon a number of

grouping principles, termed Gestalt laws of grouping. These include five categories, namely proximity, similarity, continuity, closure, and symmetry [3].

Computer vision research in grouping or segmentation has long exploited these perceptual grouping principles. For example, in image segmentation [4, 5, 6, 7], pixel visual appearance similarity and local proximity are often used to group pixels into objects. These principles are exploited by the human visual system to robustly perform perceptual grouping in diverse contexts and for diverse object categories. Exploiting them is important for a universal grouping algorithm.

We aim to develop such a universal grouper for human free-hand sketches which takes a sketch as input and groups the constituent strokes into semantic parts. Note that this is different from semantic segmentation for either photos [7] or sketches [8, 9, 10], where each segmented part is given a label, and the labels are often object category-dependent. Only the group relationship between strokes is predicted so that the grouper can *universally* be applied to any object category.

Very few existing studies [11, 12] investigate sketch perceptual grouping. These approaches compute hand-crafted features from each stroke and use the proximity and continuity principles to compute a stroke affinity matrix for subsequent clustering/grouping. They thus have a number of limitations: (i) Only two out of the five principles are exploited, while the unused ones such as closure are clearly useful in grouping human sketches which can be fragmented (see Fig. 1). (ii) How the principles are formulated is determined manually rather than learned from data. (iii) Fixed weightings of different principles are used which are either manually set [11] or learned [12]. However, for different sketches, different principles could be used by humans with different weightings. Therefore a more dynamic sketch-specific grouping strategy is preferable. Nevertheless, the existing sketch perceptual grouping datasets [8, 12] are extremely small, containing 2,000 annotated sketches at most. This prevents more powerful and flexible deep neural network models from being developed.

The first contribution of this paper is to provide the first large-scale sketch perceptual grouping (SPG) dataset consisting of 20,000 sketches with ground truth grouping annotation, i.e., 10 times larger than the largest dataset to date [12]. The sketches are collected from 25 representative object categories with 800 sketches per category. Some examples of the sketches and their annotation can be seen in Fig. 1. A dataset of such size makes the development of a deep universal grouper possible.

Even with sufficient training samples, learning a deep universal sketch grouper is non-trivial. In particular, there are two challenges: how to make the deep grouper generalisable to unseen object categories and domains without any training data from them; and how to design training losses that enforce both local (stroke pairwise) grouping consistency and global (whole sketch level) grouping consistency given variable number of strokes per sketch. Most losses used by existing deep models are for supervised classification tasks; grouping is closer to clustering than classification so few options exist.

In this paper, we propose a novel deep sketch grouping model to overcome both challenges. Specifically, treating a sketch as a sequence of strokes/segments,

our model is a sequence-to-sequence variational autoencoder (VAE). The reconstruction loss in this deep generative model forces the learned representation to preserve information richer than required for the discriminative grouping task alone. This has been proven to be useful for improving model generalisation ability [13], critical for making the grouper universal. As for the discriminative grouping learning objectives, we deploy two losses: a pairwise stroke grouping loss enforcing local grouping consistency and a global grouping loss to enforce global grouping consistency. This separation of the local and global grouping losses enables us to balance the two and makes our model more robust against annotation noise. Based on the proposed grouper we develop a simple sketch synthesis model by grouping and abstracting photo edge maps. The synthesised sketches can be used to learn a strong *unsupervised* fine-grained sketch-based image retrieval (FG-SBIR) model, i.e., using photos only.

Our contributions are as follows: (1) We contribute the largest sketch perceptual grouping dataset to date with extensive human annotation. To drive future research, we will make the dataset publicly available. (2) For the first time, a deep universal sketch grouper is developed based on a novel deep sequence-to-sequence VAE with both generative and discriminative losses. (3) Extensive experiments show the superiority of our grouper against existing ones, especially when evaluated on new categories or new domains. Its usefulness on a number of sketch analysis tasks including sketch synthesis and FG-SBIR is also demonstrated.

2 Related work

Perceptual Grouping: Humans can easily extract salient visual structure from apparent noise. Gestalt psychologists referred to this phenomenon as perceptual organisation [1, 2] and introduced the concept of perceptual grouping, which accounts for the observation that humans naturally group visual patterns into objects. A set of simple Gestalt principles were further developed, including proximity, similarity and continuity [3], with closure, connectedness and common fate introduced later, primarily for studying human vision systems [4, 14].

Sketch Groupers: Very few studies exist on grouping sketch strokes into parts. The most related studies are [11, 12]. They compute an affinity matrix between strokes using hand-crafted features based on proximity and continuity principles. The two principles are combined with fixed weights learned from human annotated stroke groups. In contrast, we assume that when humans draw sketches and annotate them into groups, all grouping principles could be used. Importantly, using which ones and by how much are dependent on the specific sketch instance. Our model is thus a deep neural network that takes the sketch as input and aims to model all principles implicitly via both generative and discriminative grouping losses. It thus has the potential to perform principle selection and weighting dynamically according to the sketch input. We also provide a much larger dataset compared to the one provided in [12]. We show that on both datasets, our model outperforms that in [12] by a big margin. Note that perceptual grouping has been modelled using a deep autoencoder in [15].

However, the objective is to group discrete graphical patterns which has richer visual cues that make them more akin to the problem of image segmentation, and thus easier than grouping line drawings in sketches.

Sketch Semantic Segmentation: A closely related problem to sketch grouping is sketch semantic segmentation [8, 9, 10]¹. The key difference is that a sketch grouper is universal in that it can be applied to any object category as it only predicts whether strokes belong to the same group rather than what group. In contrast, sketch segmentation models need to predict the label of each group. As a result, typically one model is needed for each object category. Note that although two different problems are tackled, our work can potentially benefit sketch semantic segmentation in two ways: (i) The grouping principles modelled implicitly in our grouping model could be used for semantic segmentation, e.g., by modifying/fine-tuning our model to a fully supervised one. (ii) The SPG dataset also contains group ID labels for each category so can be used for developing deep segmentation models, which has not been possible to date due to the small sizes of existing sketch segmentation datasets [8, 9, 10].

Sketch Stroke Analysis: Like our model, a number of recent sketch models are based on stroke modelling. [10] studied stroke semantic segmentation. A sequence-to-sequence variational autoencoder is used in [16] for a different purpose of conditional sketch synthesis. The work in [17] uses a sketch RNN for sketch abstraction problem by sequentially removing redundant strokes. A stroke-based model is naturally suited for perceptual grouping. Modelling Gestalt principles is harder if a sketch is treated as a 2D pixel array instead of strokes.

Fine-grained SBIR: FG-SBIR has been a recent focus in sketch analysis [?, 18, 19, 20, 21, 22]. Training a FG-SBIR model typically requires expensive photo-sketch pair collection, which severely restricts its applicability to large number of object categories. In this work, we show that our universal grouper is general enough to be applied to edge maps computed from object photos. The edge maps can then be abstracted by removing the least important groups. The abstracted edge map can be used to substitute human sketches and form synthetic sketch-photo pairs for training a FG-SBIR model. We show that the performance of a model trained in this way approaches that of the same model trained with human labelled data, and is superior to the state-of-the-art unsupervised alternative [17].

3 Free-hand Sketch Perceptual Grouping Dataset

We contribute the Sketch Perceptual Grouping dataset, the largest free-hand sketch perceptual grouping dataset to date. It contains 20,000 sketches distributed over 25 categories with each sketch manually annotated into parts.

Category Selection The sketches come from the QuickDraw dataset [16], which is by far the largest sketch dataset. It contains 345 categories of everyday objects. Out of these, 25 are selected for SPG (see Table 1) based on following criteria: (i) Complexity: the category should contain at least three semantic parts,

¹ Their relationship is analogous to that between unsupervised image segmentation [5, 6] and semantic segmentation [7].

Fig. 2: Examples to illustrate our sketch selection process. See details in text.

meaning categories such as cloud and moon are out. (ii) **Variety**: The selected categories need to be sufficiently different from each other to be appropriate for testing the grouper’s generalisation ability to unseen classes. For example, only one of the four-legged animal classes is chosen.

Sketch Instance Selection: Each QuickDraw category contains at least 100,000 sketches. Annotating all of them is not feasible. So 800 sketches are chosen from each category. First, some quality screening is performed. Specifically, since all QuickDraw sketches were drawn within 20 seconds, there are a large number of badly drawn sketches that are unrecognisable by humans, making part grouping impossible. We thus first discard sketches which could not be recognised by an off-the-shelf sketch classifier [23]. The remaining sketches are then subject to the following instance selection criteria: (i) **Majority**: Sketches in each category can form subcategories which can be visually very different from each other. Only the sketches from the majority subcategory are selected. For example, the top row of Fig. 2 shows that most sketches from the alarm clock category belong to the “with hand” subcategory, whilst a small minority depicts digital clocks without hands. Only sketches from the former are selected. (ii) **Complexity**: Over-abstract sketches with less than three parts are removed. (iii) **Ambiguity**: Finally, we eliminate sketches that contain both the target object and other objects/background to avoid ambiguity of the object category. Fig. 2 shows examples of how these criteria are enforced during instance selection.

Annotation: Now given a sketch, each annotator is asked to group the strokes into groups. Each group has a semantic meaning and typically corresponds to an object part. So apart from the grouping label, the group ID is also annotated. Even though the group ID information is not used in our perceptual grouping model, it can be used when the task is sketch semantic segmentation. To obtain consistent grouping annotation, 25 annotators are recruited and each only annotates one category. Examples of the annotation can be seen in Fig. 1.

Fig. 3: A schematic of the proposed deep perceptual grouper.

4 Deep Universal Sketch Perceptual Grouper

4.1 Model Overview

Our deep sketch grouper is a variant of the sequence-to-sequence variational autoencoder (VAE) [24, 25]. As shown in Fig. 3, it is essentially a deep encoder-decoder with both the encoder and decoder being RNNs for modelling a sketch as a set of strokes. The encoder produces a global representation of the sketch, which is used as a condition for a variational decoder that aims to reconstruct the input sketch. Such sketch synthesis is a side task here. Our main aim is for the decoder to produce a representation of each stroke useful for grouping them. Once learned, the decoder should implicitly model all the grouping principles used by the annotators in producing the grouping labels, so that the learned stroke representation can be used to compute a stroke affinity matrix indicating the correct stroke grouping. To this end, the decoder has two branches: a generative branch to reconstruct the input sketch; and a discriminative branch that produces the discriminative stroke feature/affinity matrix.

4.2 Encoder and Decoder Architecture

Traditional perceptual grouping methods treat sketches as images composed of static pixels, thus neglecting the dependency between different segments and strokes (each stroke consists of a variable number of line segments). In our dataset, all the sketches are captured in a vectorised format, making sequential modelling of sketches possible. More specifically, we first represent a sketch as a sequence of N stroke-segments $[S_1, S_2, \dots, S_N]$. Each segment is a tuple $(\Delta x, \Delta y, p)$, where Δx and Δy denote the offsets along the horizontal and vertical directions respectively, while p represents the drawing state, following the same representation used for human handwriting [26].

With these stroke segments as inputs, both the encoder and decoder are RNNs. In particular, we adopt the same architecture as in sketch-rnn [16] for conditional sketch synthesis. That is, a bi-directional RNN [27] is used as the encoder to extract the global embedding of the input sketch. The final state output of the encoder is then projected to a mean and a variance vector, to define an IID Gaussian distribution. That distribution is then sampled to produce a random vector z as the representation of the input sketch. Thus z is not a deterministic output of the encoder given a sketch, but a random vector conditional on the input. The decoder is an LSTM model. Its initial state is conditional on z via a single fully connected (FC) layer. At each time step, it then predicts the offset for each stroke segment in order to reconstruct the input sketch. For further details on the encoder/decoder architecture, please refer to [16].

4.3 Formulation

The decoder splits into two branches after the LSTM hidden cell outputs: a generative branch to synthesise a sketch and a discriminative branch for grouping. Different learning objectives are used for the two branches: in the generative branch, two losses encourage the model to reconstruct the input sketch; in the discriminative branch, the sketch grouping annotation is used to train the decoder to produce an accurate stroke affinity matrix for grouping.

Group Affinity Matrix: The grouping annotation is represented as a sparse matrix denoting the group relationship between segments $\mathbf{G} \in \mathbb{R}^{N \times N}$. Denoting the i^{th} sketch segment as $S_i, i \in [1, N]$, we have:

$$G_{i,j} = \begin{cases} 1, & \text{if } S_i, S_j \text{ are from the same group} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Where each element of the matrix indicates whether the i^{th} and j^{th} segments belong to the same group or not. A straightforward design of the discriminative learning objective is to make the affinity matrix computed using the learned stroke feature $f_i = \phi(S_i)$ as similar as possible to $\mathbf{G}$, via an l_1 or l_2 loss. However, we found that in practice this works very poorly. This is because $\mathbf{G}$ conveys two types of grouping constraints: each element enforces a binary pairwise constraint for two segments, whilst the whole matrix also enforces global grouping constraint, e.g., if S_1 and S_2 are in the same group, and S_2 and S_5 are also in the same group, then global grouping consistency dictates that S_1 and S_5 must also belong to the same group. Balancing these two is critical because pairwise grouping predictions are typically noisy and can lead to global grouping inconsistency. However, using a single loss makes it impossible to achieve a balance. We thus propose to use two losses to implement the two constraints.

Local Grouping Loss: This loss requires that the pairwise relationship between two segments are kept when the pairwise affinity is measured using the learned stroke segment feature. The decoder LSTM learns a mapping function ϕ and map the i^{th} stroke segment S_i to a 128D feature vector f_i . To measure the affinity of any two segments in the input sketch, the absolute element-wise

feature difference is computed to obtain a symmetric absolute feature difference matrix, $\mathbf{D} \in \mathbb{R}^{N \times N \times 128}$ as:

$$\mathbf{D} = \{D_{i,j} \mid i, j \in [1, N]\} = \{|f_i - f_j| \mid i, j \in [1, N]\}. \quad (2)$$

Each vector $D_{i,j} \in \mathbb{R}^{128}$ is then subject to a binary classification loss (cross-entropy) to obtain the local affinity prediction $\hat{G}_{i,j}$, between the i^{th} and j^{th} segments. The local grouping loss, $\mathcal{L}_A$, is thus computed as:

$$\mathcal{L}_A = \sum_{i=1}^N \sum_{j=1}^N [-G_{i,j} \log(\hat{G}_{i,j}) - (1 - G_{i,j}) \log(1 - \hat{G}_{i,j})]. \quad (3)$$

Global Grouping Loss: Using only a local grouping loss may lead to global grouping inconsistency. However, formulating the global grouping consistency into a loss for a deep neural network is not straightforward. Our strategy is to first derive a global grouping representation for each segment using the local affinity prediction $\hat{G}_{i,j}$. We then use a triplet ranking loss to enforce that the segments belonging the same group have more similar grouping relationships to each other, than to a segment outside the group. Although the triplet ranking loss involves three segments only, since each segment is represented by its grouping relationship to all other segments, this loss is a global grouping loss. More concretely, we first construct the local affinity prediction matrix $\hat{\mathbf{G}}$ with $\hat{G}_{i,j}$ as elements. Each row vector of $\hat{\mathbf{G}}$, $\hat{G}_{i,:}$, is then used as a global grouping relationship vector to represent S_i . The final global grouping loss $\mathcal{L}_G$ is:

$$\mathcal{L}_G = \max(0, \Delta + d(\hat{G}_{i,:}, \hat{G}_{i^+,:}) - d(\hat{G}_{i,:}, \hat{G}_{i^-,:})), \quad (4)$$

where i represents an anchor segment, i^+ a positive segment in the same group and i^- a negative segment from a different group, Δ is a margin and $d(\cdot)$ denotes a distance function between two feature inputs. Here we take the squared Euclidean distance under the l_2 normalisation.

Generative Losses: For the generative branch, we use the same generative losses as in [16]. These include a reconstruction loss $\mathcal{L}_R$ and a KL loss $\mathcal{L}_{KL}$ measuring the difference between the latent random vector z and a IID Gaussian vector with zero-mean and unit variance.

Full Learning Objective: Our full loss $\mathcal{L}_F$ can be written as:

$$\mathcal{L}_F = \lambda_a \mathcal{L}_A + \lambda_g \mathcal{L}_G + \lambda_r (\mathcal{L}_R + \mathcal{L}_{KL}) \quad (5)$$

where the hyper-parameters λ_a , λ_g , and λ_r describe the relative importance of the different losses in the full training objective.

Model Testing: During testing stage, given a sketch, the trained model is used to compute an estimated segment affinity matrix, $\hat{G}$. This affinity matrix is then used to generate the final grouping. Since the number of groups varies for different sketches, the group number also needs to be estimated. To this end, we adopt a recent agglomerative clustering method [28] to obtain the final grouping. Note that the method does not introduce any additional free parameters.

4.4 Applications to Sketch Analysis

Sketch Synthesis from Edge Map: A simple sketch synthesis method can be developed based on the proposed universal grouper. The method is based on grouping edge maps extracted from photo images and removing the least important groups. Assume that the N segments of an edge map have been grouped in K groups, denoted as $P_k, k \in [1, K]$. An importance measure is defined as:

$$I(P_k) = I_L(P_k) \cdot I_N(P_k) + I_D(P_k) \quad (6)$$

where $I_L(P_k)$, $I_N(P_k)$ and $I_D(P_k)$ measure the importance from the perspectives of length, numbers and distribution of the segments in group P_k respectively. A less important group has smaller number of segments with shorter lengths but occupies a bigger region. We thus have:

$$I_L(P_k) = \frac{\sum_{i=1}^{N_{P_k}} L_{S_i}}{\sum_{i=1}^N L_{S_i}}, \quad I_N(P_k) = \frac{N_{P_k}}{N}, \quad I_D(P_k) = \frac{\max(w, h)N_{P_k}}{\sum_{i=1}^{N_{P_k}} d(M_{P_k}, M_{S_i})} \quad (7)$$

where N_{P_k} is the number of segments in P_k , L_{S_i} is the length of segment S_i , w and h are the width and height of the object, respectively, M_{P_k} denotes the average position of group P_k in the image plane, M_{S_i} represents the average position of segment S_i , and Euclidean distance $d(\cdot)$ is used. With the importance measure $I(P_k)$ computed for each group, we can then drop the least important groups defined as those with $I(P_k) < I_\delta$ where I_δ is a threshold.

Fine-Grained Sketch Based Image Retrieval: We further develop an unsupervised FG-SBIR method following [17]. Specifically, we apply our grouper to edge maps extracted from photos to synthesise human style sketches. Three threshold values of I_δ are used for each photo to accounts for the variable levels of abstraction among human sketchers. The photos and corresponding synthesised sketches are then used as data to train an off-the-shelf FG-SBIR model [19]. During testing, the grouping and group removal processes are applied to the human sketches, again with three different thresholds. The matching scores using the three abstracted sketches plus the original query sketch are then fused to produce the final retrieval results. Note that for this unsupervised FG-SBIR model to work well, our grouper must be truly universal: it needs to work well on both human sketches which it was trained on, and photo edge maps.

5 Experiments

5.1 Datasets and Settings

Dataset Splits and Preprocessing: Among the 25 categories in the new SPG dataset, we randomly select 20 as **seen categories**, and use the remaining 5 categories as **unseen categories** to test the generalisation of our universal grouper. In each seen category, we select 650 sketches for training, 50 for validation, and 100 for testing. For the unseen categories, no data are used for training

and we randomly select 100 sketches per category for testing to have the same per-category size as the seen categories. We normalise all the sketch strokes, and augment the sketch via stroke removal and distortion [23].

Implementation Details: Our deep grouper is implemented on Tensorflow on a single Titan X GPU. For model training, we set the importance weights λ_r , λ_a and λ_g for different losses (Eq. (4)) to 0.5, 0.6, and 1, respectively. The Adam optimiser [29] is applied with the parameters $\beta_1 = 0.5$, $\beta_2 = 0.9$, $\epsilon = 10^{-8}$. The initial learning rate is set to 0.0003 with exponential weight decay. The model is trained for 22,000 iterations with a batch size of 100.

Evaluation Metrics: Sketch perceptual grouping shares many common characteristics with the unsupervised image segmentation problem [5]. We thus adopt the same metrics including variation of information (VOI), probabilistic rand index (PRI), and segmentation covering (SC) as defined in [30]. More detailed definition of these metrics in the context of sketch grouping are: (i) **Variation of Information:** In this metric, the distance between two groups in terms of their average conditional entropy is calculated. (ii) **Probabilistic Rand Index:** Rand index compares the compatibility of assignments between pairs of stroke segments in each group. (iii) **Segmentation Covering:** the overlapping between the machine grouping and human grouping is measured. For SC and PRI, higher scores are better, while for VOI, a lower score indicates better results.

Competitors: Very few sketch perceptual grouping methods exist. The state-of-the-art model **Edge-PG** [12] uses two Gestalt principles, namely proximity (spatial closeness) and continuity (slope trend) to compute an affinity matrix and feeds the matrix to a graph cut algorithm to get the groups. The weightings of the two principles are learned from data using RankSVM. This method thus differs from ours in that hand-crafted features are used and only two principles are modelled. Beyond sketch grouping, many semantic image segmentation methods have been proposed lately based on fully convolutional networks (FCN). We choose one of the state-of-the-art models, **DeepLab** [7] as a baseline. It is trained to take images as input and output the semantic grouping, i.e., each pixel is assigned a class label. A conditional random field (CRF) is integrated to the network to enforce the proximity and similarity principles. Note that: (1) DeepLab is a supervised semantic segmentation method. It thus needs not only grouping annotation as our model does, but also group ID annotation, which is not used by our model and Edge-PG. This gives it an unfair advantage. (2) It performs grouping at the pixel level whilst both our model and Edge-PG do it at the stroke/segment level.

5.2 Results on Perceptual Grouping

Results on Seen Categories: In this experiment, the model are trained on the seen category training set and tested on the seen category testing set. From Table 1, we can see that: (i) Our model achieves the best performance across all 25 categories on each metric, except the VOI metric on face where our model is slightly inferior to DeepLab. The improvement on VOI is particularly striking indicating that the groups discovered by our model in each sketch are distinctive

Category	Ours			Edge-PG [12]			DeepLab [7]		
	VOI ↓	PRI ↑	SC ↑	VOI ↓	PRI ↑	SC ↑	VOI ↓	PRI ↑	SC ↑
Airplane	0.58	0.88	0.78	0.72	0.80	0.71	1.09	0.72	0.65
Alarm clock	0.46	0.93	0.83	0.59	0.84	0.73	0.86	0.80	0.70
Ambulance	0.67	0.86	0.77	1.35	0.67	0.60	1.19	0.71	0.63
Ant	0.86	0.83	0.69	1.32	0.68	0.62	1.38	0.69	0.60
Apple	0.25	0.92	0.91	0.54	0.88	0.79	0.82	0.83	0.72
Backpack	0.57	0.88	0.79	1.29	0.70	0.61	1.59	0.67	0.59
Basket	0.76	0.84	0.74	1.27	0.71	0.59	1.37	0.69	0.61
Butterfly	0.83	0.76	0.65	1.30	0.69	0.58	1.58	0.66	0.58
Cactus	0.51	0.90	0.83	0.86	0.82	0.71	0.90	0.79	0.68
Calculator	0.50	0.86	0.83	0.98	0.77	0.68	1.17	0.72	0.64
Camp fire	0.28	0.95	0.91	1.05	0.71	0.65	0.77	0.85	0.74
Candle	0.89	0.78	0.69	1.47	0.65	0.57	1.54	0.67	0.60
Coffee cup	0.38	0.91	0.86	0.85	0.83	0.68	0.98	0.79	0.66
Crab	0.69	0.81	0.74	1.29	0.69	0.56	1.58	0.67	0.60
Duck	0.86	0.83	0.69	0.95	0.74	0.68	1.63	0.65	0.57
Face	0.81	0.84	0.74	1.24	0.69	0.61	0.80	0.82	0.73
Ice-cream	0.41	0.94	0.85	0.79	0.82	0.71	1.40	0.68	0.62
Pig	0.63	0.84	0.78	1.55	0.63	0.50	0.98	0.77	0.67
Pineapple	0.50	0.93	0.82	0.63	0.83	0.72	1.05	0.74	0.65
Suitcase	0.54	0.89	0.83	0.58	0.82	0.75	1.10	0.73	0.64
Average	0.59	0.87	0.79	1.03	0.75	0.65	1.20	0.73	0.65

Table 1: Comparative grouping results on seen categories.

Category	Ours			Edge-PG [12]		
	VOI ↓	PRI ↑	SC ↑	VOI ↓	PRI ↑	SC ↑
Angel	0.70	0.87	0.73	1.19	0.69	0.60
Bulldozer	0.81	0.85	0.73	1.37	0.65	0.58
Drill	0.67	0.78	0.77	1.45	0.61	0.53
Flower	0.39	0.90	0.84	0.79	0.75	0.64
House	0.46	0.91	0.83	0.85	0.77	0.69
Average	0.64	0.86	0.77	1.13	0.69	0.61

Table 2: Perceptual grouping results on unseen categories.

to each other. In contrast, the two compared models tend to split a semantic part into multiple groups resulting similar groups (see Fig. 4). (ii) Edge-PG is much worse than our method because it is based on hand-crafted features for only two principles, while our model implicitly learns the features and combination strategy based on end-to-end learning from human group annotation. (iii) Although DeepLab also employs a deep neural network and uses additional annotations, its result is no better than Edge-PG. This suggests that for sketch perceptual grouping, it is important to treat sketches as a set of strokes rather than pixels, as strokes already grouping pixels. These constraints are ignored by the DeepLab types of models designed for photo image segmentation.

Some examples of the grouping results are shown in Fig. 4. As expected, ignoring the stroke level grouping constraint on pixels, each stroke is often split into multiple groups by DeepLab [7]. Edge-PG [12] does not suffer from that problem. However, it suffers from the limitations on modelling only two principles. For example, to group the clock contour (second column) into one group, the closure principle needs to be used. It is also unable to model even the two principles effectively due to the limited expressive power of hand-crafted features: in the airplane example (first column), the two wings should be grouped

Fig. 4: Qualitative grouping results on seen categories.

Fig. 5: Qualitative grouping results on unseen categories.

together using the continuity principle, but broken into two by Edge-PG. In contrast, our model produces more consistent groupings using multiple principles dynamically. For instance, both DeepLab and our model successfully deploy the similarity principle to group the two legs of both the alarm clock (second column) and duck (third column) together. But DeepLab does so by explicitly encoding the principle in its CRF layer, while our model does it implicitly. In the cactus example (last column), to produce the correct grouping of those spikes, both continuity, similarity and less prevalent principles such as common fate need to be combined. Only our model is able to do that because it has implicitly learned to model all the principles used by humans to annotate the groupings.

Results on Unseen Categories: In this experiment, models learned using seen categories are tested directly on unseen categories without any fine-tuning. This experiment is thus designed to evaluate whether the grouper is indeed universal, i.e., can be applied to any new object category. Note that as a supervised segmentation method, DeepLab cannot be applied here because each category has a unique set of group IDs. The results of our model and Edge-GP are shown in Table 2. It can be seen that our model outperforms Edge-GP by a big margin. Importantly, comparing Table 2 with Table 1, our model’s performance on PRI and SC hardly changed. In contrast, the Edge-PG’s performance on the unseen categories is clearly worse than that on the seen categories. This suggests that our grouper is more generalisable and universal. Some qualitative results are shown in Fig. 5. It again shows that the lack of powerful feature learning and limitation on only two principles contribute to the weaker results of Edge-GP.

Method	VOI ↓	PRI ↑	SC ↑
Edge-PG [12]	1.69	0.62	0.53
Ours	0.96	0.78	0.71

Table 3: Compare Edge-PG [12] with ours on Edge-PG [12]’s datasets.

Method	Seen Categories			Unseen Categories		
	VOI ↓	PRI ↑	SC ↑	VOI ↓	PRI ↑	SC ↑
Ours - A - G	1.45	0.65	0.59	1.53	0.64	0.56
Ours - R - G	1.12	0.71	0.64	1.36	0.68	0.59
Ours - R - A	1.27	0.69	0.63	1.48	0.64	0.57
Ours - G	0.63	0.86	0.78	0.71	0.84	0.73
Ours - A	0.75	0.80	0.72	0.95	0.78	0.67
Ours - R	0.68	0.83	0.76	0.86	0.78	0.69
Ours + l_2	2.68	0.58	0.49	2.63	0.59	0.49
Ours full model	0.59	0.87	0.79	0.64	0.86	0.77

Table 4: Performance of different variants of our model on seen and unseen categories.

Results on Unseen Dataset: To further demonstrate the generalisation ability of our universal grouper, we test the trained model on a different dataset. Specifically, we choose 10 categories from the dataset in [32] including 5 categories overlapping with our dataset and 5 new categories. Note that the sketches in this datasets are from the database proposed in [33], which are drawn without the 20 second constraint, thus exhibiting much more details with better quality in general. This dataset thus represents a different domain. Table 3 shows that our model again demonstrates better generalisation ability than Edge-PG.

Ablation Study: Our model is trained with a combination of generative and discriminative losses (Sec. 4.3). These include the local grouping loss $\mathcal{L}_A$, global grouping loss $\mathcal{L}_G$, reconstruction loss $\mathcal{L}_R$ and KL loss $\mathcal{L}_{KL}$. Among them, all but the KL loss can be removed, leading to six variants of our model, e.g., Ours - A - G is obtained by removing $\mathcal{L}_A$ and $\mathcal{L}_G$. In addition, we implement Ours + l_2 which is having an l_2 loss on the predicted affinity matrix $\hat{G}$ w.r.t. the ground truth matrix G to examine the importance of having separate local and global grouping losses. The results are shown in Table 4. Clearly all three losses contribute to the performance of our model. The poorest result was obtained when an l_2 loss is added on the predicted affinity matrix, suggesting that balancing the local and global grouping losses is critical for learning a good grouper. We further show that the improvement of our full model over Ours - R on the unseen categories (0.64 vs 0.86) is bigger on the seen categories (0.59 vs. 0.68). This indicates that the generative loss helps the model generalise to unseen categories.

5.3 Applications on Sketch Synthesis and FG-SBIR

One application of our grouper is to use it as an abstraction model so that edge maps extracted from photos can be grouped and abstracted to synthesise human-like sketches. Fig. 6 shows some examples of edge map grouping results and the

Method	Shoe-V2		Chair-V2	
	Top1	Top10	Top1	Top10
Scribbler [31]	8.86%	32.28%	31.27%	78.02%
LDSA [17]	21.17%	55.86%	41.80%	84.21%
Ours	26.88%	61.86%	45.57%	88.61%
Upper Bound	34.38%	79.43%	48.92%	90.71%

Table 5: FG-SBIR performance on Shoe-V2 and Chair-V2 dataset

Fig. 6: Applying our grouper to synthesise abstract sketches from photo edge map. **(a)** columns show the photos; **(b)** columns give the edge maps extracted from the photos and the grouping results; **(c)** columns provide synthesised abstract sketches.

synthesised sketches. It can be seen that our grouper is generalisable to photo edges and our abstraction method produces visually appealing sketches. The synthesised sketches are then used to train a state-of-the-art FG-SBIR model [19] without using any real human sketches. We use the largest FG-SBIR datasets QMUL Shoe-V2 and Chair-V2 [34]. We first compare with the same FG-SBIR model trained using synthesised sketches from the deep conditional GAN network in [31] (denoted as Scribbler). As can be seen in Table 5, our model performs much better. This suggests that our edge abstraction model, albeit simple, synthesises more realistic sketches from edge maps. We further compare with a recently proposed unsupervised FG-SBIR model LDSA [17] which is also based on abstracting photo edge maps to synthesise sketches. Table 5 shows that our model outperforms LDSA model by 5.71% and 3.63% on top 1 accuracy on Shoe-V2 and Chair-V2, respectively. The results are not far off the upper-bound which is obtained using the same FG-SBIR model trained with the real sketch-photo pairs in Shoe-V2 and Chair-V2. This shows that our method enables FG-SBIR to be used without the expensive collection of sketch-photo pairs.

6 Conclusion

We have proposed an end-to-end sketch perceptual grouping model. This is made possible by collecting a new large-scale sketch grouping dataset SPG. Our grouper is trained with generative losses to make it generalisable to new object categories and datasets/domains. Two grouping losses were also formulated to balance the local and global grouping constraints. Extensive experiments showed that our model significantly outperforms existing groupers. We also demonstrated our grouper’s application to sketch synthesis and FG-SBIR.

References

1. J. Wagemans, J. H. Elder, M. Kubovy, S. E. Palmer, M. A. Peterson, M. Singh, and R. von der Heydt, “A century of gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization.” *Psychological bulletin*, 2012.
2. J. Wagemans, J. Feldman, S. Gepshtein, R. Kimchi, J. R. Pomerantz, P. A. van der Helm, and C. van Leeuwen, “A century of gestalt psychology in visual perception: II. Conceptual and theoretical foundations.” *Psychological bulletin*, 2012.
3. M. Wertheimer, *Laws of organization in perceptual forms*. London, England: Kegan Paul, Trench, Trubner & Company, 1938.
4. X. Ren and J. Malik, “Learning a classification model for segmentation,” in *ICCV*, 2003.
5. X. Xia and B. Kulis, “W-net: A deep model for fully unsupervised image segmentation,” *ArXiv e-prints*, 2017.
6. C. Wang, B. Yang, and Y. Liao, “Unsupervised image segmentation using convolutional autoencoder with total variation regularization as preprocessing,” in *ICASSP*, 2017.
7. L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *arXiv:1606.00915*, 2016.
8. Z. Sun, C. Wang, L. Zhang, and L. Zhang, “Free hand-drawn sketch segmentation,” in *ECCV*, 2012.
9. Z. Huang, H. Fu, and R. W. Lau, “Data-driven segmentation and labeling of free-hand sketches,” *TOG*, 2014.
10. R. G. Schneider and T. Tuytelaars, “Example-based sketch segmentation and labeling using crfs,” *TOG*, 2016.
11. Y. Qi, J. Guo, Y. Li, H. Zhang, T. Xiang, and Y.-Z. Song, “Sketching by perceptual grouping,” in *ICIP*, 2013.
12. Y. Qi, Y. Z. Song, T. Xiang, H. Zhang, T. Hospedales, Y. Li, and J. Guo, “Making better use of edges via perceptual grouping,” in *CVPR*, 2015.
13. G. Hinton and R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, 2006.
14. A. Amir and M. Lindenbaum, “A generic grouping algorithm and its quantitative analysis,” *TPAMI*, 1998.
15. Z. Lun, C. Zou, H. Huang, E. Kalogerakis, P. Tan, M.-P. Cani, and H. Zhang, “Learning to group discrete graphical patterns,” *TOG*, 2017.
16. D. Ha and D. Eck, “A neural representation of sketch drawings,” *arXiv preprint arXiv:1704.03477*, 2017.
17. U. R. Muhammad, Y.-Z. Song, T. Xiang, and T. Hospedales, “Learning deep sketch abstraction,” in *CVPR*, 2018.
18. Y. Li, T. M. Hospedales, Y.-Z. Song, and S. Gong, “Fine-grained sketch-based image retrieval by matching deformable part models,” in *BMVC*, 2014.
19. Q. Yu, F. Liu, Y.-Z. Song, T. Xiang, T. M. Hospedales, and C. C. Loy, “Sketch me that shoe,” in *CVPR*, 2016.
20. J. Song, Y.-Z. Song, T. Xiang, T. Hospedales, and X. Ruan, “Deep multi-task attribute-driven ranking for fine-grained sketch-based image retrieval,” in *BMVC*, 2016.
21. K. Pang, Y.-Z. Song, T. Xiang, and T. M. Hospedales, “Cross-domain generative learning for fine-grained sketch-based image retrieval,” in *BMVC*, 2017.

22. J. Song, Y. Qian, Y.-Z. Song, T. Xiang, and T. Hospedales, “Deep spatial-semantic attention for fine-grained sketch-based image retrieval,” in *ICCV*, 2017.
23. Q. Yu, Y. Yang, F. Liu, Y.-Z. Song, T. Xiang, and T. M. Hospedales, “Sketch-a-net: A deep neural network that beats humans,” *IJCV*, 2017.
24. S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio, “Generating sentences from a continuous space,” *Computer Science*, 2015.
25. D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *ArXiv e-prints*, 2013.
26. A. Graves, “Generating sequences with recurrent neural networks,” *arXiv preprint arXiv:1308.0850*, 2013.
27. M. Schuster and K. K. Paliwal, “Bidirectional recurrent neural networks,” *TSP*, 1997.
28. J. Yang, D. Parikh, and D. Batra, “Joint unsupervised learning of deep representations and image clusters,” in *CVPR*, 2016.
29. D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
30. P. Arbellez, M. Maire, C. Fowlkes, and J. Malik, “Contour detection and hierarchical image segmentation,” *TPAMI*, 2011.
31. P. Sangkloy, J. Lu, C. Fang, F. Yu, and J. Hays, “Scribbler: Controlling deep image synthesis with sketch and color,” *CVPR*, 2017.
32. Y. Qi, J. Guo, Y. Z. Song, T. Xiang, H. Zhang, and Z. H. Tan, “Im2sketch: Sketch generation by unconflicted perceptual grouping,” *Neurocomputing*, 2015.
33. M. Eitz, J. Hays, and M. Alexa, “How do humans sketch objects?” *TOG*, 2012.
34. Q. Yu, Y.-Z. Song, T. Xiang, and T. M. Hospedales, “SketchX – Shoe/Chair fine-grained SBIR dataset,” <http://sketchx.eecs.qmul.ac.uk>, 2017.