Tracing Versus Freehand for Evaluating Computer-Generated Drawings

ZEYU WANG, Yale University
SHERRY QIU, Yale University
NICOLE FENG, Carnegie Mellon University
HOLLY RUSHMEIER, Yale University
LEONARD MCMILLAN, University of North Carolina at Chapel Hill and Mental Canvas, Inc.
JULIE DORSEY, Yale University and Mental Canvas, Inc.

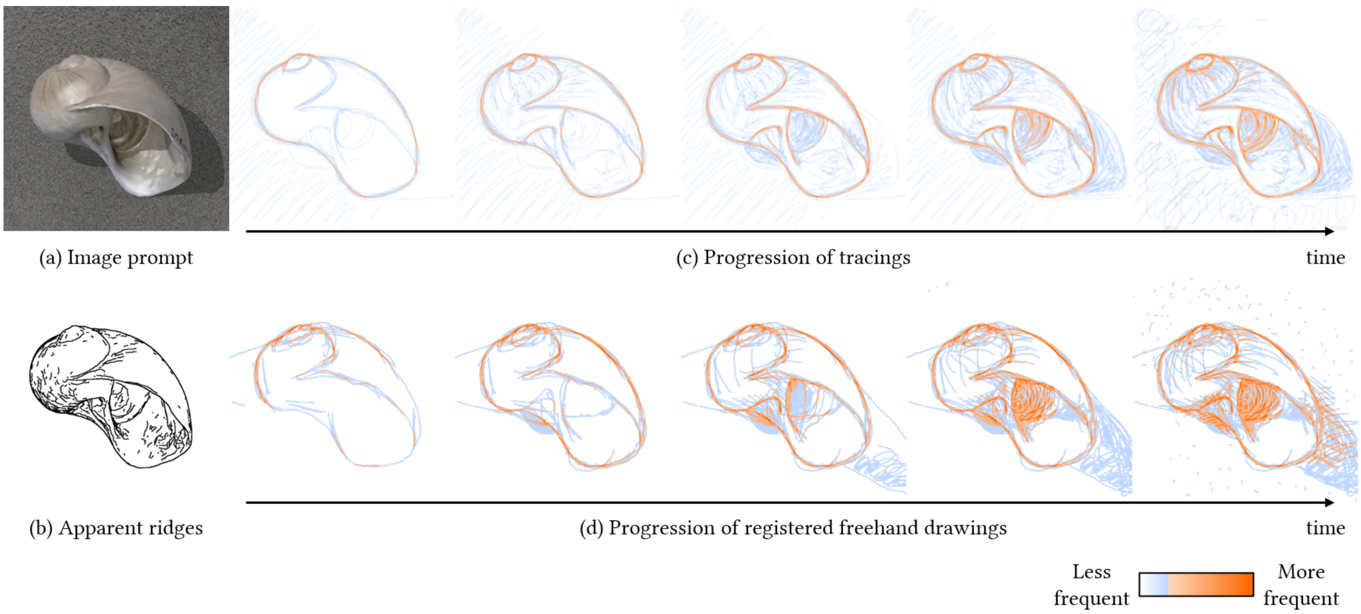

Fig. 1. Progression of drawings. (a) One of 100 image prompts in our dataset. (b) Apparent ridges rendered from the geometry. (c) Progression of density map of tracings binned into temporal quintiles. (d) Progression of density map of registered freehand drawings binned into temporal quintiles. All drawings of the prompt are superimposed. Note the spatio-temporal similarities between tracing and freehand drawing. Contours are drawn early and tone details later. Fossil model by Digital Atlas of Ancient Life on Sketchfab (CC0 1.0).

Non-photorealistic rendering (NPR) and image processing algorithms are widely assumed as a proxy for drawing. However, this assumption is not well assessed due to the difficulty in collecting and registering freehand drawings. Alternatively, tracings are easier to collect and register, but there is no quantitative evaluation of tracing as a proxy for freehand drawing. In this paper, we compare tracing, freehand drawing, and computer-generated drawing approximation (CGDA) to understand their similarities and differences. We collected a dataset of 1,498 tracings and freehand drawings by 110 participants for 100 image prompts. Our drawings are registered to the prompts and include vector-based timestamped strokes collected via stylus input. Comparing tracing and freehand drawing, we found a high degree of similarity in stroke placement and types of strokes used over time. We show that tracing can serve as a viable proxy for freehand drawing because of similar correlations between spatio-temporal stroke features and labeled stroke types. Comparing hand-drawn content and current CGDA output, we found that 60% of drawn pixels corresponded to computer-generated pixels on average. The overlap tended to be commonly drawn content, but people's artistic choices and temporal tendencies remained largely uncaptured. We present an initial analysis to inform new CGDA algorithms and drawing applications, and provide the dataset for use by the community.

CCS Concepts: • **Applied computing** → *Arts and humanities*; • **Human-centered computing** → *User studies*; • **Computing methodologies** → *Non-photorealistic rendering*; *Perception.*

Authors' addresses: Zeyu Wang, Yale University, 51 Prospect St, New Haven, CT, 06511, zeyu.wang@yale.edu; Sherry Qiu, Yale University, 51 Prospect St, New Haven, CT, 06511, sherry.qiu@yale.edu; Nicole Feng, Carnegie Mellon University, 500 Forbes Ave, Pittsburgh, PA, 15213, nfeng@andrew.cmu.edu; Holly Rushmeier, Yale University, 51 Prospect St, New Haven, CT, 06511, holly.rushmeier@yale.edu; Leonard McMillan, University of North Carolina at Chapel Hill and Mental Canvas, Inc. 201 S Columbia St, Chapel Hill, NC, 27599, mcmillan@cs.unc.edu; Julie Dorsey, Yale University and Mental Canvas, Inc. 51 Prospect St, New Haven, CT, 06511, julie.dorsey@yale.edu.

0730-0301/2021/8-ART52 $15.00
https://doi.org/10.1145/3450626.3459819

Additional Key Words and Phrases: sketch dataset, drawing process, stroke analysis

ACM Reference Format:
Zeyu Wang, Sherry Qiu, Nicole Feng, Holly Rushmeier, Leonard McMillan, and Julie Dorsey. 2021. Tracing Versus Freehand for Evaluating Computer-Generated Drawings. *ACM Trans. Graph.* 40, 4, Article 52 (August 2021), 12 pages. https://doi.org/10.1145/3450626.3459819

1 INTRODUCTION

The computer graphics and image processing communities have long been interested in developing methods that resemble freehand drawing. These computer-generated drawing approximations (CGDA) are often used as a proxy for hand drawings. For example, several recent deep learning-based applications [Isola et al. 2017; Li et al. 2018; Su et al. 2018] have relied exclusively on CGDA for training data. However, the extent to which current CGDA results actually resemble real drawings remains unclear. In particular, most CGDA methods only approximate the end result of drawing, neglecting the drawing process which includes stroke and temporal information. In this paper, we compare current CGDA methods with both freehand drawing and tracing, in order to better model the process and content of drawing and inform new algorithms.

Tracing is a form of hand drawing commonly used in art and design that has not been compared with freehand drawing or CGDA output. Tracings are easier to obtain than freehand drawings and are registered to image prompts by definition. Like freehand drawing, tracing can also provide temporal information and stroke demarcation. Comparisons of tracing to freehand drawing have been long discussed in the art community [Jamieson 2019], but there is a lack of quantitative analysis. Therefore, we consider tracing as a proxy for freehand drawing and also compare it to CGDA, thus enabling new paradigms for data collection and algorithm evaluation.

This study compares tracing, freehand drawing, and CGDA output based on a collection of tracings and freehand drawings. Our data have vector and temporal information and are registered to image prompts. We observed that in both freehand drawing and tracing, the types of strokes drawn over time are similar; common contours are drawn early and tone details occur later in the drawing process (Fig. 1). We found that in both forms of drawing, similar correlations exist between stroke features and time as well as between stroke features and drawing intention. This suggests that tracing can serve as a viable proxy for freehand drawing and justifies this efficient method of drawing data collection. In addition, we evaluated current CGDA methods by comparing their output to both freehand drawing and tracing and found a 60% overlap. This reveals aspects of CGDA methods that require more attention, such as the drawing process and different types of strokes. We expect our dataset and analysis to be useful for many applications, including data-driven CGDA, stroke type classification, drawing process simulation, sketch-based modeling, and sketch beautification.

This paper makes the following contributions:

- A dataset of 1,498 tracings and freehand drawings of 100 image prompts by 110 participants, with temporal, vector, and pressure information that are registered to the prompts.
- A comparison between tracing and freehand drawing through a spatio-temporal analysis, in which we demonstrate their similarities in stroke placement and drawing progression.
- A comparison of CGDA results with hand-drawn data, in which we show that current CGDA methods capture about 60% of pixels drawn and suggest areas for improvement.

2 RELATED WORK

2.1 A Taxonomy of Hand-Drawn Datasets

There has been enormous work on datasets and algorithms for various drawing applications. As shown in Table 1, large-scale datasets like TU-Berlin [Eitz et al. 2012], Sketchy [Sangkloy et al. 2016], and QuickDraw [Ha and Eck 2018] are intended for sketch recognition and image retrieval. The use of text prompts in these datasets leaves room for participants' interpretations and encourages symbolic sketches. While valuable for machine learning applications, many of these sketches do not reflect the complexity of observational drawing. In addition to freehand sketches, tracings, which are directly registered to the prompt, have been collected to study image segmentation and contour extraction. Tracing datasets like BSDS500 [Arbeláez et al. 2011] and PhotoSketch [Li et al. 2019] illustrate people's perception of edges and contours, which is essential to creating a good drawing [Suwa and Tversky 1997].

More related to our motivation are the drawing datasets used to understand how artists and designers draw, including Princeton [Cole et al. 2008], Portrait [Berger et al. 2013], OpenSketch [Gryaditskaya et al. 2019], and GMU [Yan et al. 2020]. The Portrait and OpenSketch datasets include temporal information to analyze the process of portrait drawing and technical drafting. The OpenSketch and GMU datasets include contours and additional types of strokes that reveal people's drawing intention. Because it is hard to collect high-quality data, these datasets are relatively small in scope (portrait photos or diffuse renderings) and size (one or two dozen prompts and usually at most ten artists per prompt). Therefore, it remains necessary to create a dataset of timestamped representational drawings with a wider variety of prompts, artists, and stroke types in order to further understand how people trace and draw.

2.2 Computer-Generated Drawing Approximation

The computer graphics community has been interested in creating NPR techniques to render various artistic styles. Different algorithms for generating lines have been proposed in order to imitate human-drawn strokes. Suggestive contours [DeCarlo et al. 2003] extend true contours of shapes; ridges and valleys [Ohtake et al. 2004] and apparent ridges [Judd et al. 2007] depict salient surface features; neural contours incorporate multiple line extractors [Liu et al. 2020]; and hatching strokes [Kalogerakis et al. 2012; Praun et al. 2001] convey material, tone, and form. These NPR techniques take a 3D model as input, examine differential properties such as surface curvatures, and produce lines in the form of static bitmaps. In the image domain, edge detection can generate lines at the boundary of high-contrast regions. Various image processing techniques produce edge maps to imitate drawing, such as the extended difference-of-Gaussians [Winnemöller et al. 2012], holistically-nested edge detection [Xie and Tu 2015], and contour extraction from images [Li et al. 2019].

Table 1. A taxonomy of drawing datasets. Our dataset includes both tracings and registered freehand drawings composed of vector-based timestamped strokes. The wide variety of prompts, participants, and stroke types presented in our dataset allows us to study general tendencies in drawing.

Dataset	Temporal data	Registration	Shading & texture	Prompts	People /prompt	Tracings	Freehand drawings
TU-Berlin [Eitz et al. 2012]	✓	✗	✗	250	~80	0	20k
Sketchy [Sangkloy et al. 2016]	✓	✗	✗	125	~600	0	75k
QuickDraw [Ha and Eck 2018]	✓	✗	✗	345	~145k	0	50m
BSDS500 [Arbeláez et al. 2011]	✗	✓	✗	500	4–9	2,696	0
PhotoSketch [Li et al. 2019]	✓	✓	✗	1,000	5	5,000	0
Princeton [Cole et al. 2008]	✗	✓	✗	12	1–11	208	208
Portrait [Berger et al. 2013]	✓	✓	✓	24	7	0	672
OpenSketch [Gryaditskaya et al. 2019]	✓	✓	✓	12	7–15	0	417
GMU [Yan et al. 2020]	✗	✓	✓	141	3–5	526	281
SpeedTracer (ours)	✓	✓	✓	100	11–21	1,210	288

Recently, CGDA techniques have become widely used to generate training data for sketching applications based on deep learning. For example, many use convolutional neural networks (CNN) to learn the inverse mapping from NPR to shape, such as reconstructing normal maps [Su et al. 2018] and dense 3D point clouds [Lun et al. 2017], predicting parameters for procedural modeling [Huang et al. 2017], and creating caricature face models [Han et al. 2017] and other freeform surfaces [Delanoy et al. 2018; Li et al. 2018] from drawings. While CGDA is commonly used as a proxy for drawing, this assumption lacks evaluation. In particular, although some CGDA techniques simulate the process of observational drawing [Liu et al. 2014], most of them neglect the order in which people draw and artists' drawing intention. Our work aims to understand the process of tracing and freehand drawing, which could inform new methods that better emulate drawing.

2.3 Tracing vs. Freehand Drawing

Tracing is a common tool and practice in art and design. Art students learn to observe proportions by looking at an object through a glass viewfinder with grids and drawing over it [Edwards 2012]. Architects also trace and selectively reuse lines for communicating and iterating design ideas [Johannessen and Van Leeuwen 2017]. Drawing interfaces such as ShadowDraw [Lee et al. 2011], DrawAFriend [Limpaecher et al. 2013], ShadowModel [Fan et al. 2013], and EZ-Sketching [Su et al. 2014] leverage a tracing metaphor to help users create more realistically proportioned drawings and models. Although tracing is easier to produce than freehand drawing, the art community has long discussed the distinction between these two forms. While tracing helps artists understand structure and perspective, it can also discourage detailed analysis of the work and become a crutch [Jamieson 2019]. Quantitative comparisons between tracing and freehand drawing mainly take place in psychology, focusing on eye-hand interactions, brain activities, and drawing skills using simple shapes [Gowen and Miall 2006, 2007; Ostrofsky et al. 2012]. However, they are not readily applicable to observational drawing and computer graphics.

This paper aims to advance understanding of how people trace and draw by collecting timestamped drawings and including a wider variety of prompts, participants, and stroke types. Our dataset allows the study of general consistency and variation between tracing and freehand drawing. Our detailed analysis provides insight into the drawing process and informs future drawing applications.

3 DATA COLLECTION

3.1 Study Design

3.1.1 Prompt selection. Past studies have used text prompts to collect sketches for recognition and image retrieval. Text prompts are less specific and encourage varying interpretations. As a result, sketches are often symbolic or caricatured. To encourage more realism and to support registration, we used visual prompts in a similar manner to the analysis of portrait drawing [Berger et al. 2013]. The image prompt was visible throughout a drawing session, so that the artist could observe and make various drawing decisions.

We included a total of 100 photorealistic images from various sources, with wide variations in appearance (unicolor diffuse vs. textured and specular) and content (single objects vs. complex scenes). We chose diverse shapes with abundant geometric and visual features rather than simple uniform shapes. Specifically, there were 33 rendered images from publicly available 3D models with appearance assets, 21 images from photometric stereo datasets [Shi et al. 2016; Toler-Franklin et al. 2007; Xiong et al. 2015; Zhang 2012] with associated normal maps, 30 images used for raster pencil drawing production [Lu et al. 2012], 9 images of 3D-printed objects from multiple viewpoints [Slavcheva et al. 2018], and 7 photographs of sculptures. The image content ranged from clay animals to furniture pieces to landscapes. Overall, this collection covers a wide range of shapes and scenes with diverse reflectance properties.

3.1.2 Constraints. Many previous studies instructed participants to draw only contour strokes. However, drawings typically employ a wider range of strokes including those that represent tone, shadows, textures, and those accentuating important details. In order to capture these different types of strokes, we did not limit the types that participants could draw. On the other hand, we found time limits used in previous studies helpful for encouraging participants to prioritize strokes and for collecting normalized data. We imposed a time limit of two minutes based on observations from a pilot study—most people could complete the overall content within two minutes and still had time to add details such as shading and texture.

3.2 Implementation

3.2.1 Drawing interface. We implemented SpeedTracer, a web-based drawing interface using Vue.js [You 2020]. Our drawing interface is composed of a square canvas with an image prompt and an optional blank canvas with user options on the sides of the webpage. We use square canvases with images scaled and padded to 800×800 as this is close to the prompts' original resolutions. Participants were instructed to trace over the image or draw on the blank canvas using a digital stylus. The interface provides the ability to select stroke width, color, and opacity, undo or redo a stroke, and clear the canvas. Stroke width is determined by multiplying user-specified width and pen pressure, allowing varying widths and tapering. Participants could change stroke color for artistic purposes and better tracing visibility. Strokes were stored in vector format in a database with timestamps, pressure, and appearance properties. Each participant traced or drew a prompt only once but could trace or draw multiple prompts. The prompts were ordered such that each prompt would be completed by a similar number of participants.

3.2.2 Participants. A total of 110 participants contributed to our dataset. Among them, 85 were professional illustrators, experienced students from art and architecture schools, students enrolled in a college drawing class, or amateur drawing enthusiasts. To enrich the dataset with easier-to-collect crowdsourced data used in machine learning research, the remaining 25 participants were recruited from Amazon Mechanical Turk for the tracing task. 39 participants contributed tracings using a 28-inch Microsoft Surface Studio with the Surface Pen in our lab. Others used their own tablets with stylus support such as an iPad with the Apple Pencil and Wacom tablets. All freehand drawings were collected remotely.

In all, there were 46 males, 62 females, and two who identified as non-binary gender. Their ages ranged from 18 to 61 years, with an average age of 27. Of these, 59 had two or more years of art training, 20 had one year, and 31 had none. Participants reported an average of four years of art training with the most experienced reporting 38 years. Each participant contributed 1–50 tracings or freehand drawings, with an average of 14. 13 participants were left-handed. Most participants were volunteers and not compensated. We paid each of the 25 turkers \$0.50 for 10 tracings, and each of the seven artists recruited from mihuashi.com \$7.68 for 15 freehand drawings.

3.2.3 Dataset. We excluded 187 tracings mostly by turkers because they included only an outline or content unrelated to the prompt. We also excluded all 15 freehand drawings by one participant who had no art training and drew only partial outlines. The resulting dataset contains a total of 1,210 tracings from 96 participants and 288 freehand drawings from 19 participants. Five artists contributed both tracings and freehand drawings. We collected fewer freehand drawings because they required more skill. We did not collect freehand drawings for 30 prompts of complex scenes because they would be hard to register. Each of the remaining 70 prompts of single objects has about 12 tracings and 4 freehand drawings.

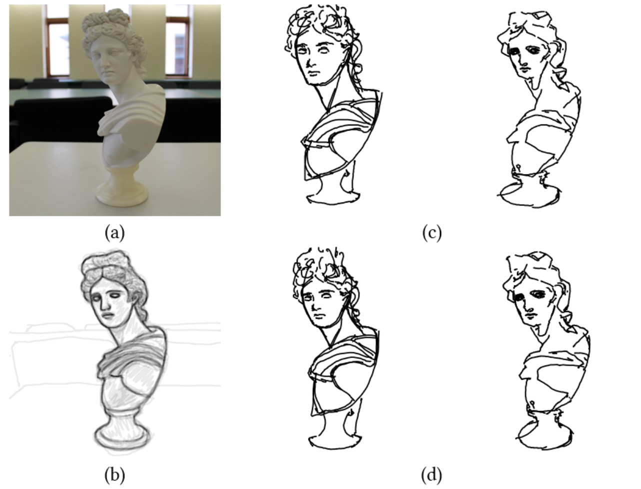

Fig. 2. Registration of freehand drawings using a composite of collected tracings as a guide. (a) Image prompt. (b) Composite of tracings. (c) Sample freehand drawings. (d) Same freehand drawings after registration.

3.3 Drawing Registration

In order to compare freehand drawings to tracings and CGDA results, we need to first register them to the image prompt. We use tracings to aid in registration by maximizing the correlation between each freehand drawing and a composite of all tracings of the same prompt over a displacement field. We implemented a coarse-to-fine optimization framework using `ImageRegistrationMethod()` in the Insight Segmentation and Registration Toolkit (ITK) [Kitware 2020]. The first step was to initialize a reasonable displacement field at full resolution. Our automatic initialization used a global affine transformation, followed by B-spline transformations at three scales from an image pyramid. Freehand drawings that included many background strokes occasionally led to a bad initialization. To address these cases, we manually labeled 18 to 42 fiducials in both the freehand drawing and the image prompt, and initialized the displacement field by fitting a thin plate spline model. The second step was to maximize the correlation between the freehand drawing and the tracing composite using gradient descent to obtain the optimal displacement field. We used the following parameters: 256×256 resolution, `SetMetricAsANTSNeighborhoodCorrelation(radius=16)`, and `SetOptimizerAsGradientDescent(learningRate=1, numberOfIterations=300, estimateLearningRate=EachIteration)`. Overall, the registered drawings share contours at nearby pixel locations (Fig. 2), which allows us to perform a spatio-temporal analysis on tracings and freehand drawings simultaneously.

4 DRAWING ANALYSIS

To compare tracing and freehand drawing, we examine similarity in content drawn, stroke usage, and progression of drawing over time. We then evaluate CGDA methods by finding the overlap between their output and both tracing and freehand drawing. Our analysis includes 70 prompts of single objects with 851 tracings and 288 freehand drawings.

4.1 Do People Draw Similar Content?

Our analysis begins with an exploration of the content and the way people trace and draw. We rasterized collected tracings and registered freehand drawings to enable pixelwise comparisons across drawings. We also leveraged temporal information to understand the order pixels are covered during the drawing process.

4.1.1 What people draw in common. Inspired by previous analysis [Cole et al. 2008], we explore common content within tracings, within freehand drawings, and between tracings and freehand drawings. We computed a histogram of pairwise distances between drawings for both tracing and registered freehand drawing (Fig. 3c). For every pixel in each drawing rasterized using one pixel-wide strokes, we recorded its closest chessboard distance in pixels to every other drawing of the same prompt. For both forms of drawing across all prompts, over 50% of the distances are close to a pixel in another drawing (within four pixels) and about 15% of the distances are far from any other drawing (more than 20 pixels).

To demonstrate similarities between tracing and freehand drawing, we first generated density maps and quantified their overlap (Fig. 3a). We superimposed all tracings of the same image prompt to create a density map for tracing, and similarly with all registered freehand drawings of the same prompt. To compensate for imperfect registration, we dilated the raster drawings for two iterations before superimposition so two strokes four pixels apart can overlap. We thresholded the density map using half the number of participants to obtain a binary commonly drawn region (CDR). We computed a histogram of pairwise distances between CDR pixels in tracing and freehand drawing, and found that over 80% of the distances are within four pixels.

Findings: In both tracing and freehand drawing, over half of the pixels are common content. People also draw similar common content between tracing and freehand drawing, and we observed that this mostly occurs at silhouettes and interior contours of objects.

4.1.2 When people draw common pixels. Drawing analysis should encompass both the end result and the drawing process. After exploring what people draw in common, we examine when they draw the common content. We used the interpolated timestamps for each pixel to analyze how many pixels people draw over time. For each drawing, we equally divided the temporal span into 25 bins and counted the number of all pixels drawn within each temporal bin. We also counted the number of CDR pixels based on their average timestamp for each prompt (Fig. 3d). We computed the Pearson correlation coefficient with a p-value between the number of pixels and the temporal bin index to understand how it changes over time. All analyses in this paper use a significance level of $p<0.05$.

We observed that the number of all pixels drawn in tracing increases over time in 18% of the prompts, decreases over time in 13%, and has no significant correlation with time in 69%. Similarly, we observed 20%, 14%, and 66% in freehand drawing. In contrast, the number of CDR pixels in tracing increases over time in 0% of the prompts, decreases over time in 93%, and has no significant correlation with time in 7%. Similarly, we observed 0%, 66%, and 34% in freehand drawing.

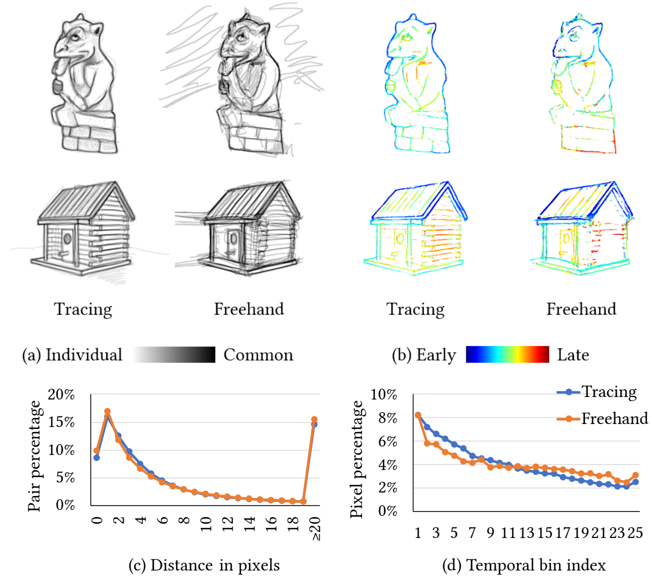

Fig. 3. Similarities between tracing and freehand drawing. (a) Density maps of two sample prompts. (b) Commonly drawn regions (CDR) colored by the average time at which each pixel is drawn. (c) Histograms of pairwise closest distances between pixels for all prompts. (d) Histograms of the number of CDR pixels over time for all prompts. Note the decreasing trend.

Findings: In both tracing and freehand drawing, the number of all pixels drawn is mostly uncorrelated with time. However, people tended to draw more common pixels early, and this tendency is stronger in tracing.

4.1.3 Where people draw over time. After exploring when the common content is drawn, we examine changes in focus over time as it relates to spatial locality. For CDR pixels of each prompt, we recorded their (x, y) coordinates relative to the top left corner and their distance to the barycenter of the drawing (*distc*). We then computed the Pearson correlation coefficient with a p-value between these properties and time (Fig. 4a). We found that among all prompts 90% in tracing have a significantly negative correlation between *distc* and time, and 84% in freehand drawing have a significantly positive correlation between y and time.

To understand this tendency at a finer level, we divided the temporal span of each drawing into ten equal bins and took snapshots of the drawing in progress. We then computed the convex hull area of each snapshot relative to that of the final drawing (Fig. 4b). The difference in the convex hull area between tracing and freehand drawing is significant in all temporal bins except for the first and the last, as suggested by a one-way analysis of variance (ANOVA) [De Veaux et al. 2020]. People tended to cover larger areas earlier in tracing, which reflects their behavior of drawing outside-to-inside.

Findings: People preferred finishing the entire outline first in tracing, as opposed to decomposing the prompt into parts in freehand drawing. This difference is more obvious when the prompt contains facial features catching artists' attention (Fig. 3b). A one way

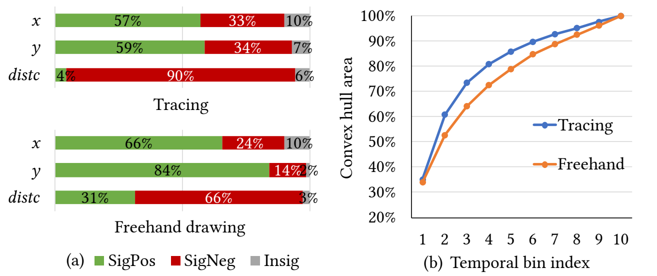

Fig. 4. Stroke placement over time. (a) Percentage of prompts with correlations between pixel location (x, y) in the CDR and time. *distc*: distance to the barycenter of the drawing. SigPos: significant positive correlation, SigNeg: significant negative correlation, Insig: insignificant correlation. (b) Convex hull area over time. In tracing, a larger area is covered earlier due to a stronger tendency to draw outside-to-inside than in freehand drawing.

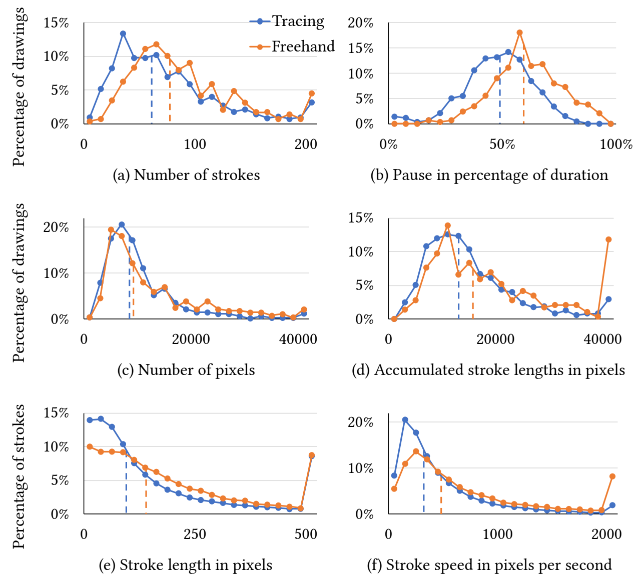

Fig. 5. Distributions across all drawings: (a) the number of strokes in each drawing, (b) total duration of pauses in each drawing, as percentage of time used for entire drawing, (c) the number of pixels in each rasterized drawing, (d) accumulated stroke lengths in each drawing, (e) the arc length of each stroke, and (f) the speed of each stroke. The dashed vertical line represents the median. The last point is an overflow bin except in (b).

ANOVA shows that the temporal difference of commonly drawn content between tracing and freehand drawing is significantly smaller if the prompt does not contain any facial features ($p = 0.0099$).

4.2 Do People Use Similar Strokes?

While a raster representation shows the consistency across drawings, information about individual strokes is lost. To understand whether people use similar strokes in tracing and freehand drawing, we analyze the distribution of vector stroke statistics, stroke ordering, and how stroke features change over time.

4.2.1 Distribution of basic statistics. We computed four statistics for each tracing and each registered freehand drawing: number of strokes, pauses as percentage of total drawing duration, number of pixels, and accumulated stroke lengths. We also computed two statistics for each stroke: arc length and speed. Fig. 5 shows their distributions in tracing and freehand drawing. In tracing, the number of strokes has a median of 61, pause duration 49%, number of pixels 8382, accumulated stroke lengths 13105, stroke length 96, and stroke speed 323. The corresponding statistics in freehand drawing are 78, 59%, 9246, 15769, 140, and 483. Since each prompt has a comparable number of collected drawings, these distributions can reflect general trends in tracing and freehand drawing.

Findings: We observed more flexible drawing behaviors in freehand drawing than tracing. In freehand drawing, people tended to draw more strokes and overdrawing is more common. People also tended to pause more possibly because they needed more time to consider the next part and its proportions when working with a blank canvas. Stroke length is more evenly distributed in freehand drawing than in tracing. This agrees with our observation—when tracing, people tended to draw long strokes at silhouettes followed by much shorter shading strokes, whereas people tended to use many connecting strokes at silhouettes in freehand drawing.

4.2.2 Comparing stroke ordering. To examine trends in the drawing process, we compare stroke ordering in tracing and freehand drawing by evaluating a list of heuristics used to assign order to vector strokes. Specifically, finding a Hamiltonian path in a graph of strokes can generate a plausible drawing order by minimizing an energy function using heuristics such as simplicity, proximity, and collinearity [Fu et al. 2011]. The energy function captures unary and binary costs and is defined as

$$E = w \sum_{i=0}^{n-1} c_{\text{ind}}(l_i)\theta(i) + \sum_{i=0}^{n-2} c_{\text{tra}}(l_i, l_{i+1}),$$

where l_i is the i^{th} stroke after ordering. $c_{\text{ind}}(l) = \eta c_{\text{str}}(l) + c_{\text{cir}}(l)$ captures the complexity of an individual stroke, where $c_{\text{str}}(l) = 1 - \textit{straightness}$ captures the deviation from a straight line and $c_{\text{cir}}(l) = \text{std}(\kappa)$ captures the deviation from a constant-curvature circle. $\theta(i) = 1 - i/n$ is used to enforce the heuristic of simplicity, i.e., simple strokes should be drawn first. $c_{\text{tra}}(l_i, l_j) = w_{\text{p}} c_{\text{pro}}(l_i, l_j) + (1 - w_{\text{p}}) c_{\text{col}}(l_i, l_j)$ captures the transition cost from stroke l_i to l_j, where the proximity term $c_{\text{pro}}(l_i, l_j)$ is the distance between the closest points on two strokes, and the collinearity term $c_{\text{col}}(l_i, l_j)$ is the positive angular difference between two endpoint tangents. We used the exact definitions and weights $w = 1, \eta = 0.1, w_{\text{p}} = 1/9$ by Fu et al., and we refer the reader to their paper for more details.

To understand if these heuristics play a similar role in tracing and freehand drawing, we performed a Monte Carlo sampling of all possible orderings of common strokes (defined in Section 4.4) in each drawing. For each ordering, we computed each term of the energy function using strokes before registration. A heuristic is considered important if the energy function takes a low value at the ground-truth ordering compared to other random orderings. As shown in the inset table, we computed a score for each heuristic and their combination as the percent of random orderings whose energy was above that of the ground-truth ordering.

Findings: Overall, these heuristics characterize stroke ordering better in freehand drawing than tracing. Proximity is most important for both, suggesting that temporally neighboring strokes tend to be spatially close. Collinearity comes second and is less important in tracings, which could be due to single long strokes in tracings versus shorter connecting strokes in freehand drawings. Simplicity does not characterize stroke ordering effectively.

	Tracing	Freehand
$c_{\mathrm{pro}}(l_i, l_j)$	97.89%	99.81%
$c_{\mathrm{col}}(l_i, l_j)$	90.55%	97.62%
E	81.35%	93.04%
$c_{\mathrm{str}}(l)$	44.16%	66.68%
$c_{\mathrm{cir}}(l)$	31.42%	51.33%

4.2.3 Stroke features over time. To further understand similarities between strokes used in tracing and freehand drawing, we examine how stroke features evolved throughout the drawing process. We extracted features [Mahoney 2018] from each stroke before registration and found their correlation to the *time* the stroke was completed. Stroke features included arc *length*, distance between two endpoints (*dist*), *duration*, *speed*, average *pressure*, stroke *width*, *opacity*, the number of points, screen-space angle of the vector connecting the two endpoints, average distance to image center, bounding box coordinates, positions of two endpoints, curvature approximated by nine uniformly sampled points on the stroke, and straightness defined as *dist*/*length*. For every drawing, we computed the Pearson correlation coefficient between each feature and *time* with a p-value. We excluded drawings in which the correlation coefficient could not be computed, e.g., missing pressure or strokes of a fixed width. For each feature, we computed the percentage of drawings with a significantly positive, significantly negative, or insignificant correlation. We reported all correlation results to avoid data dredging.

As shown in Fig. 6, in both tracing and freehand drawing, *opacity*, *duration*, *dist*, and *length* tended to decrease over time, whereas *width* and *speed* tended to increase over time. An inconsistency was *pressure*, as it mainly decreased over time only in tracing—people tended to apply more pressure when they started tracing contours than drawing them on a blank canvas.

Findings: In both tracing and freehand drawing, early strokes such as silhouettes tended to be long and drawn slowly using opaque colors. In contrast, late strokes such as hatching tended to be short and drawn quickly using translucent colors and wider strokes.

4.3 How Do Drawings Develop Over Time?

Changes in stroke features over time indicate that people might employ multiple types of strokes with different drawing intention. Therefore, it is important to understand if drawing intention evolves similarly in tracing and freehand drawing.

4.3.1 Labeling drawing intention. The first author manually labeled 10,832 strokes in 130 tracings and 3,400 strokes in 42 freehand drawings for 10 prompts. All strokes fall into two main categories: geometric contours and tone details. We further divided geometric contours into silhouettes (foreground-background separation), interior contours (depth discontinuity), and ridges and valleys (normal discontinuity). We also divided tone details into hatching (strokes

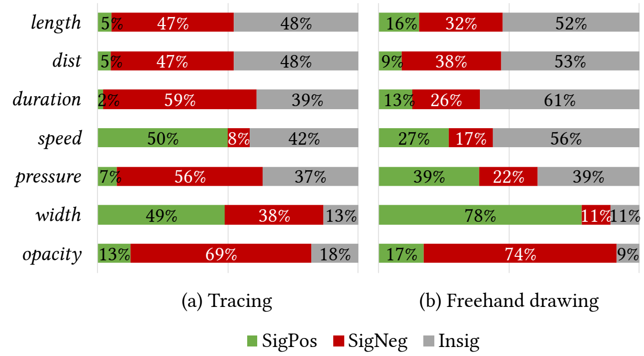

Fig. 6. Percentage of drawings with a significantly positive, significantly negative, or insignificant correlation between stroke features and *time*.

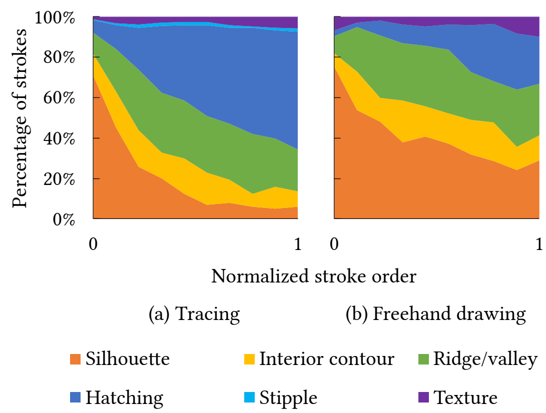

Fig. 7. Distribution of stroke types over time. Note the percentage of silhouettes at the beginning and hatching strokes towards the end.

for shading), stipple (dots for shading), and texture (color difference). Fig. 7 shows the percentage of strokes for each type over time.

Findings: Overall, we observed that drawings progressed similarly in tracing and freehand drawing. Silhouettes were dominant at the beginning and hatching strokes became more prominent towards the end. Interior contours and ridges and valleys appeared mainly in the middle of the drawing process. Stipple and texture strokes were drawn infrequently. People drew more hatching in tracing than in freehand drawing, possibly because freehand drawing required more skill and people had less time to depict shading.

4.3.2 Building stroke classifiers. Trends in how stroke features and labels change over time suggest the possibility of stroke classification. We built random forest classifiers to predict stroke type using features from Section 4.2 with 80% data for training and 20% for testing. We first performed robust standardization on the features using the 5$^{\mathrm{th}}$ and 95$^{\mathrm{th}}$ percentiles of each feature. During training, we used four-fold cross-validation to select optimal hyperparameters. We split data in three different ways—splitting prompts, drawings, or strokes. Table 2 summarizes the ten-time average test accuracy

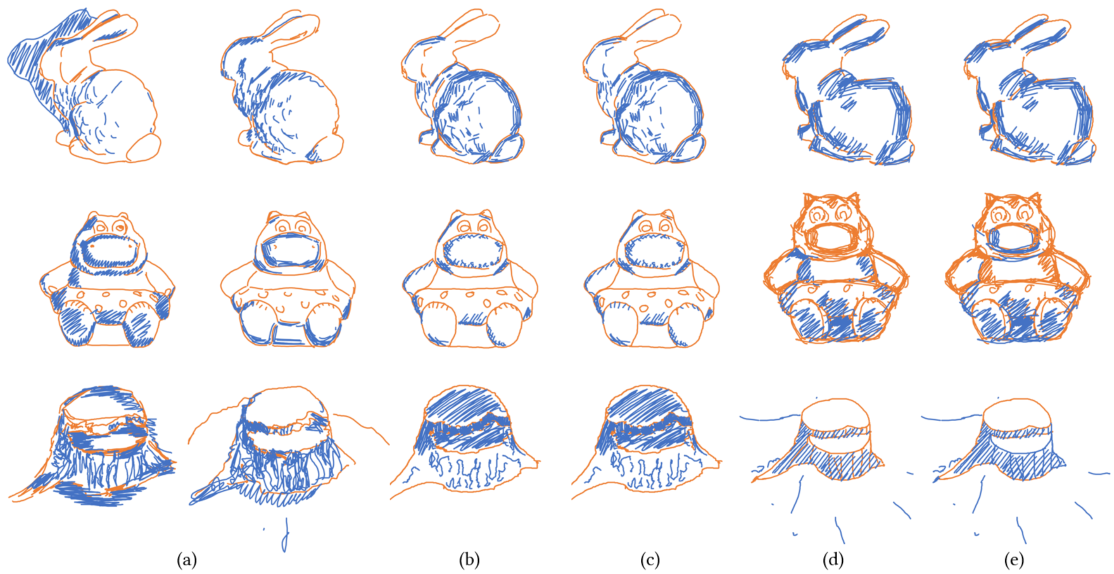

Fig. 8. Sample results of two stroke classifiers both predicting geometric contours (orange) and tone details (blue). The first classifier used 80% tracings as training data and 20% tracings as test data. The second classifier used all tracings as training data and all freehand drawings as test data. (a) Labeled tracings used in training for both classifiers. (b) Labeled test tracings for the first classifier. (c) Predicted labels by the first classifier. (d) Labeled test freehand drawings for the second classifier. (e) Predicted labels by the second classifier, which shows models learned from tracings can be transferred to freehand drawings.

Table 2. Average test accuracy of random forest classifiers for predicting stroke labels. The classifiers were trained using two labels and six labels, and with different splits of data.

	Tracing		Freehand drawing	
	2 labels	6 labels	2 labels	6 labels
Split prompts	75.57%	59.49%	84.24%	48.96%
Split drawings	80.16%	66.67%	78.30%	53.52%
Split strokes	90.11%	82.99%	89.06%	69.40%

using two (contours, details) or six (silhouettes, interior contours, ridges and valleys, hatching, stipple, texture) labels.

The classifiers distinguished geometric contours from tone details with an accuracy of about 80% on splits of prompts and drawings, and an accuracy of about 90% on splits of strokes. The most important features for classification were *time*, *length*, *duration*, *speed*, and *straightness*. Fig. 8 shows sample training data, ground truth for testing, and predictions using the classifier with 80% drawings as training data. In another experiment, we trained a classifier using all tracings and tested it on freehand drawings, which achieved an average accuracy of 76.32% over all prompts.

Findings: A stroke's temporal features can effectively characterize drawing intention in both tracing and freehand drawing. Furthermore, the high accuracy of classifying strokes in a freehand drawing with a classifier trained on tracings suggests that models learned from tracings can be transferred to freehand drawings.

4.4 How Much Do CGDA and Drawing Overlap?

After comparing tracing and freehand drawing, we revisit current methods for computer-generated drawing approximation as they are commonly used as a proxy for hand drawings. CGDA methods fall into two classes: NPR methods that assume an underlying 3D representation, and image processing methods that operate directly on images. We compare both tracings and freehand drawings to CGDA results to quantitatively establish whether their use is a reasonable approximation to either form of drawing.

4.4.1 Precision and recall. Inspired by the analysis approach of Cole et al. [2008], we start with a pixelwise comparison between rasterized hand-drawn data and CGDA output. We computed precision and recall, where precision is the fraction of pixels in a CGDA output that fall in a neighborhood of any pixel in the drawing. We rasterized each stroke in the drawing and considered it to be captured by CGDA output if over half of its pixels fall in a neighborhood of any pixel in the CGDA image. We then measured recall as the fraction of pixels on such strokes captured by the CGDA output. We generated one-pixel wide CGDA output using five widely used methods: suggestive contours (SC) [DeCarlo et al. 2003], ridges and valleys (RV) [Ohtake et al. 2004], and apparent ridges (AR) [Judd et al. 2007] for 33 prompts rendered from 3D models as well as Canny edges [Canny 1986] and holistically-nested edges (HED) [Xie and Tu 2015] for all 70 prompts (Fig. 10). We chose thresholds for the algorithms such that the number of output pixels was closest to the median number of pixels in all tracings and freehand drawings of the same prompt. We used a 9×9 neighborhood based on observations from the pairwise distances and a hybrid raster-vector representation to compensate for imperfect registration.

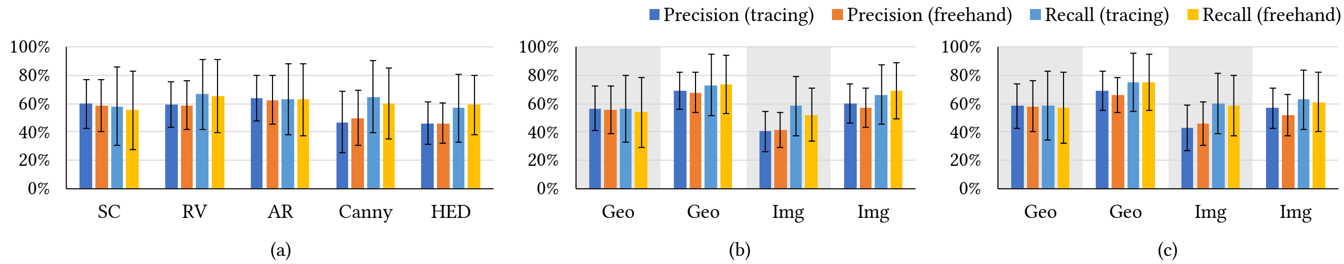

Fig. 9. Precision and recall of NPR and image processing output compared with tracing and freehand drawing. (a) Average precision and recall for five types of computer-generated output. (b) Average precision and recall on images with (gray) and without background. (c) Average precision and recall on images with (gray) and without texture. Obj: object-space NPR output, i.e., SC+RV+AR. Img: image processing output, i.e., Canny+HED.

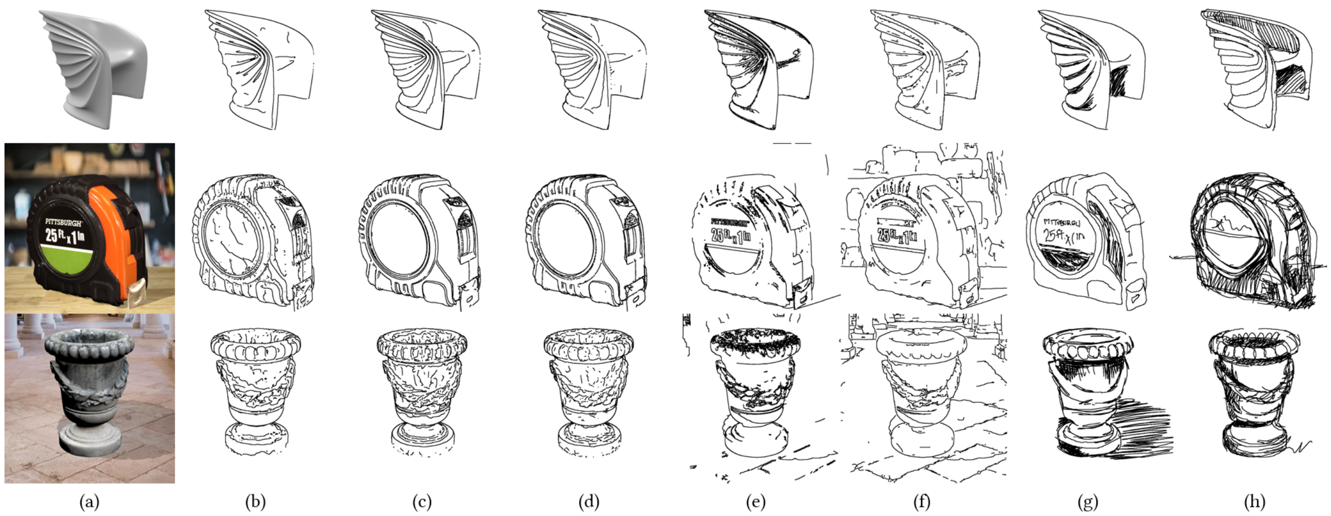

Fig. 10. Computer-generated results and drawings. (a) Image prompts. (b) Suggestive contours (SC). (c) Ridges and valleys (RV). (d) Apparent ridges (AR). (e) Canny edges. (f) Thinned holistically-nested edges (HED). (g) Sample tracings. (h) Sample registered freehand drawings. A 3D model is required for (b)–(d). Chair model by KarloBalboa on TurboSquid (royalty-free). Tape model by Artec Group, Inc. (CC BY 3.0). Vase model by 3dhdscan on Sketchfab (CC BY 4.0).

Findings: As shown in Fig. 9a, all CGDA methods used in our analysis achieved comparable precision and recall on tracing and freehand drawing—about 60% of drawn content is consistent with computer-generated output. A one-way ANOVA suggests that NPR algorithms achieved higher precision than image processing algorithms although the difference in recall was not as significant. SC and RV achieved comparable precision whereas AR achieved higher precision. RV and AR achieved comparable recall whereas SC achieved lower recall. Canny and HED achieved comparable precision but Canny achieved higher recall than HED. AR is the best proxy for drawing among these methods.

4.4.2 Image attributes. We observed that CGDA algorithms achieved different performance on different images. To better understand the disparity, we analyze how image attributes relate to the performance of CGDA algorithms. We divided the image prompts based on their attributes and computed the average precision and recall for CGDA results (Fig. 9bc). Among all 70 prompts, 39 have a background and 45 contain texture. Using one way ANOVA, we found that for both tracing and freehand drawing, NPR output achieved lower precision and recall on image prompts containing background or texture. This agrees with our intuition because NPR typically does not consider image background or object texture. Image processing-based CGDA output achieved lower precision and recall on prompts with background or texture for both forms of drawing, although the difference in recall is not as prominent. This can be explained by the fact that image processing methods often include extraneous edges from background and texture areas that people do not necessarily draw.

Findings: The disparity between precision and recall on different prompts shows a gap between drawings and computer-generated results, and reveals that people employ diverse drawing strategies when depicting complex prompts, e.g., non-diffuse and textured objects. Such prompts are lacking in previous drawing studies, and therefore the richness of hand-drawn data was not captured. This analysis shows our dataset's value for identifying extensions to NPR and image processing algorithms that better approximate drawing.

4.4.3 Correlation with common strokes. In Section 4.1 (Fig. 3), we discovered that people tended to depict a common set of contours, and here we examine how these common strokes correspond to computer-generated output. We used three density maps (tracing, freehand drawing, CGDA) to compute two per-stroke metrics: common score and CGDA score. We rasterized each stroke, recorded the maximum tracing or freehand drawing density in a 5×5 neighborhood of each pixel along the stroke, and averaged the density values over the entire stroke to compute the common score. Similarly, we computed the CGDA score with a density map generated by combining computer-generated results of all NPR and image processing algorithms. Then, we computed the Pearson correlation coefficients with a p-value between both scores and stroke order, as well as between the two scores for all strokes in each drawing.

As shown in Fig. 11d, the common score and the CGDA score tend to decrease over time. This decrease is more pronounced in tracing. There is a significantly positive correlation between the common score and the CGDA score in 85% or more of the tracings and freehand drawings. On average, strokes captured by CGDA (score≥0.5) account for 62% pixels in tracing and 60% pixels in freehand drawing, suggesting that about 40% of hand-drawn content remains uncaptured.

Findings: CGDA methods can capture content that people commonly draw, but not the diverse artistic choices that people employ.

4.4.4 Comparison with NPR hatching. We observed that in addition to contours, people often depicted tone details using hatching. Therefore, we compare drawing with NPR hatching [Praun et al. 2001] to determine how well NPR captures drawn shading. We first created a composite of all drawings for each prompt and blurred the composite using a 25×25 box filter to get areas of dense stroke coverage. Then, we generated NPR hatching using the same configurations as the prompts and varied the light intensity so that the hatching image blurred with the same box filter had a similar level of overall response as the blurred drawing composite (Fig. 12). Similar to the previous definition of precision and recall, we computed the fraction of pixels in the blurred NPR hatching whose response was within 10% of total response range at the same pixel in the blurred drawing composite and vice versa.

In both tracing and freehand drawing, less than half of drawn shading is consistent with the hatching algorithm. For example, the algorithm generates uniform hatching in the hair and the face, whereas people place strokes primarily in the hair to demonstrate texture. Furthermore, people tend to emphasize the contrast near highlights, whereas the hatching algorithm tends to avoid them.

Findings: The hatching algorithm does not understand the semantics of the prompt, resulting in hatching patterns different from what people draw. The diversity of shading techniques presented in our dataset can inform new NPR hatching algorithms.

4.4.5 Local properties. Since computer-generated output is derived from local properties, it is necessary to examine which ones contribute most to stroke placement. Following notations in previous analysis [Cole et al. 2008], we computed four image-space properties for all prompts, including image luminance (ImgLuminance), gradient magnitude (ImgGradMag), and maximum and minimum eigenvalues of the image Hessian (ImgMaxCurv and ImgMinCurv).

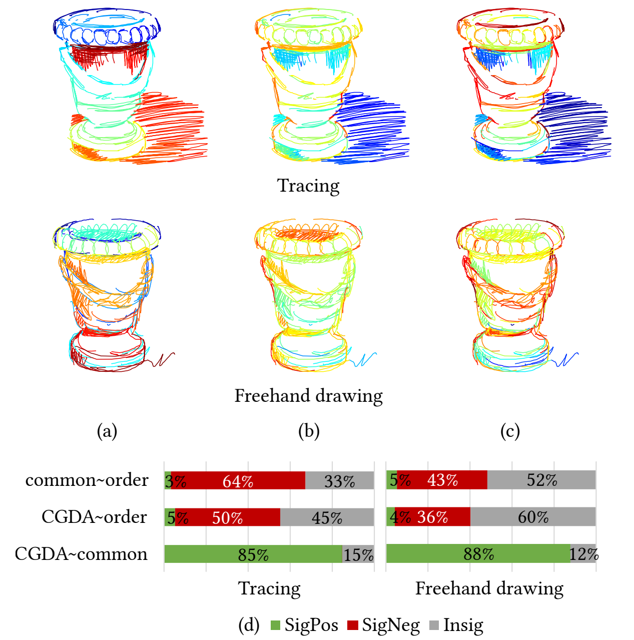

Fig. 11. Correlation between stroke order, common score, and CGDA score. (a) A drawing pseudocolored by stroke order. Warmer color represents later drawn strokes. (b) A drawing pseudocolored by common score. Warmer color represents more commonly drawn strokes. (c) A drawing pseudocolored by CGDA score. Warmer color represents strokes better captured by CGDA. (d) Percentage of drawings with a significantly positive, significantly negative, or insignificant correlation between common score, CGDA score, and stroke order. Note that common strokes tend to be captured by CGDA.

We computed the dot product between normal and view vectors ($N \cdot V$) for 21 photometric stereo and 33 rendered prompts. We also computed nine object-space properties for the rendered prompts, including the maximum, minimum, mean, and Gaussian curvatures, the derivative of the maximum curvature in the corresponding direction, the largest view-dependent principal curvature (ViewDepCurv) and its derivative in the corresponding direction, the radial curvature and its derivative in the radial direction. We used a random forest regressor to construct a hierarchy of conditionals on local properties to predict the probability of stroke placement for each prompt and understand feature importance of each local property.

The six local properties most important for predicting stroke placement are shown in the inset table. ImgGradMag, $N \cdot V$, and ImgLuminance are most important; ImgGradMag is not dominantly important. ViewDepCurv is the fourth important local property, which agrees with our observation that apparent ridges better approximate drawing than other CGDA methods.

	Tracing	Freehand
ImgGradMag	33.92%	29.36%
$N \cdot V$	30.10%	33.15%
ImgLuminance	21.86%	25.50%
ViewDepCurv	17.53%	14.07%
ImgMinCurv	4.76%	4.39%
ImgMaxCurv	4.38%	3.75%

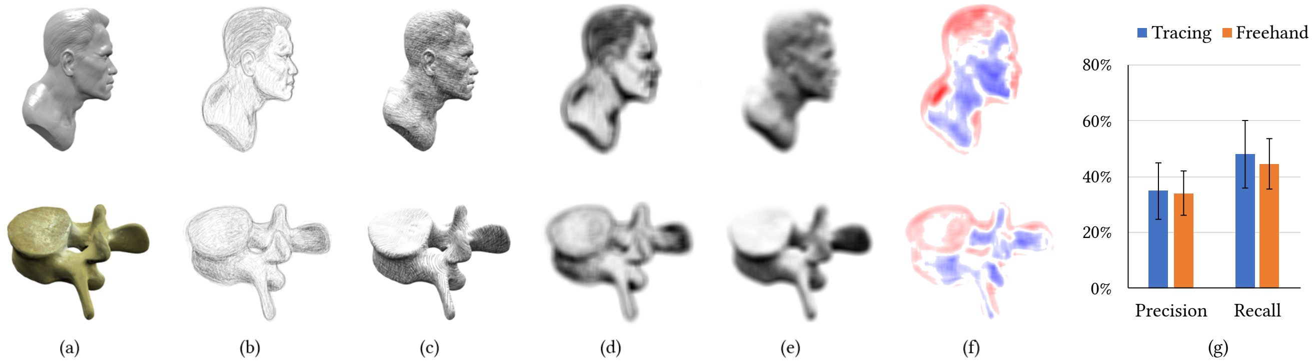

Fig. 12. Comparison between drawn and computer-generated hatching. (a) Image prompts. (b) Composites of all tracings of the prompt. (c) Results of the hatching algorithm. (d) Blurred version of (b). (e) Blurred version of (c). (f) Differences between (d) and (e). White indicates the difference is within 10% of the response range. Blue indicates a stronger response in the NPR hatching. Red indicates a stronger response in the drawing composite. (g) Average precision and recall of the hatching algorithm on all 33 rendered images. Bust model by atanazy on TurboSquid (royalty-free). Bone model by Artec Group, Inc. (CC BY 3.0).

Findings: Each local property has comparable importance when used to predict stroke placement in tracing and freehand drawing, indicating similarity between the two forms. The high importance of $N \cdot V$ and ImgLuminance shows the value of strokes for shading in our dataset compared to those containing only contours.

5 CONCLUSION

In this paper, we present an analysis of a new dataset that reveals the similarities and differences between freehand drawing, tracing, and computer-generated drawing approximation. Our wide range of image prompts and different types of strokes provide a valuable supplement to previous sketch datasets and informs new areas of drawing research. From our analysis, we find that tracing is similar to freehand drawing in terms of temporal tendencies and representation choices, and thus can serve as a viable proxy for drawing. An examination of both tracing and freehand drawing suggests that people's intention evolves over time, which can be characterized by similar spatio-temporal stroke features. By comparing hand-drawn data with computer-generated output, we find that current NPR and image processing methods only capture 60% of drawn pixels on average, highlighting the great value of collecting hand-drawn data.

Our study has several limitations. There might exist potential bias due to our setup of data collection—drawing habits may differ between digitized strokes from stylus input versus pen on paper, as well as when no time limit is imposed. Our registration of freehand drawings to tracings may introduce registration errors and this form of normalization might mask artistic styles. We did not consider undo operations that might reveal useful information about revision and refinement. Our findings in comparing drawing to CGDA are specific to the commonly-used algorithms that we considered.

Our dataset and analysis are useful and have implications in a variety of applications, such as training data for data-driven NPR methods that better emulate the drawing process, as well as customized treatment and recognition of different types of strokes in sketching interfaces. For example, the temporal information and shading strokes from our dataset can be used to learn order assignment and drawing density through image translation networks. Our stroke classifiers are also useful for understanding drawing intention, making it possible to infer shape from strokes of different styles. Other interesting future work includes analyzing construction lines and junctions, comparing artists of varied expertise, and understanding distortion in freehand drawing. The dataset and analysis code are available at https://github.com/zachzeyuwang/tracing-vs-freehand.

ACKNOWLEDGMENTS

This work was partially supported by National Science Foundation award #1942257. We thank Justine Luo for her help in rendering image prompts. We thank Theodore Kim and Benedict Brown for their valuable feedback. We also thank the reviewers for their helpful suggestions and all the participants who contributed to this dataset.

REFERENCES

Pablo Arbeláez, Michael Maire, Charless Fowlkes, and Jitendra Malik. 2011. Contour Detection and Hierarchical Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 5 (May 2011), 898–916.
Itamar Berger, Ariel Shamir, Moshe Mahler, Elizabeth Carter, and Jessica Hodgins. 2013. Style and Abstraction in Portrait Sketching. *ACM Trans. Graph.* 32, 4, Article 55 (July 2013), 12 pages.
John Canny. 1986. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8, 6 (Nov 1986), 679–698.
Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz. 2008. Where Do People Draw Lines? *ACM Trans. Graph.* 27, 3, Article 88 (Aug 2008), 11 pages.
Richard D. De Veaux, Paul F. Velleman, and David E. Bock. 2020. *Stats: Data and Models, 5th Edition.* Pearson.
Doug DeCarlo, Adam Finkelstein, Szymon Rusinkiewicz, and Anthony Santella. 2003. Suggestive Contours for Conveying Shape. *ACM Trans. Graph.* 22, 3 (July 2003), 848–855.
Johanna Delanoy, Mathieu Aubry, Phillip Isola, Alexei A. Efros, and Adrien Bousseau. 2018. 3D Sketching Using Multi-View Deep Volumetric Prediction. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1, Article 21 (July 2018), 22 pages.
Betty Edwards. 2012. *Drawing on the Right Side of the Brain.* Penguin Group.
Mathias Eitz, James Hays, and Marc Alexa. 2012. How Do Humans Sketch Objects? *ACM Trans. Graph.* 31, 4, Article 44 (July 2012), 10 pages.
Lubin Fan, Ruimin Wang, Linlin Xu, Jiansong Deng, and Ligang Liu. 2013. Modeling by Drawing with Shadow Guidance. *Computer Graphics Forum* 32, 7 (2013), 157–166.
Hongbo Fu, Shizhe Zhou, Ligang Liu, and Niloy J. Mitra. 2011. Animated Construction of Line Drawings. *ACM Trans. Graph.* 30, 6, Article 133 (Dec 2011), 10 pages.

Emma Gowen and R. Chris Miall. 2006. Eye-Hand Interactions in Tracing and Drawing Tasks. *Human Movement Science* 25, 4 (2006), 568–585.
Emma Gowen and R. Chris Miall. 2007. Differentiation Between External and Internal Cuing: An fMRI Study Comparing Tracing with Drawing. *NeuroImage* 36, 2 (2007), 396–410.
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, Article 232 (Nov 2019), 16 pages.
David Ha and Douglas Eck. 2018. A Neural Representation of Sketch Drawings. In *Proceedings of the International Conference on Learning Representations*.
Xiaoguang Han, Chang Gao, and Yizhou Yu. 2017. DeepSketch2Face: A Deep Learning Based Sketching System for 3D Face and Caricature Modeling. *ACM Trans. Graph.* 36, 4, Article 126 (July 2017), 12 pages.
Haibin Huang, Evangelos Kalogerakis, Ersin Yumer, and Radomir Mech. 2017. Shape Synthesis from Sketches via Procedural Models and Convolutional Networks. *IEEE Transactions on Visualization and Computer Graphics* 23, 8 (Aug 2017), 2003–2013.
Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-To-Image Translation with Conditional Adversarial Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
David Jamieson. 2019. Why Learn to Draw When You Can Trace? https://vitruvianstudio.com/why-learn-to-draw-when-you-can-trace/. Accessed Aug 9, 2020.
Christian Mosbæk Johannessen and Theo Van Leeuwen. 2017. *The Materiality of Writing: A Trace Making Perspective*. Routledge.
Tilke Judd, Frédo Durand, and Edward Adelson. 2007. Apparent Ridges for Line Drawing. *ACM Trans. Graph.* 26, 3, Article 19 (July 2007), 8 pages.
Evangelos Kalogerakis, Derek Nowrouzezahrai, Simon Breslav, and Aaron Hertzmann. 2012. Learning Hatching for Pen-and-Ink Illustration of Surfaces. *ACM Trans. Graph.* 31, 1, Article 1 (Feb 2012), 17 pages.
Kitware. 2020. Insight Toolkit. https://itk.org/. Accessed Oct 20, 2020.
Yong Jae Lee, C. Lawrence Zitnick, and Michael F. Cohen. 2011. ShadowDraw: Real-Time User Guidance for Freehand Drawing. *ACM Trans. Graph.* 30, 4, Article 27 (July 2011), 10 pages.
Changjian Li, Hao Pan, Yang Liu, Xin Tong, Alla Sheffer, and Wenping Wang. 2018. Robust Flow-Guided Neural Prediction for Sketch-Based Freeform Surface Modeling. *ACM Trans. Graph.* 37, 6, Article 238 (Dec 2018), 12 pages.
Mengtian Li, Zhe Lin, Radomir Mech, Ersin Yumer, and Deva Ramanan. 2019. Photo-Sketching: Inferring Contour Drawings from Images. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*. 1403–1412.
Alex Limpaecher, Nicolas Feltman, Adrien Treuille, and Michael Cohen. 2013. Real-Time Drawing Assistance through Crowdsourcing. *ACM Trans. Graph.* 32, 4, Article 54 (July 2013), 8 pages.
Difan Liu, Mohamed Nabail, Aaron Hertzmann, and Evangelos Kalogerakis. 2020. Neural Contours: Learning to Draw Lines from 3D Shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
Jingbo Liu, Hongbo Fu, and Chiew-Lan Tai. 2014. Dynamic Sketching: Simulating the Process of Observational Drawing. In *Proceedings of the Workshop on Computational Aesthetics*. ACM, New York, NY, USA, 15–22.
Cewu Lu, Li Xu, and Jiaya Jia. 2012. Combining Sketch and Tone for Pencil Drawing Production. In *Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering*. Eurographics Association, Goslar, Germany, 65–73.
Zhaoliang Lun, Matheus Gadelha, Evangelos Kalogerakis, Subhransu Maji, and Rui Wang. 2017. 3D Shape Reconstruction from Sketches via Multi-view Convolutional Networks. In *Proceedings of the International Conference on 3D Vision*. 67–77.
James M Mahoney. 2018. *The V-Sketch System: Machine Assisted Design Exploration in Virtual Reality*. Master's thesis. Cornell University.
Yutaka Ohtake, Alexander Belyaev, and Hans-Peter Seidel. 2004. Ridge-Valley Lines on Meshes via Implicit Surface Fitting. *ACM Trans. Graph.* 23, 3 (Aug 2004), 609–612.
Justin Ostrofsky, Aaron Kozbelt, and Angelika Seidel. 2012. Perceptual Constancies and Visual Selection as Predictors of Realistic Drawing Skill. *Psychology of Aesthetics, Creativity, and the Arts* 6, 2 (2012), 124.
Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. 2001. Real-time Hatching. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*. ACM, New York, NY, USA, 579–584.
Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. 2016. The Sketchy Database: Learning to Retrieve Badly Drawn Bunnies. *ACM Trans. Graph.* 35, 4, Article 119 (July 2016), 12 pages.
Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. 2016. A Benchmark Dataset and Evaluation for Non-Lambertian and Uncalibrated Photometric Stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3707–3716.
Miroslava Slavcheva, Wadim Kehl, Nassir Navab, and Slobodan Ilic. 2018. SDF-2-SDF Registration for Real-Time 3D Reconstruction from RGB-D Data. *International Journal of Computer Vision* 126, 6 (June 2018), 615–636.
Qingkun Su, Wing Ho Andy Li, Jue Wang, and Hongbo Fu. 2014. EZ-Sketching: Three-Level Optimization for Error-Tolerant Image Tracing. *ACM Trans. Graph.* 33, 4, Article 54 (July 2014), 9 pages.
Wanchao Su, Dong Du, Xin Yang, Shizhe Zhou, and Hongbo Fu. 2018. Interactive Sketch-Based Normal Map Generation with Deep Neural Networks. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1, Article 22 (July 2018), 17 pages.
Masaki Suwa and Barbara Tversky. 1997. What Do Architects and Students Perceive in Their Design Sketches? A Protocol Analysis. *Design Studies* 18, 4 (1997), 385–403.
Corey Toler-Franklin, Adam Finkelstein, and Szymon Rusinkiewicz. 2007. Illustration of Complex Real-World Objects Using Images with Normals. In *Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering*. ACM, New York, NY, USA, 111–119.
Holger Winnemöller, Jan Eric Kyprianidis, and Sven C. Olsen. 2012. XDoG: An EXtended Difference-of-Gaussians Compendium Including Advanced Image Stylization. *Computers & Graphics* 36, 6 (2012), 740–753.
Saining Xie and Zhuowen Tu. 2015. Holistically-Nested Edge Detection. In *Proceedings of the IEEE International Conference on Computer Vision*. 1395–1403.
Ying Xiong, Ayan Chakrabarti, Ronen Basri, Steven J. Gortler, David W. Jacobs, and Todd Zickler. 2015. From Shading to Local Shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37, 1 (Jan 2015), 67–79.
Chuan Yan, David Vanderhaeghe, and Yotam Gingold. 2020. A Benchmark for Rough Sketch Cleanup. *ACM Trans. Graph.* 39, 6, Article 163 (Nov 2020), 14 pages.
Evan You. 2020. Vue.js. https://vuejs.org/. Accessed Oct 20, 2020.
Li Zhang. 2012. Photometric Stereo. http://pages.cs.wisc.edu/~lizhang/courses/cs766-2012f/projects/phs/index.htm. Accessed Apr 12, 2019.