Supplemental Materials of "Detecting Viewer-Perceived Intended Vector Sketch Connectivity"

JERRY YIN* and CHENXI LIU*, University of British Columbia, Canada
REBECCA LIN, University of British Columbia, Canada
NICHOLAS VINING, University of British Columbia, Canada and NVIDIA, Canada
HELGE RHODIN, University of British Columbia, Canada
ALLA SHEFFER, University of British Columbia, Canada

ACM Reference Format:
Jerry Yin, Chenxi Liu, Rebecca Lin, Nicholas Vining, Helge Rhodin, and Alla Sheffer. 2022. Supplemental Materials of "Detecting Viewer-Perceived Intended Vector Sketch Connectivity". 1, 1 (May 2022), 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 ADDITIONAL RESULTS

Five additional representative results and comparisons are shown in Fig. 1. All comparison methods produce noticeable artifacts on these inputs (highlighted in Fig. 1), while our method generates artist-intended results. These results demonstrate that our method outperforms existing methods and can be successfully applied both to realistic vector line drawings in the wild, and to sketches produced by upstream sketch processing methods.

2 IMPLEMENTATION DETAILS

2.1 List of Features

A list of the classifier features we use and their Gini importance values are listed in Table 1.

2.2 Preprocessing

Hook Removal. When finishing a stroke, if the artist rapidly switches direction before the pen raise is registered by the drawing device, an unintentional hook can appear. These can interfere with tangent and distance computations. Handling hooks robustly remains an open problem and is not a contribution of our paper. Let vertices that remain after applying the Ramer-Douglas-Peucker algorithm be corners. We use a simple heuristic: the section from the endpoint to the first corner is a hook if the distance from the endpoint p_1 to the line segment between its two nearest corners is shorter than $W_1 f$, where f is a parameter, and the hook is shorter than $\min(1.5W_1 f, \frac{1}{2}L_1)$. We use $f = 3$ for Company et al. [2019]; Gryaditskaya et al. [2019]; Ha and Eck [2018]; Qi et al. [2021] and $f = 1$ for all other sources. We preserve near-connections by only de-hooking an endpoint if it would not increase the envelope distance to its nearest stroke by more than a factor of 2. Across all processed inputs, we further manually removed 10 hooks not found by the criteria above.

Consolidation. Most of the inputs we process were not consolidated previously, and thus may contain overdrawing. To align stroke length and local context feature computations to human perception, we weakly consolidate our inputs using a simple heuristic that accounts for the presence of varying stroke widths and light oversketching in our data. We locate all pairs of partially side-by-side strokes and densely sample their side-by-side sections using orthogonal cross-sections. If the sections have roughly similar lengths (ratio $\in [1/1.2, 1.2]$), have similar tangent directions at corresponding cross-section points ($< 20°$) that are themselves close enough (all $< \frac{3}{2}\max(W_1, W_2)$), the endpoints are close enough ($< \max(W_1, W_2)$), and if either the endpoints both overlap with the other stroke or the substrokes are both longer than their pen widths, we replace the two strokes with a single stroke fitted to both. We repeat this on all pairs until no more strokes can be merged. Finally, we chain strokes by merging them into a single stroke when their endpoints overlap and the endpoint tangents align within 20°.

Dangling Endpoints. A subset of stroke endpoints in our drawings already overlap other strokes, and can be connected without consulting our classifier. These endpoints are non-dangling. In the case of an overshot connection, where the two stroke centerlines intersect, we define the intersection diameter d as the largest stroke width at the intersection. An endpoint is non-dangling if the length of the overshot portion is shorter than 15% of the larger of the stroke length and d, and the Euclidean distance from the intersection to the edge of the endpoint cap is less than $1.5d$.

Self-Connections. We define a connection between points $\mathrm{p}_1, \mathrm{p}_2$ on the same stroke to be valid if they form a loop—if $\max(3\|\mathrm{p}_1 - \mathrm{p}_2\|, \pi W_1)$ is less than the distance between $\mathrm{p}_1, \mathrm{p}_2$ along the stroke. We then define the projection of an endpoint onto its own stroke as the closest valid point, if it exists. From there, feature computations work as described in the main paper.

3 STUDY DESIGN

Junction Annotation Study. To validate our final connection decisions, we collected additional manual annotations of 91 potential end-to-end and T-junctions across 10 drawings (see our supplementary materials for the full question set). In this study (Fig. 2), for a

*Both authors contributed equally to this research.

Authors' addresses: Jerry Yin, jerryyin@student.ubc.ca; Chenxi Liu, chenxil@cs.ubc.ca, University of British Columbia, Vancouver, Canada; Rebecca Lin, ryelin@student.ubc.ca, University of British Columbia, Vancouver, Canada; Nicholas Vining, nvining@cs.ubc.ca, University of British Columbia, Vancouver, Canada, NVIDIA, Toronto, Canada; Helge Rhodin, rhodin@cs.ubc.ca, University of British Columbia, Vancouver, Canada; Alla Sheffer, sheffa@cs.ubc.ca, University of British Columbia, Vancouver, Canada.

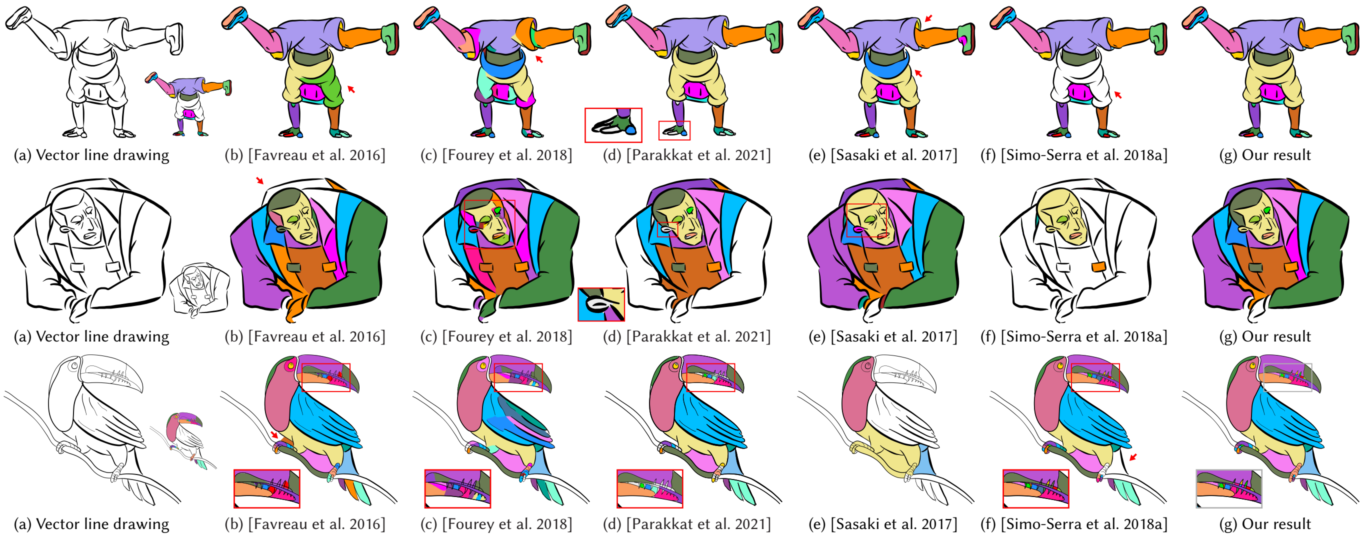

Fig. 1. Additional results and comparisons. Insets on the left show loops formed by raw intersecting strokes. These inputs are from different data sources (from top to bottom: two Blender Art Gallery drawings [2021] by different artists, and a consolidated StrokeStrip sketch [2021]), showing our method's ability to support different styles and stroke widths. Light gray spots in the outputs of [Parakkat et al. 2021] correspond to pixels unassigned by their method. Input drawings (shown with modifications in this work) from top to bottom ©Lien-ze Tsao under CC BY 4.0; ©The "Hero" artist Team under CC BY 4.0; ©Elinor Palomares.

Table 1. Gini importances of junction classifier features.

Endpoint-endpoint features	Gini importance
Envelope distance $d^E/(0.5(W_1+W_2))$	0.246
Envelope distance $d^E/\max(L_1, L_2)$	0.171
Envelope distance $d^E/\min(L_1, L_2)$	0.160
Junction type	0.001
$\max(\theta_1, \theta_2)$	0.015
$\min(\theta_1, \theta_2)$	0.011
Larger step-away ratio $\max(d_1^S/d^C, d_2^S/d^C)$	0.031
Smaller step-away ratio $\min(d_1^S/d^C, d_2^S/d^C)$	0.027
Larger projection ratio	0.001
Smaller projection ratio	0.009
Larger relative location	0.009
Smaller relative location	0.0004
Larger distance to nearest other	0.190
Smaller distance to nearest other	0.128
T-junction feature	**Gini importance**
Envelope distance $d^E/(0.5(W_1+W_2))$	0.250
Envelope distance d^E/L_1	0.174
Envelope distance d^E/L_2	0.270
θ_1	0.067
Step-away ratio d_1^S/d^C	0.077
Relative location	0.012
Distance to nearest other	0.136
Endpoint density b	0.015

given question, participants were shown a full line drawing with colors indicating potential endpoint-endpoint connections (endpoints coloured pink and green with gradients) and T-junctions (the endpoint coloured orange with a gradient and the other stroke flatly painted in blue), as well as a zoomed-in view around the potential junction in question. Participants were shown "a series of magnified views where one or two strokes are highlighted at the region of interest." Participants were then asked to "identify whether the strokes were intended by the artist to form a junction at the highlighted region of interest" and to answer the question "Did the artist intend for the two highlighted endpoints (pink and green) to form a junction?" or "Did the artist intend for the highlighted endpoint (orange) and the highlighted stroke (blue) to form a T-junction, with the highlighted stroke as the top of the 'T'?", depending on the shown junction type. We recruited 16 non-expert participants (nine males and seven females, split between two sessions with 43 and 48 questions respectively), resulting in each potential junction labelled by eight participants. Two examples of the study question are shown in Fig. 2; the full instructions and annotation questions are included in our supplementary materials.

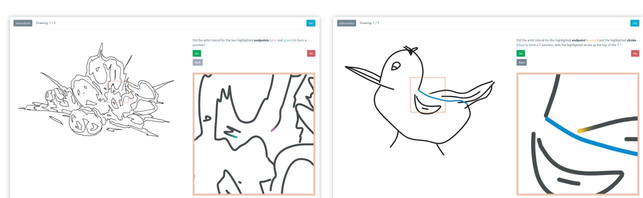

(a) End-end junction question example

(b) T-junction question example

Fig. 2. Junction annotation study example questions and interface. The full line drawing is shown on the left, the zoomed-in view of the potential junction in question and the corresponding question is shown on the right. In both views, we use color gradients to indicate endpoints (pink and green in (a) and orange in (b)) and a solid blue to indicate the non-endpoint stroke of a T-junction (b). Images left and right ©Nahu under CC BY 4.0; [Ge et al. 2020].

Comparative Study. We conducted a comparative study to evaluate human perceptual preference between our method and existing gap closure methods. Participants were shown an input line drawing on top (A), and colorizations of this drawing obtained using our method and an alternative method, randomly assigned to (B) and (C) below. Participants were asked to “envision which strokes in [the input] drawings are intended by the artist to form closed loops,” to “Identify the differences between the two [shown] colorings (ignore small color bleedings),” and then to answer “Which of the images on the bottom. (B) or (C), better corresponds to the partition you envisioned?” by selecting from “(B),” “(C),” “Both,” and “Neither.” We recruited 30 participants (19 male, 11 female), resulting in six responses per question for 27 inputs from 11 data sources with diverse authors and styles. We used the same 27 inputs for all five comparison methods and ensured that no drawing was shown more than once in the same questionnaire. The questionnaires and response counts per question are included in our supplementary materials.

We generated results for comparison methods by providing participants with inputs rasterized at 600px for Favreau et al. [2016], Fourey et al. [2018], Sasaki et al. [2017], and Simo-Serra et al. [2018] and Parakkat et al. [2021]. The full set of results is included in our supplementary materials. We generated the colorizations of resulting closed loops by recoloring the output colorizations from Fourey et al. [2018] and Parakkat et al. [2021] (setting their “unassigned” pixels to light gray); by identifying and coloring closed loops in the vector outputs of Favreau et al. [2016]; and by first binarizing the output grayscale raster images with a threshold of 0.5 (for pixel intensities in [0, 1]) and then flood-filling with a size of 1 px for Sasaki et al. [2017] and Simo-Serra et al. [2018]. We chose colors for each result pair in a question such that the corresponding closed loops between the two results were assigned the same color, and different closed loops within a partition were given different colors.

REFERENCES

Blender. 2021. Blender Cloud. https://cloud.blender.org/p/gallery/5b642e25bf419c1042056fc6

Pedro Company, Raquel Plumed, Peter A. C. Varley, and Jorge D. Camba. 2019. Algorithmic Perception of Vertices in Sketched Drawings of Polyhedral Shapes. *ACM Trans. Appl. Percept.* 16, 3 (Aug. 2019), 18:1–18:19.

Jean-Dominique Favreau, Florent Lafarge, and Adrien Bousseau. 2016. Fidelity vs. Simplicity: A Global Approach to Line Drawing Vectorization. *ACM Trans. Graph.* 35, 4 (July 2016), 120:1–120:10.

Sébastien Fourey, David Tschumperlé, and David Revoy. 2018. *A Fast and Efficient Semi-Guided Algorithm for Flat Coloring Line-Arts.* The Eurographics Association.

Songwei Ge, Vedanuj Goswami, Larry Zitnick, and Devi Parikh. 2020. Creative Sketch Generation. In *International Conference on Learning Representations.*

Yulia Gryaditskaya, Mark Sypesteyn, Jan Willem Hoftijzer, Sylvia Pont, Frédo Durand, and Adrien Bousseau. 2019. OpenSketch: A Richly-Annotated Dataset of Product Design Sketches. *ACM Trans. Graph.* 38, 6 (Nov. 2019), 232:1–232:16.

David Ha and Douglas Eck. 2018. A Neural Representation of Sketch Drawings. In *International Conference on Learning Representations.* https://openreview.net/forum?id=Hy6GHpkCW

Amal Dev Parakkat, Marie-Paule Cani, and Karan Singh. 2021. Color by Numbers: Interactive Structuring and Vectorization of Sketch Imagery. In *CHI ’21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.*

Anran Qi, Yulia Gryaditskaya, Jifei Song, Yongxin Yang, Yonggang Qi, Timothy M. Hospedales, Tao Xiang, and Yi-Zhe Song. 2021. Toward Fine-Grained Sketch-Based 3D Shape Retrieval. *IEEE Trans. Image Process.* 30 (2021), 8595–8606.

Kazuma Sasaki, Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. 2017. Joint Gap Detection and Inpainting of Line Drawings. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR).* 5768–5776.

Edgar Simo-Serra, Satoshi Iizuka, and Hiroshi Ishikawa. 2018. Mastering Sketching: Adversarial Augmentation for Structured Prediction. *ACM Trans. Graph.* 37, 1 (Jan. 2018), 11:1–11:13.

Dave Pagurek Van Mossel, Chenxi Liu, Nicholas Vining, Mikhail Bessmeltsev, and Alla Sheffer. 2021. StrokeStrip: Joint Parameterization and Fitting of Stroke Clusters. *ACM Trans. Graph.* 40, 4 (July 2021), 50:1–50:18.