DrawMyPhoto: Assisting Novices in Drawing from Photographs

Blake Williford
Sketch Recognition Lab
Texas A&M University
bwilliford@gmail.com

Abhay Doke
Microsoft
abhay.doke@microsoft.com

Michel Pahud
Microsoft Research
mpahud@microsoft.com

Ken Hinckley
Microsoft Research
kenh@microsoft.com

Tracy Hammond
Sketch Recognition Lab
Texas A&M University
hammond@tamu.edu

ABSTRACT

We present *DrawMyPhoto*, an interactive system that can assist a drawing novice in producing a quality drawing by automatically parsing a photograph in to a step-by-step drawing tutorial. The system utilizes image processing to produce distinct line work and shading steps from the photograph, and offers novel real-time feedback on pressure and tilt, along with grip suggestions as the user completes the tutorial. Our evaluation showed that the generated steps and real-time assistance allowed novices to produce significantly better drawings than with a more traditional grid-based approach, particularly with respect to accuracy, shading, and details. This was confirmed by domain experts who blindly rated the drawings. The participants responded well to the real-time feedback, and believed it helped them learn proper shading techniques and the order in which a drawing should be approached. We saw promising potential in the tool to boost the confidence of novices and lower the barrier to artistic creation.

CCS Concepts

•Human-centered computing → Human computer interaction (HCI); Interactive systems and tools; Interaction techniques; •Applied computing → Computer-assisted instruction;

Author Keywords

Human-Computer Interaction; Drawing; Sketching; Art Education; Sketch Recognition; User Experience Design; Image Processing

INTRODUCTION

Drawing can be a valuable lifelong skill but it remains difficult for many people to learn. Research has shown that many children and adolescents struggle with self-efficacy [3] regarding their drawing ability [31, 20], and that there tends to be a loss of interest in artistic and creative pursuits by early adolescence [16]. There is therefore a need for lowering the barrier to entry with respect to drawing, and giving such people a much-needed boost in drawing self-efficacy.

Photographs are ubiquitous in modern society and nearly everyone has treasured photos of their pets, friends, family, and travels. While many applications make it easy to use image processing to give photos a "sketched" or rendered look (e.g. [42, 28]), this bypasses a learning opportunity and may not give the person the same sense of pride or sense of value in the ensuing result. Previous work has shown that when people invest time and effort in to something they build or create, it holds significantly more value to them, such as the "Ikea Effect" explored by Norton et al. [29] and the notion of "digital possessions" from Belk [4].

The goal of this work was to develop a system for converting a photo in to a set of steps that guide a user in drawing the photo in a way that approaches how professional artists draw [13, 12, 32]. In doing so, we believed that the user could gain a strong boost in drawing self-efficacy, which is an important factor when learning drawing.

To accomplish this, we used image processing to automatically generate a line work step, three shading steps, and a final details step from the original photograph along with novel real-time feedback on pressure and tilt, and pen grip suggestions. We also provide grids inspired by a popular grid-based approach common in art education, along with an underlay feature which can assist novices who prefer to trace.

We conducted a user study with 20 participants split into two groups (a control group with minimal assistance and an experimental group with the full *DrawMyPhoto* system) to determine the efficacy of the system in achieving our goals and providing a rewarding experience for novices. The methodology for answering our research questions included expert ratings, a questionnaire and semi-structured interviews, a log analysis, and observations.

C&C '19, June 23–26, 2019, San Diego, CA, USA

 ISBN 978-1-4503-5917-7/19/06…$15.00

DOI: https://doi.org/10.1145/3325480.3325507

Our contributions include:

- A description of the system, detailing how readily-available image processing techniques can be used in a unique way to provide the automatic generation of step-by-step guidance in drawing a photo. This approach mirrors the order and manner in which many professional artists draw.
- The design rationale behind novel real-time feedback mechanisms for pressure and tilt which guide users in proper shading technique. Such feedback can be expanded upon by other designers and researchers for similar applications.
- Evidence that the system was a rewarding experience for novices and allowed them to produce quality drawings. Expert ratings were significantly higher ($p < 0.01$) for the group with full assistance with respect to overall quality, accuracy, shading, and details. Many of the participants who used the system also self-reported they had learned proper shading techniques and the order in which to approach drawings.

Together, these contributions illustrate how *DrawMyPhoto* offers a promising and rewarding experience for novices to learn how to draw using their own photos—photos they already have a strong emotional investment in.

RELATED WORK

There have been a variety of interactive systems for both teaching drawing skills and providing assistance for drawing, as well as painting. In general, these systems sit on a spectrum from intelligent tutoring systems (ITS) [34] to creativity support tools (CST) [33] and many systems combine aspects of both. *DrawMyPhoto* can be considered an ITS, but draws inspiration from CSTs in this space, many of which are powerful tools, but have questionable educational outcomes. We will discuss the strengths and weaknesses of these systems in this section, and how our system builds on this prior work by placing more emphasis on shading technique.

The Drawing Assistant [22] processes photographs to provide guidance and feedback for drawing, and is in principle very similar to *DrawMyPhoto*. We build on this work to offer assistance more focused on shading and details, not just proportion, alignment, and line work.

ICanDraw [10] is a step-by-step tutorial system for drawing human faces that utilizes sketch recognition and template-matching to guide the user in properly proportioning faces as they draw them. *ICanDraw* is notable for how well it encourages accuracy, an important aspect of quality drawings, and particularly important for the nuances of the human face. However it does not provide any assistance for shading, and its feedback is limited. It is also limited to the domain of faces and portraits. *EyeCanDraw* [8] is a very similar system focused on the human eye, but has many of the same limitations.

Sketch Sketch Revolution [15], is a tutorial system driven by recording strokes of an artist and then allowing a user to replicate the drawing stroke-by-stroke. It offers unique feedback on stroke speed and accuracy, along with beautification, but no feedback on pressure or tilt. This can be a powerful tool for learning, however the user is merely replicating an existing drawing instead of learning to make their own drawing from a reference, which was the goal of *DrawMyPhoto*.

SketchTivity [25, 37, 38, 17, 21] is an ITS for teaching design sketching, a more conceptual form of drawing commonly used by industrial designers and engineers. It does not cover shading, with a focus solely on line work and perspective sketching. It does, however, provide real-time feedback on accuracy, line quality, and speed of strokes. Since this conforms users to a specific style of drawing, we wanted to avoid that type of feedback and instead focus on pressure and tilt feedback which can help users shade properly.

PortraitSketch [39] is a tool that assists users in drawing portraits. As the user traces the photograph, the system beautifies their strokes and adjusts them to be more matched with the contours of the underlying image. This includes shading strokes, which are automatically straightened to make the shading look cleaner and more professional. Tracing can be a great way to boost confidence for novices. Users from their study were very pleased with the results, and the tool is a strong approach to lowering the barrier to art for novices. *Painting with Bob* [5] is a similar system that allows for novices to paint their photographs. By utilizing the underlying image, novices can attain impressive results while also maintaining their own style. We build on this work by also aiming to allow users to preserve their own style, but with no beautification.

How2Sketch [19] is an excellent system for generating step-by-step drawing tutorials from a 3D model. However, the system does not offer any interactive feedback and it only produces static steps. The step-by-step approach is important in drawing, and we build on this work with a system that is fully interactive.

ShadowDraw [26] is a novel system built for computer-assisted tracing. The system provides underlays from a large pool of drawings as a user begins drawing. Machine learning is used to recognize what the user seem to be drawing and relevant underlay drawing(s) are provided. This approach is innovative, but users may not learn much from tracing these underlays. It also provides no assistance with value or shading.

The differentiation of *DrawMyPhoto* is that it is designed to be more educational and is particularly focused on the order in which a drawing should be approached. It also focuses more on how to properly shade and add value to a drawing than previous systems. It draws inspiration from and builds on previous ITS and CST work, while introducing novel feedback mechanisms for pressure and tilt.

SYSTEM DESIGN

DrawMyPhoto is a progressive web app (PWA) built primarily in Javascript and utilizing Windows Ink API for rendering strokes based on applied pressure and tilt. The image processing component is built in Python using the OpenCV library. We will describe in more detail our **goals** for the system, as well as the nuances of its implementation.

G1—*Scaffolded*—Instructional scaffolding is a technique used in education to keep learners in their proximal zone of development [36]. This technique ensures the learner is adequately

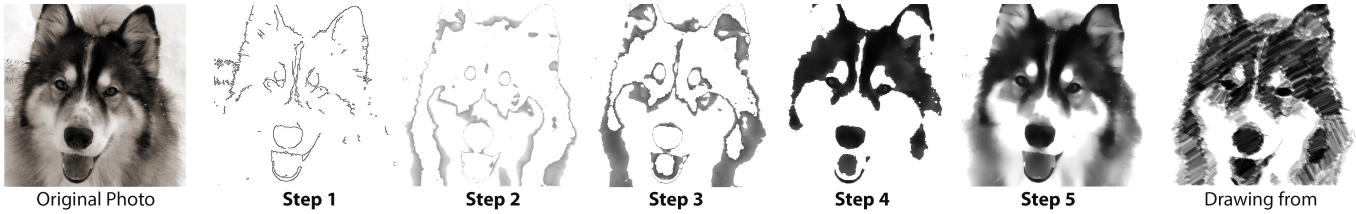

Figure 1. The DrawMyPhoto system generates an interactive drawing tutorial from a photograph by utilizing image processing. It generates 5 distinct steps and offers novel real-time feedback on pressure and tilt, allowing even novices to produce quality drawings as well as learn how stylus pressure and tilt affects shading.

challenged for where their skill level is, but not so challenged that they are prone to giving up. Placing training wheels on a bicycle while learning to ride one is a classic example of scaffolding. Forms of scaffolding for the system include the grid as well as the underlay feature (Figure 2). Through piloting the system we found that many novices were still uncomfortable with the grid approach, and so the underlay was introduced as a means to scaffold it further. When the user gains more confidence they can turn off the underlay feature or use it less.

G2—*Easy to use*—We wanted novices to be able to focus on their drawing without too many tools and power getting in the way. We accomplished this through a very minimal user interface. Differing stroke darkness and width is achieved through pressure and tilt of the stylus rather than complex menus and on-screen interactions. The tutorial itself is very linear, with five very clear steps.

G3—*Intelligent feedback in real-time*—We wanted the system to be able to detect performance of the user as they produce the line work and shading in a way such that real-time feedback can be provided that is similar to what a human could provide (Figure 3). This is accomplished through detecting pressure and tilt of the stylus and is described in more detail later in this section.

G4–*Preservation of personal style*—Among the goals of many previously designed systems is allowing the preservation of personal style [39, 5]. This system was no exception, as we wanted the users to produce more accurate drawings with better shading, but not in a manner that restricted their style.

G5—*Higher Quality Drawings*—Our overall goal was to allow novices to quickly produce a reasonably high-quality drawing independently of their experience and expertise in drawing. In doing so, we believed the system could provide a strong boost to their drawing self-efficacy.

Pedagogical Basis

We consulted with a domain expert in drawing as well as referenced the leading literature in art education [13, 32, 12, 9] while designing the system. When many professional artists draw in pencil, they begin with line work, ensuring proper proportion and alignment of the drawing, as well as composition. Artists lightly add strokes, building darker shades over time, generally beginning from light values and working towards darker values as they fill in the details of the drawing. This ensures they can recover from mistakes. Our goal was to process photographs and produce guidance images from them that mirror this process.

One exception is that we do not provide proportion guidelines, which many artists utilize when drawing the human face and human body particularly. This is in part because it is unnecessary with the underlay hint feature which we describe later in this section.

Image Processing

Figures 1 and 6 can be referenced for the five distinct steps the system automatically generates. For a photograph a user uploads, the system converts the image to grayscale and then increases the contrast and brightness by 50% which pushes white values forward. This is an important step, because without doing so the entire image may be gray values, meaning the user will shade the entire image. This often results in poor drawings and is not generally how trained artists draw (they will typically leave plenty of white negative space for areas not in shadow [13]). This initial pre-processing of the image results in the final "details" step which the user can reference to fill in any details they may have missed as they were shading in the individual steps. In order to produce the preceding line work and shading steps, we perform some additional operations on the image.

For the initial line work step we use Canny edge detection [6]. This suffices to produce a strong template for the line work of the drawing. The low threshold is the median of the single channel pixel intensities (a value between 0 and 255) multiplied by 0.66, while the upper threshold is the median of the single channel pixel intensities multiplied by 1.33. These thresholds allow for strong results on most photos, however some photos with poor lighting conditions or complex backgrounds can have unwanted noise.

For the shading steps we use k-means clustering [24] with k = 4 to find four clusters based on value. One of the clusters is inevitably the white values, which we discard. The remaining clusters are the light, medium, and dark gray value regions of the image. These correspond to the light, medium, and dark shading steps respectively. We can increase this number and produce more shading steps if needed, but we found three to be a reasonable number for a drawing tutorial for novices. While this technique works well for most photographs, naturally, some photographs with poor lighting, poor distinction between edges, etc. do not produce perfect results. Edge detection and segmentation are ongoing problems in computer vision and

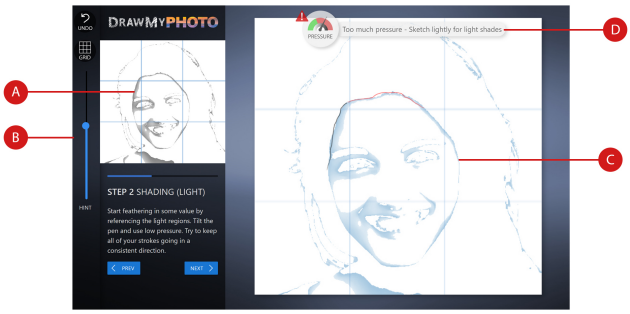

Figure 2. (A) is the sidebar with reference image for each step of the tutorial. It also includes some basic instructions for each step. (B) is a toolbar designed for bimanual interaction that allows user to adjust underlay feature, undo, as well as toggle the grid on and off by using the non-dominant hand. (C) is the drawing canvas itself which is always on the side of the dominant hand (in this case the right hand). Lastly (D) is where the real-time intelligent feedback appears to guide the user.

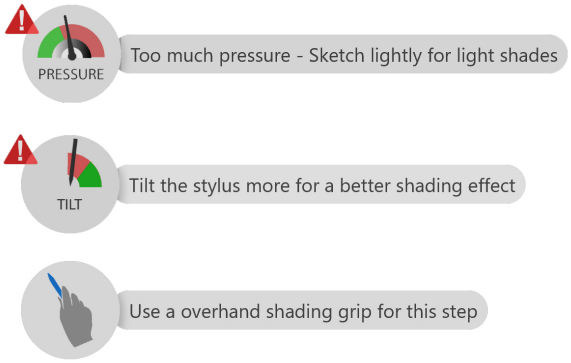

Figure 3. Examples of pressure and tilt feedback, along with grip suggestions. After piloting with users we found it was better to only show the feedback when necessary so that users would take notice. Constantly persistent feedback was often ignored. The "meters" and the associated ideal ranges in green were quickly grasped by most users. The values chosen allow for red-green colorblind people to still detect a difference.

some photographs may not have clear boundaries between foreground and background, or may have noisy backgrounds which influence the results. Likewise, some photos may be heavier in certain values, making certain shading steps either negligible or completely blank.

In terms of performance, the total computation time for a 400x400 pixel image which is adequate resolution for the tutorial is less than one second. Canny edge detection via OpenCV has already been shown to be very efficient with a computation time of 22 milliseconds for even a 2000x2000 pixel image [18]. For k-means clustering, we timed the computation for a variety of images and averaged the values on a consumer PC with an Intel i7-7660U 2.55 GHz CPU. It ranges from 236 milliseconds on average for a 400x400 pixel image to 1036 milliseconds on average for a 1000x1000 pixel image.

It is important to emphasize that our goal was not to advance the state-of-the-art in image processing or computer vision, but rather advance a novel application of existing algorithms and explore novel interactions and feedback associated with the educational goals of the system.

User Interface

The *DrawMyPhoto* user interface supports both right-handed and left-handed users by adjusting the canvas to be on the side of the dominant hand. A continuously visible reference image along with step instructions is always visible in a sidebar (Figure 2). The sidebar also includes a toolbar, giving the ability to undo strokes, as well as toggle a grid on and off and adjust a hint underlay. It was designed for bimanual input (e.g. "Pen + Thumb" work of Pfueffer et al. [30]), so that the user's dominant hand can be drawing while they adjust the grid or underlay with their non-dominant hand.

We chose to not give options for stroke width and color in an effort to keep the experience as simple as possible for novices. Stroke width and darkness can also change dynamically with pressure and tilt which is what is encouraged in this system.

Both the reference image and the canvas have a 3x3 grid. This is inspired by a popular grid-based method for drawing that has been used by artists for many years and was popularized by art educators like Betty Edwards [13]. It reduces cognitive load by allowing for the person drawing to only focus on each piece of the grid at a time. This can allow them to produce drawings that are more accurate and have proper proportions.

We also built a hint underlay feature which allows the users to trace for each step. We found from piloting the system that many novices were still intimidated by the grid approach. We decided to add this as a form of scaffolding, since our primary goal was to help novices produce high quality drawings and build up their confidence. The underlays mirror the reference images but are rendered in a light blue color so that the user can see a difference between where they have drawn and where they need to draw. "Non-photo blue" has been a popular underlay color in disciplines like architectural drafting and industrial design for many years.

Real-time Pressure and Tilt Feedback

Many modern devices and styluses detect both pressure and tilt, which not only allows for more realistic digital rendering, but can be used as raw input data towards novel interactions and feedback. The work of Xin et al. explored novel interactions using tilt [40, 41], however this data has never been used for help in a drawing application. This data is particularly useful in an educational drawing application because both pressure and tilt are important for various drawing techniques, particularly when using a versatile medium like pencil.

We designed a novel feedback system for pressure and tilt that gives appropriate feedback to the user depending on their performance and the step they are on (Figures 3 and 4). For example, in the initial line work step, low pressure is encouraged so that the line work is light and not overpowering the drawing. In the shading steps, varying levels of pressure are encouraged for the light, medium, and dark values. Additionally, low tilt is encouraged for a more realistic shading effect. The range of threshold values was determined through iterative testing and feedback from both users and domain experts. The ranges are generous and allow for enough flexibility so as to be useful to the user, but not constantly disruptive.

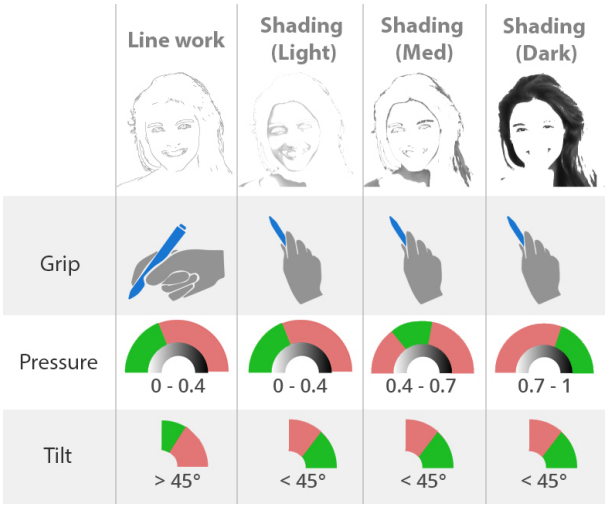

Figure 4. This shows how the feedback and ideal ranges change from step to step, and includes threshold values. Note that for Step 5 (Details) no feedback is given to the user because at that point they may use any range of pressure or tilt to finish the drawing, along with any grip they choose.

The pressure data is used raw and is a value from 0 to 1, with 1 being the highest detectable pressure. The "tilt" is actually an altitude value computed from tiltX and tiltY raw values, which are the angle of the stylus in relation to the screen in the X and Y planes respectively. This computed altitude value is the angle of the stylus in relation to the screen regardless of its orientation in the X and Y planes. We still label it "Tilt" for the user because it is a more easily understandable term.

Initially we had pressure and tilt feedback continuously visible, however through iterative testing and piloting we found that many users did not notice it as they focused on drawing. We discovered it was best for the feedback to appear only when the user was performing poorly and then to fade out. In this way, the user can notice the feedback, adjust accordingly, and then resume drawing. In addition to the dial and message, the stroke that was poorly executed turns red temporarily, provoking the user to be aware of which stroke and erase it if they choose to.

Additionally, as the user progresses steps, a feedback message depicting how to grip the stylus appears. This is important for the shading steps as it can be easier to tilt the stylus properly with an overhand grip, and this grip is often used by professional artists when they are shading [13, 12].

EVALUATION

We conducted a user study with novices to evaluate the efficacy of *DrawMyPhoto* in achieving its design goals, as well as to answer the following **research questions** related to those goals:

R1—*Can the system allow novices to produce high quality drawings with respect to accuracy, value (shading), and detail?*

R2—*What are novices able to learn from using the system?*

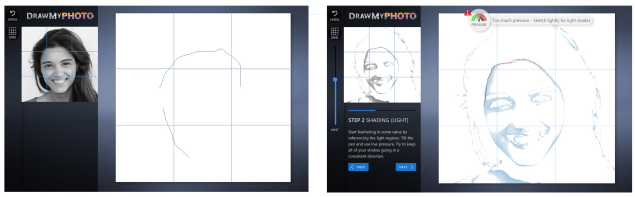

Control (Minimal Assistance) **Experiment** (Full Assistance)

Figure 5. The control group experienced a minimal interface with just the grid (barely visible in this image) and the unprocessed reference photo to help. The experimental group experienced the full interface with grid, generated steps, underlay feature, and real-time feedback on pressure, tilt, and grip. We set up the study this way to avoid learning and history effects.

R3—*Does the assistance provided by the system affect the number of strokes, average pressure / tilt of the user, or time spent drawing?*

Methodology

We chose to conduct a between-subjects study design in order to see the effects of the full *DrawMyPhoto* system in relation to a minimally-assisted drawing experience (Figure 5). We designed the study this way to compare the grid-based reference approach traditionally used in art education to this more modern technology-assisted approach. We hypothesized that the fully assistive system would result in much higher quality drawings, particularly with respect to accuracy and shading.

We recruited 20 participants of different ages and gender (18 to 59, average age 38; 14 males, 6 females) and assigned them to two equally sized quasi-random groups, ensuring they were demographically balanced for gender, handedness, and novice ability. Each group had at least one left-handed user and at least two female users. All participants were drawing novices. The average self-rated drawing ability in the control group and the *DrawMyPhoto* group was 1.7 ($\pm$ 0.64) out of 5, and 1.8 ($\pm$ 0.64) out of 5, respectively. 16 of the 20 participants expressed an interest in getting better at drawing, citing reasons ranging from personal enjoyment, e.g. *"peace of mind,"* to more practical reasons, e.g. *"It would be useful to convey ideas."*

All participants used the software on a Surface Pro device with the latest Surface Pen which has 4,096 points of pressure and tilt sensing from 90° (vertical) to approximately 26°. They were encouraged to use the device in a naturalistic manner, since the variable being tested was the assistive and educational experience, not how they use the device. Each participant answered some initial demographic questions, were given a brief tutorial of the interface, and were encouraged to draw the reference photo for up to 20 minutes. Subsequently, each participant answered some follow-up questions and a semi-structured interview was conducted to gain qualitative data on their experience.

We chose a portrait (see Figure 6) as the reference photo because many people want to learn to draw portraits [39], there is a lot of attention on portraits in the educational literature [13], and many previous research projects described earlier utilize portraits for their studies [39, 11]. This may allow for

easier comparison to previous systems in the future, although that was not the main focus of this research. The picture we chose has distinct light, medium, and dark values, allowing for distinct shading steps. Also, the portrait being a female with long hair allowed us to observe how participants approach adding the hair texture, which yielded interesting results.

Expert Ratings

We pursued expert ratings primarily to **answer R1**. We used an approach similar to Consensual Assessment Technique [1, 2], but more focused on ratings of quality versus creativity. We found two domain experts and qualified drawing instructors who have taught drawing and sketching for more than ten years at the university level. We had them rate the 20 drawings from the 20 participants blindly. They had no knowledge of which group the drawings belonged to, and could reference the original photo which all participants referenced when drawing. They rated each drawing on a 1 to 5 scale (1 being very poor, 5 being excellent) with respect to four different categories—overall quality, accuracy, shading (value), and details. As an additional set of data, they also ranked all 20 drawings in overall quality from 1 to 20.

Questionnaire and Interviews

In order to **answer R2**, the questionnaire and semi-structured interview questions were targeted at the quality and nuances of the learning experience, and the overall effect of the experience with regards to their drawing ability and confidence in drawing. We also encouraged participants to offer thoughts and suggestions for improvements to the system, particularly with respect to the real-time feedback.

Log Analysis

For **R3**, we hypothesized that the experimental group might produce more strokes and spend more time on the drawing because they feel less "stuck," as well as strokes of lower pressure on average because of the guidance to vary the pressure. Novices tend to give up on a drawing very quickly [13], and we expected to see this behavior less in the experimental group which had full assistance. To test this, we logged all stroke data for all participants, capturing information such as total number of strokes, stroke average pressure, stroke average tilt, and the total time spent to complete their drawings.

RESULTS

The drawings that resulted from the study can be seen in Figure 6. In this section we will discuss and analyze the results and how they relate to our design goals and research questions.

Expert Ratings

The expert blind ratings are shown in Figure 7. We conducted a Wilcoxon Signed-Rank Test between the ratings of each group for each category and found that the experts gave statistically significant higher ratings to the drawings in the experimental group in every category ($p < 0.001$ for overall quality, accuracy, shading, and details). This suggests design goal G5 (higher quality drawings) was reached, confirms our hypothesis of R1, and suggests the system was very effective at helping the novices produce high quality drawings.

While the higher accuracy can likely be attributed to the underlay feature that most novices in the experimental group used, the higher shading and details ratings are what we find most interesting. It can be readily seen from Figure 6 that many of the drawings in the experimental group have value that much more closely resembles the reference photo. This can especially be seen beneath the chin, where medium values were encouraged in step 3 of the tutorial, as well as in the hair, where dark values were encouraged in step 4 of the tutorial. Regarding details, one may notice much more attention to the nuances of the eyes and mouth in the experimental group drawings, while the control group tended to gloss over details or otherwise depict noses and mouths incorrectly.

It's worth noting one possible negative effect of the fully assisted experience is that some participants neglected detail in the textures of the hair, instead quickly shading it all in with dark value. Meanwhile in the control group, many participants attempted to replicate the texture of the hair, while neglecting the actual value of the hair, which is encouraged in step 4 of the tutorial they did not have access to. This suggests a potential failure to achieve the design goal G4 (preservation of personal style) because the system neglects to give feedback on depicting specific textures.

Questionnaire and Interviews

Participants rated their own drawings, the learning experience, and their own ability again as part of the post-study questions (See Figure 8). We conducted a Wilcoxon Signed-Rank test between the groups and found statistical significance for both ratings of their own drawings ($p < 0.02$), and for the learning experience ($p < 0.01$), but not self-rated ability ($p > 0.16$). We certainly did not expect to see an effect on self-rated ability after just one drawing, but we found it interesting that both groups rated their ability higher on average after the study.

Since the *DrawMyPhoto* group was exposed to the full educational and assistive experience, the qualitative data gained from this group was more rich, and is the primary focus of this section. A content analysis found that the most common comments participants made regarding the learning experience itself (R2) related to the proper order in which techniques should be used in drawing as well as shading techniques.

> *"I learned something about the order you draw lines, a better sense of shading"—P6*
>
> *"Maintaining certain pressures when shading, the order in which drawing techniques should be applied"—P7*
>
> *"The order of things, I really didn't know the process of drawing"—P9*
>
> *"The different levels of shading. How to look for different shades in a picture"—P11*
>
> *"To see shadow more than I did and address with varying pressure"—P13*

Many of the participants grasped the feedback quickly and expressed that they liked it suggesting G2 (easy to use) G3 (intelligent feedback in real-time) were reached, but some suggested it had diminishing returns, and offered ways to

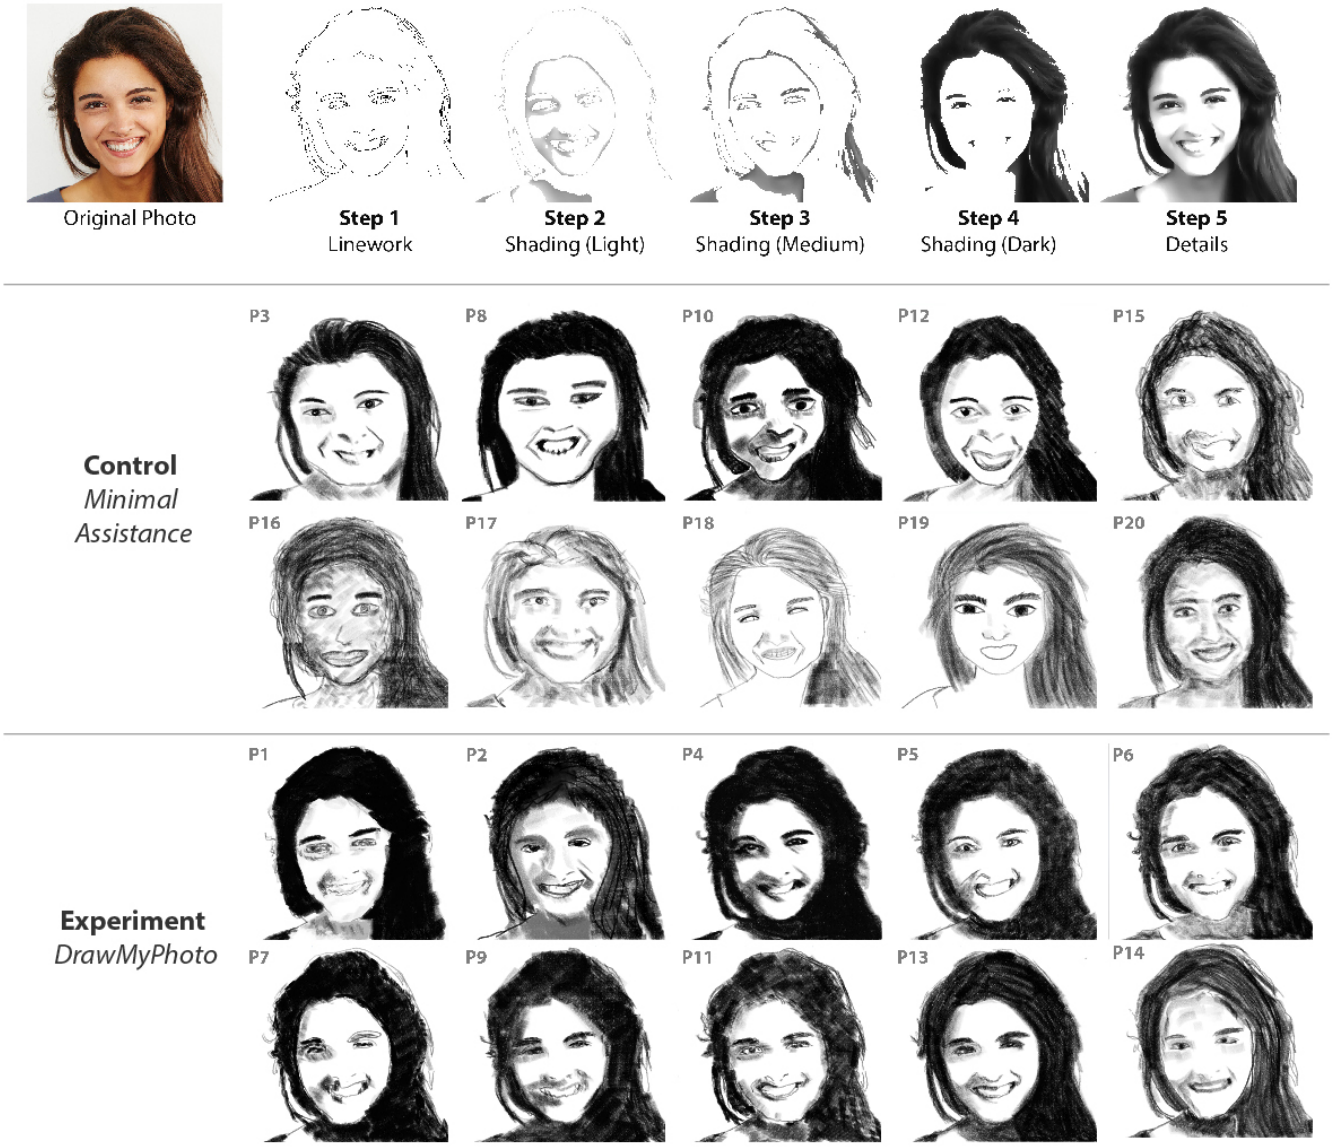

Figure 6. It can be readily seen that the assisted group which used the full application produced drawings that are not only more accurate, but have more consistent shading and attention to value, as well as attention to detail. This was confirmed by the expert ratings (See Figure 7). Note the common novice mistakes from the control group including improper proportioning, lack of attention to value, and drawing details that aren't there. However, also note that many of the novices put more effort in the texture of the hair

improve it such as combining the notifications with persistent visibility (instead of one or the other) or providing the feedback in a less frequent but more impactful manner.

"I really liked the tilt feedback. It was really fluid and cool"—P1

"I liked the highlighting in red, that makes it clear where the feedback is, what we're talking about"—P13

"Feedback was valuable at first but its value tapered and became less meaningful. Sometimes it was trying to nag me in a way that wasn't appropriate"—P13

Some of the participants surprised themselves, producing drawings that were beyond what they expected they were capable of. Most notably, P4 had a self-rated drawing ability of 1 out of 5, but produced the highest ranked drawing (average ranking of 1.5 $\pm$0.71) as well as one of the highest rated drawings (average overall quality 4.5 $\pm$0.71, average accuracy 4.5 $\pm$0.71, average shading 4.5 $\pm$0.71, and average details 4 $\pm$0.00). P13, who had a self-rated drawing ability of 2 out of 5, was also impressed with his drawing, noting that it helped with his confidence.

"I really did pretty well there, I'm pleased with the result of that. It certainly helped with the confidence. The

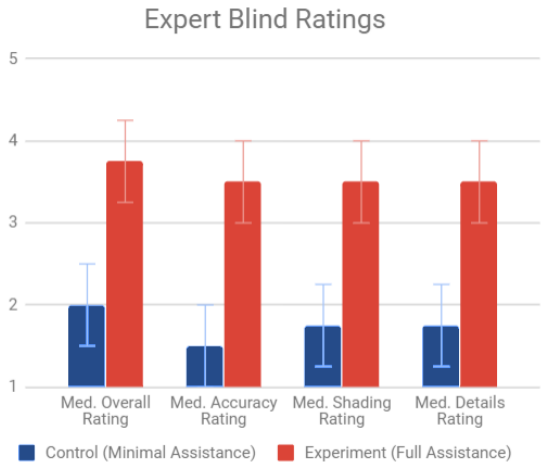

Figure 7. The experts gave statistically significant higher ratings to the experimental group in every category ($p < 0.003$ for overall quality, $p < 0.003$ for accuracy, $p < 0.005$ for shading, and $p < 0.006$ for details).

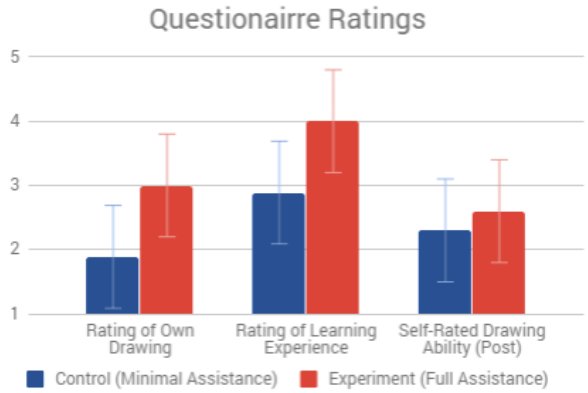

Figure 8. Participants in the fully assisted group rated their drawings higher, considered the learning experience better, and rated their drawing ability slightly higher than the control group.

different layers, the shading, the lines, breaking it down like that, that really helps immensely"—P13

Meanwhile, participants in the control condition were more likely to express how difficult it was to draw, even with the grid provided.

"Even with a grid drawing is hard"—P3

"Drawing is difficult when I get to shading and details"—P12

For this reason, we believe the tool can provide a strong boost in confidence and drawing self-efficacy [3] for users. This is important with respect to motivation to practice and continuing to improve drawing skills.

Log Analysis

There was large variance in the number of strokes participants drew in both groups, and very little variance between the groups (See Table 1). There was no statistical evidence that the fully assisted experience changed the average number of strokes, average pressure, average tilt, or time spent drawing ($p > 0.05$ in all cases).

As an example, P13 shaded with lots of short choppy strokes, resulting in 1245 total strokes while P7 produced more deliberate continuous strokes, resulting in only 216 total strokes. Both of those participants were in the *DrawMyPhoto* group and produced quality drawings. This is reflected in the very high standard deviations for Avg. Number of Strokes for both groups in Table 1. While the Avg. Pressure and Avg. Tilt are lower in the *DrawMyPhoto* group, as we hypothesized, their is marginal difference and it is not statistically significant. Surprisingly, the control group spent more time on their drawings on average, but again with a marginal difference.

This nullifies our hypothesis for R3, but we find it to be an interesting discovery. It suggests that the average number of strokes, average pressure and tilt, and drawing completion time by people is more related to personal style of drawing and

Table 1. Log Analysis

Measurement	Control	DrawMyPhoto
Avg. # of Strokes	520 ($\pm$ 294)	488 ($\pm$ 296)
Avg. Pressure	0.46 ($\pm$ 0.10)	0.41 ($\pm$ 0.09)
Avg. Tilt	45.87° ($\pm$ 25.89)	51.24° ($\pm$ 6.26)
Avg. Time Spent	16:20 ($\pm$ 4:15)	15:45 ($\pm$ 2:40)

approach to drawing. It is known that different people develop different, and sometimes very unique grips and approaches to handwriting and drawing [14]. It is not unlikely that they also develop very different ways of applying strokes in terms of speed, pressure, finesse, etc.

This also suggests that the application, at least in some capacity, allows users to preserve some of their own personal style, which was one of our design goals (G4). Granted, personal style is subjective and there are many more factors that contribute to it than what we were able to measure. As was discussed previously, the novices in the control group seemed to focus more on replicating textures than those in the *DrawMyPhoto* group.

Observations

All participants were able to complete the tutorial within 20 minutes with very little confusion, suggesting G2 (easy to use) was reached. While participants were encouraged to use the device in a naturalistic manner, most opted to keep the tablet at a low angle. In fact, only three participants (P2, P11, and P14) adjusted the angle of the tablet at any point during their drawing, all three of whom laid the tablet down flat (like a piece of paper). Only two of the participants laid the tablet down flat to begin with.

We also observed how often participants adjusted the hint underlay feature. 8 of the 10 participants adjusted the underlay at least once, generally to see if the drawing was turning out well. Two of the participants, P2 and P11, adjusted the underlay frequently (10 and 8 times respectively) as they checked their drawings for any details they had missed. None of the participants turned it off completely, suggesting it was a very helpful form of scaffolding (G1), but could perhaps be a crutch

in the long-term if they really want to improve their perception skills.

DISCUSSION AND FUTURE WORK

The following sections acknowledge the promising aspects of this work, as well as the limitations, and future directions are considered.

Lowering the Barrier to Artistic Creation

We believe the results of this study suggest *DrawMyPhoto* may offer a promising way to lower people's barrier to artistic creation. Shneiderman described the grand challenge of creativity support as enabling "more people to be more creative more of the time" [33]. The fact that experts rated the *DrawMyPhoto* drawings higher, and the participants in that group themselves rated their own drawings higher, suggests that the system has an advantage over the traditional grid-reference approach, particularly for novices. A system that can immediately allow a novice to produce a very high-quality drawing in just 15-20 minutes can go a long way in cultivating self-efficacy and motivation with respect to drawing.

We also saw that the system preserves at least certain aspects of personal style (G4) like average number of strokes drawn, and average pressure and tilt applied. However it is important to note that these factors are only a fraction of what constitutes personal style, and more continued creative practice is needed to really cultivate a distinct style. The system also constrains creativity to some extent, due to its scaffolded and step-by-step structure. That said, it does not micromanage users with stroke-by-stroke instructions, leaving some room for creativity and stylistic choices.

Improving the Educational Aspect

We believe the results showed preliminary evidence that we were successful in achieving our design goals, however there is always room for improvement.

We want to emphasize that an important aspect of learning to draw is improving *perception* [13, 12]. We believe it would be important to fade the scaffolding away in the application over time [23], such that users wouldn't continue to rely on tracing the underlay. While the feature helps improve the quality of their drawings, likely boosting confidence for many users, the grid will help them more to actually improve their perception in the long-term.

Many participants grasped the pressure and tilt feedback quickly and adjusted their strokes accordingly, but the feedback had diminishing returns and could become more of a nuisance once users understood how to shade properly. Our approach was one of many possible ways to provide feedback to influence user behavior and transfer knowledge in this domain. Other approaches could include summative feedback between steps, more positive and encouraging feedback, audio feedback, or different symbolic representations such as bars instead of meters.

Improving the Image Processing

We would like to reiterate that cutting-edge image processing and computer vision was not the main focus of this research, but with the overall user experience now defined, the techniques used could be improved. For example, while Canny edge detection sufficed to produce the line work step for most general photos, we could build on the work of Son et al. [35] and use similar techniques to produce line work that is much more human-like, with tapering lines of various thicknesses and reduction of noise. Granted, this might reduce preservation of personal style by influencing users to have a specific style. Additionally, the most state-of-the-art segmentation approaches (e.g. [7, 27]) could allow noisy backgrounds to be automatically parsed out of a photograph if a user is only interested in drawing the foreground.

It's worth noting that these more advanced techniques could result in the tutorial taking much longer to generate, therefore a balance between sophistication and computation time must be found.

Future Evaluations

We acknowledge that this was a pilot study and more rigorous evaluations could determine just how effective *DrawMyPhoto* is as an educational tool, as well as a tool for promoting drawing self-efficacy in novices.

An interesting avenue of future research would be to conduct a study involving participants own photographs. We avoided this initially due to privacy concerns and more focus on the general efficacy of the system, however using one's own photographs could produce even more of a sense of achievement and pride in one's drawings.

We would also like to determine if the learning experience provided by the system can transfer to traditional media like pencil and paper in which the same pressure and tilt adjustments are ideal for producing shading of different value and quality. This would confirm that the participants are indeed learning from the experience and can transfer that knowledge in to their own future drawings. Such an evaluation would require a more in-depth study with multiple drawings.

CONCLUSION

We've presented *DrawMyPhoto*, an interactive system that assists novices in learning to draw by automatically parsing any photograph in to a step-by-step drawing tutorial. We've discussed its novel features including real-time pressure and tilt feedback. We've shown evidence through our evaluation that the system was designed well, allowing novices to produce drawings of considerably higher quality than the traditional approach of utilizing a grid and referencing a photo. This was confirmed by expert ratings. Many participants who used the full system believed they had learned how to properly shade and approach shading by varying pressure and tilt. We largely achieved our design goals, and attest that the system is a promising approach to lowering the barrier of entry to drawing and artistic creation.

ACKNOWLEDGEMENTS

The authors would like to thank drawing instructors Stephen Chininis and Wayne Li from Georgia Tech who contributed their time and expertise to rate and rank the drawings.

REFERENCES

[1] Teresa M Amabile. 1982. Children's artistic creativity: Detrimental effects of competition in a field setting. *Personality and Social Psychology Bulletin* 8, 3 (1982), 573–578.

[2] Teresa M Amabile. 1996. *Creativity in context: Update to the social psychology of creativity*. Hachette UK.

[3] Albert Bandura. 1997. *Self-efficacy: The exercise of control* (1st ed.). Worth Publishers, New York, New York, USA.

[4] Russell W Belk. 2013. Extended self in a digital world. *Journal of Consumer Research* 40, 3 (2013), 477–500.

[5] Luca Benedetti, Holger Winnemöller, Massimiliano Corsini, and Roberto Scopigno. 2014. Painting with Bob: Assisted Creativity for Novices. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. ACM, New York, New York, USA, 419–428.

[6] John Canny. 1986. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence* 6 (1986), 679–698.

[7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2018. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2018), 834–848.

[8] Danielle Cummings, Francisco Vides, and Tracy Hammond. 2012. I Don't Believe My Eyes!: Geometric Sketch Recognition for a Computer Art Tutorial. In *Proceedings of the International Symposium on Sketch-Based Interfaces and Modeling (SBIM '12)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 97–106.

[9] Rudy De Reyna. 1996. *How to draw what you see*. Watson-Guptill.

[10] Daniel Dixon, Manoj Prasad, and Tracy Hammond. 2010a. iCanDraw: using sketch recognition and corrective feedback to assist a user in drawing human faces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 897–906.

[11] Daniel Dixon, Manoj Prasad, and Tracy Hammond. 2010b. iCanDraw: Using Sketch Recognition and Corrective Feedback to Assist a User in Drawing Human Faces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 10)*. ACM, New York, New York, USA, 897–906.

[12] Bert Dodson. 1985. *Keys to drawing*. North Light Ohio.

[13] Betty Edwards. 1999. *The New Drawing on the Right Side of the Brain* (2 ed.). Tarcher, New York, New York, USA.

[14] Sandra J. Edwards, Donna J. Buckland, and Jenna D. McCoy-Powlen. 2002. *Developmental & Functional Hand Grasps* (1 ed.). SLACK Incorporated, Thorofare, New Jersey, USA.

[15] Jennifer Fernquist, Tovi Grossman, and George Fitzmaurice. 2011. Sketch-Sketch Revolution: an engaging tutorial system for guided sketching and application learning. In *Proceedings of the* 24th *annual ACM symposium on User interface software and technology*. ACM, 373–382.

[16] Kathleen A Flannery and Malcolm W Watson. 1991. Perceived competence in drawing during the middle childhood years. *Visual arts research* (1991), 66–71.

[17] Tracy Hammond, Shalini Priya Ashok Kumar, Matthew Runyon, Josh Cherian, Blake Williford, Swarna Keshavabhotla, Stephanie Valentine, Wayne Li, and Julie Linsey. 2018. It's Not Just about Accuracy: Metrics That Matter When Modeling Expert Sketching Ability. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8, 3 (2018), 19.

[18] Batuhan Hangün and Önder Eyecioğlu. 2017. Performance Comparison Between OpenCV Built in CPU and GPU Functions on Image Processing Operations. *International Journal of Engineering Science and Application* 1, 2 (2017), 34–41.

[19] James W Hennessey, Han Liu, Holger Winnemöller, Mira Dontcheva, and Niloy J Mitra. 2016. How2Sketch: generating easy-to-follow tutorials for sketching 3D objects. *arXiv preprint arXiv:1607.07980* (2016).

[20] Richard Hickman and Sarah Lord. 2009. An Examination of Adolescents' Self-Efficacy, Engagement and Achievement in Representational Drawing. *Australian Art Education* 32, 2 (2009), 73–85.

[21] Ethan Hilton, Blake Williford, Wayne Li, Erin McTigue, Tracy Hammond, Julie Linsey, and others. 2016. Consistently evaluating sketching ability in engineering curriculum. In *DS86: Proceedings of The Fourth International Conference on Design Creativity, Georgia Institute of Technology, Atlanta, GA, USA*.

[22] Emmanuel Iarussi, Adrien Bousseau, and Theophanis Tsandilas. 2013. The Drawing Assistant: Automated Drawing Guidance and Feedback from Photographs. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, New York, USA, 183–192.

[23] Shari L Jackson, Joseph Krajcik, and Elliot Soloway. 1998. The design of guided learner-adaptable scaffolding in interactive learning environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press/Addison-Wesley Publishing Co., 187–194.

[24] Tapas Kanungo, David M Mount, Nathan S Netanyahu, Christine D Piatko, Ruth Silverman, and Angela Y Wu. 2002. An efficient k-means clustering algorithm:

Analysis and implementation. *IEEE Transactions on Pattern Analysis & Machine Intelligence* 7 (2002), 881–892.

[25] Swarna Keshavabhotla, Blake Williford, Ethan Hilton, Paul Taele, Wayne Li, Julie Linsey, and Tracy Hammond. 2017. Conquering the Cube: Learning to Sketch Primitives in Perspective with an Intelligent Tutoring System. In *SBIM '17 Proceedings of the Symposium on Sketch-Based Interfaces and Modeling*, Vol. 2. ACM, Los Angeles, CA, USA, 1–11. ISBN: 978-1-4503-5079-2, DOI: 10.1145/3092907.3092911.

[26] Yong Jae Lee, C Lawrence Zitnick, and Michael F Cohen. 2011. Shadowdraw: real-time user guidance for freehand drawing. In *ACM Transactions on Graphics (TOG)*, Vol. 30. ACM, 27.

[27] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.

[28] Cewu Lu, Li Xu, and Jiaya Jia. 2012. Combining sketch and tone for pencil drawing production. In *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*. Eurographics Association, 65–73.

[29] Michael I Norton, Daniel Mochon, and Dan Ariely. 2012. The IKEA effect: When labor leads to love. *Journal of consumer psychology* 22, 3 (2012), 453–460.

[30] Ken Pfeuffer, Ken Hinckley, Michel Pahud, and Bill Buxton. 2017. Thumb+ Pen Interaction on Tablets.. In *CHI*. 3254–3266.

[31] R Richards. 2003. My drawing sucks!: Children's belief in themselves as artists. In *NZARE/AARE Conference*, Vol. 29.

[32] Scott Robertson and Thomas Bertling. 2013. *How to Draw: Drawing and Sketching Objects and Environments from your Imagination* (1^{st} ed.). Design Studio Press, Culver City, California, USA.

[33] Ben Shneiderman. 2002. Creativity support tools. *Commun. ACM* 45, 10 (2002), 116–120.

[34] Derek Sleeman and John Seely Brown. 1982. *Intelligent tutoring systems*. London: Academic Press.

[35] Minjung Son, Henry Kang, Yunjin Lee, and Seungyong Lee. 2007. Abstract line drawings from 2D images. In *Computer Graphics and Applications, 2007. PG'07. 15th Pacific Conference on*. IEEE, 333–342.

[36] Lev Vygotsky. 1987. Zone of proximal development. *Mind in society: The development of higher psychological processes* 5291 (1987), 157 pages.

[37] Blake Williford. 2017. SketchTivity: Improving Creativity by Learning Sketching with an Intelligent Tutoring System. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 477–483.

[38] Nelligan Li Linsey Hammond Williford, Taele. 2016. Persketchtivity: an intelligent pen-based educational application for design sketching instruction. (2016), 115–127.

[39] Jun Xie, Aaron Hertzmann, Wilmot Li, and Holger Winnemöller. 2014. PortraitSketch: Face Sketching Assistance for Novices. In *Proceedings of the* 27^{th} *Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. ACM, New York, New York, USA, 407–417.

[40] Yizhong Xin, Xiaojun Bi, and Xiangshi Ren. 2011. Acquiring and pointing: an empirical study of pen-tilt-based interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 849–858.

[41] Yizhong Xin, Xiaojun Bi, and Xiangshi Ren. 2012. Natural use profiles for the pen: an empirical exploration of pressure, tilt, and azimuth. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 801–804.

[42] Mingtian Zhao and Song-Chun Zhu. 2013. Artistic rendering of portraits. In *Image and Video-Based Artistic Stylisation*. Springer, 237–253.