EXPLORING METHODS FOR HOLISTICALLY IMPROVING DRAWING ABILITY WITH ARTIFICIAL INTELLIGENCE

A Dissertation

by

BLAKE WILLIFORD

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Tracy Hammond
Committee Members,	Andruid Kerne
	Frank Shipman
	Jeffrey Liew
Head of Department,	Scott Schaefer

May 2020

Major Subject: Computer Science

ABSTRACT

Drawing is a highly useful skill that can make people better at solving problems, communicating ideas to others, collaborating, and producing more creative and novel ideas. It can be a difficult skill to master for many people, however. Like any learned skill, it requires many hours of practice for noticeable improvement, and sufficient motivation is also necessary to keep practicing consistently over a period of time.

Utilizing sketch recognition and other forms of artificial intelligence to assist in learning to draw may facilitate the necessary improvements in self-efficacy and motivation students need to improve their drawing ability. While similar tools have been explored, there has been little to no effort at designing a truly holistic approach for teaching drawing skills that includes the basic fundamentals and building blocks for drawing any 3-dimensional object.

This dissertation explored the potential of an intelligent tutoring system for teaching drawing skills called *SketchTivity* along with various other technology probes focused on drawing. We found evidence that individuals could build confidence, build motivation, make measurable improvements to drawing ability, and reduce fixation when ideating concepts through the various studies we conducted. We developed a flexible perspective accuracy recognition algorithm that can help individuals learn perspective. In interviews with students and teachers who used *SketchTivity* we discovered benefits and limitations of the system. Students were engaged by the interactive lessons, motivated by the gameplay, and saw it as a great warm-up tool. Meanwhile instructors loved that the system could offload grading tasks for them.

We hope the nuances of this potential will inform the future development and promise of the approaches described in this dissertation along with similar approaches to impact education at large.

DEDICATION

To everyone who has ever said "I can’t draw".

You can, and you will.

ACKNOWLEDGMENTS

First and foremost, a big thank you to Matthew Runyon, a close collaborator of mine on much of this work. I don't think I would have made it through this Ph.D. without you Matt. I hope to keep working with you for a long time. Let's change the world buddy.

Thank you to all the great teachers and mentors who have made me who I am today. I have received 22 years of public education. To name a few, Wayne Li, John Lau, Michelle Berryman, and Dr. Julie Linsey from Georgia Tech, my advisor Dr. Tracy Hammond and committee members Dr. Andruid Kerne, Dr. Jeffrey Liew, and Dr. Frank Shipman from Texas A&M. Thanks for the opportunities and the support during this long journey. You all made me believe in the impossible.

Thanks to all my lab mates in the Sketch Recognition Lab for your support including Josh Cherian, Paul Taele, Seth Polsley, Raniero Lara-Garduno, Jung In Koh, Samantha Ray, Duc Huang, and Xien Thomas. You have been like a family to me.

Lastly, thanks to my real family for your support on this mission. I'm very blessed and grateful to be the first Ph.D. graduate in my family.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Dr. Tracy Hammond, Dr. Andruid Kerne, and Dr. Frank Shipman of the Computer Science and Engineering (CSCE) Department, along with Dr. Jeffrey Liew of the Education Psychology (EPSY) Department.

All work described in this dissertation was led by Blake Williford unless otherwise noted.

Funding Sources

The work was funded in part by the National Science Foundation under Grant Numbers 1441331 and 1441291. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

TABLE Page

1. INTRODUCTION

Drawing is a valuable skill for creativity, communication, collaboration, and problem-solving. It is a practical skill that can benefit internal thought processes [1] while also reducing cognitive load and facilitating perceptually-based reasoning [2]. It can elevate awareness of and attention to details of surroundings [3], hone analytical skills [4], and stimulate both halves of the brain [5].

In academic settings, drawing can improve students' general academic achievement and problem-solving thinking [6], improve peripheral skills in writing, critical thinking, and brainstorming, boost self-confidence in other academics from successful artistic pursuits [7] and even improve three-dimensional spatial recognition skills [8].

In professional settings, drawing is widely used to explore and refine concepts and communicate them to others. Painters regularly use drawing as a means to explore compositions and underlays before committing paint [9]. In disciplines like architecture and industrial design, freehand sketching is considered essential in exploring design ideas in the early stages of the design process [10, 11, 12]. Even in modern HCI disciplines such as user experience design and software engineering, sketching is increasingly valued for many of the same reasons [13, 14, 15, 16].

Drawing is a difficult skill to master, however. The studio setting is the dominant approach for drawing instruction in settings such as high schools or universities. In this setting students receive direct feedback on their progress from an experienced instructor and fellow classmates. However, valuable expert feedback from instructors is less available as classroom size increases and difficult to access beyond classroom and instructor office hours [17]. An additional concern is that for studio environments that involve students sharing their work to the rest of the class (which is very common), those students who have low self-efficacy [18] are less motivated to participate and complete tasks given when they are learning such as in the classroom [19]. Learning to draw in general is significantly affected by self-efficacy such that students with low-efficacy are less confident to improve, unmotivated to practice, and discouraged by more skilled peers with their own drawings [20].

This problem affects people of all ages and may be rooted in childhood as many children seem to decline in spontaneous art from around age 7 [21] and this has been found to be caused by a decrease in self-efficacy as it relates to sketching and artistic pursuits in general [22].

Intelligent Tutoring Systems (ITS) hold much promise for freeing the bandwidth of instructors, allowing for more personalized feedback and support as well as building more self-efficacy and motivation in learners. Students are increasingly seeing more intelligent tutoring systems that can provide instant feedback, but most systems only tell the student if the student's answer is right and wrong. Providing personalized feedback on the process is incredibly valuable, but takes significant instructor resources [23]. An additional growing concern among educators is that students are losing both the critical skill of sketched diagrams and the ability to think through problems via sketching [24].

We have built such an ITS for drawing fundamentals called *Sketchtivity*. We have also built a series of technology probes centered around learning drawing which are the basis of various studies in this dissertation.

By utilizing sketch recognition, game-based learning, data visualization, and real-time feedback it is possible to guide learners in drawing fundamentals and facilitate more confidence in their drawing ability. This improved self-efficacy can in theory lead to more motivation, making individuals more motivated to practice. Improvements in both self-efficacy and motivation form additional feedback loops with drawing ability itself. Over time, drawing ability could begin to effect creativity and creative self-efficacy, as the more capable and confident learner begins to use drawing as an applied skill for ideation.

The model in Figure 1.1 forms the basis of the research questions surrounding this dissertation and what is explored in it. The various studies described in this study are all centered around helping individuals *holistically* improve their drawing ability. In order to do so we focus on methods for building self-efficacy, motivation, and creativity, in addition to ability.

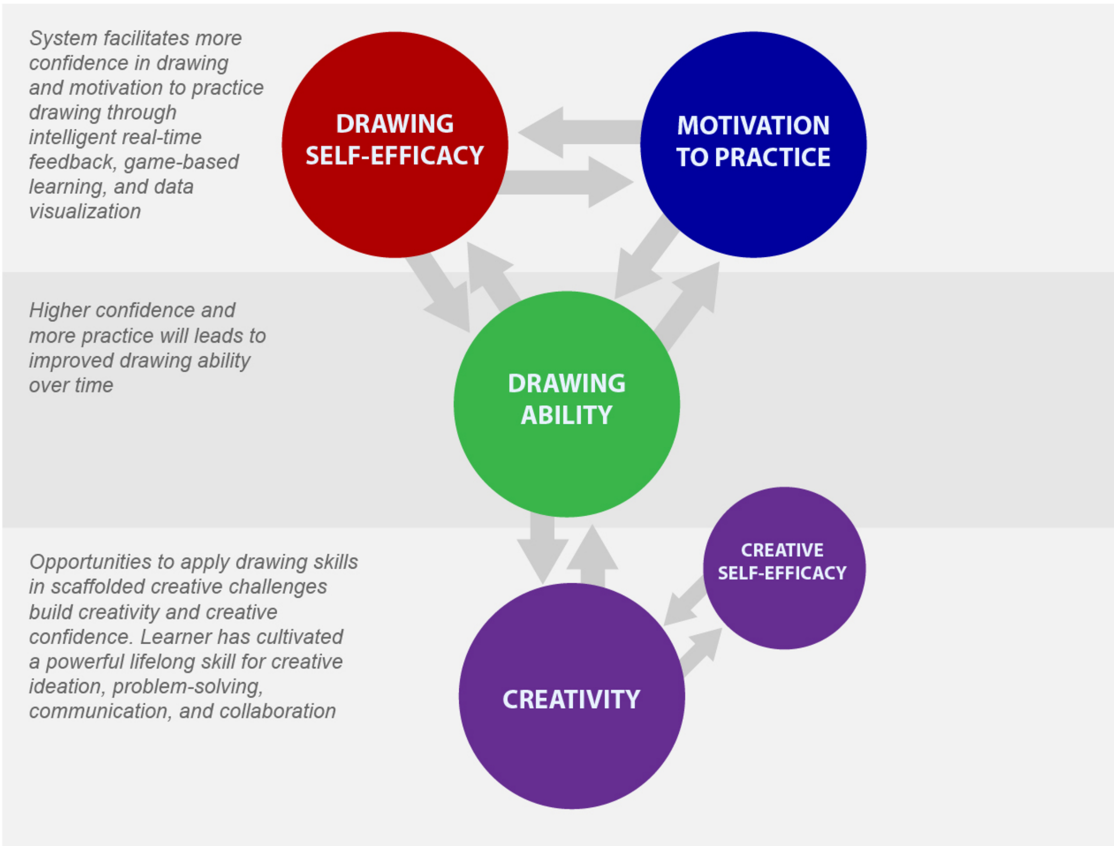

IMPROVEMENT	SELF-EFFICACY	MOTIVATION	DRAWING ABILITY	CREATIVITY
METHOD USED IN ITS	Real-time feedback Game-based learning Visualizing performance	Real-time feedback Game-based learning Visualizing performance	Increased practice (due to self-efficacy and motivation) Real-time feedback	Increased ability Increased creative self-efficacy Positive feedback
HOW IT CAN BE MEASURED	Drawing Self-Efficacy Scale Semi-structured interviews	System activity Semi-structured interviews	Extracting sketch data Domain expert panel Semi-structured interviews	Creative Self-Efficacy Scale Idea fludity / generation Domain expert panel

Figure 1.1: A high-level model of how the *Sketchtivity* system should facilitate improvements in drawing ability and creativity. This dissertation explores this model and its efficacy.

2. INTELLECTUAL MERIT

Raising global literacy has been a long and arduous multinational effort that has resulted in 86.3% of individuals aged 15 or higher in the world now being able to read and write as of 2019 [25]. In developed nations such as the United States it is closer to 99%. However, there has been little to no effort to raise global drawing ability, despite the many benefits drawing has in terms of communication, collaboration, and problem solving. As a picture can "speak 1000 words," drawing holds many benefits that when combined with literacy can make individuals far more prepared for the highly creative and collaborative 21st century.

Unfortunately, the teaching of visualization skills, in the form of sketching and drawing instruction, was dropped from most US engineering curricula in the 1980s and 1990s [26, 27] and replaced with CAD. Even in professions like architecture, sketching has become a "lost art" and is no longer valued as much as it has been in the past [28] even if there is some light focus on the discipline.

This may have unintended consequences regarding problem solving ability and possibly creativity, which can have career-long impacts on engineers, designers, and society at large. We hypothesize that reintroducing free hand sketching, along with more innovative teaching methods, can improve students' problem solving skills, creativity, and their self-efficacy towards tasks that involve visualization and the communication of ideas. Below are examples of the potential impact of this research:

- This dissertation will impact (and is already impacting) hundreds of high school and university students as well as instructors who will participate in studies and drive the technology and research forward.
- A comprehensive set of lessons and exercises that teach the fundamentals of drawing and visual communication will be built, in addition to creative challenges that encourage transfer of knowledge. Users of the system will gain mastery of valuable skills that will benefit their

careers as industry becomes increasingly collaborative and creative.

- The web-based educational platform is not device-dependent and can therefore be used on a wide variety of devices including Cintiqs, iPads, Android tablets, Surface Pros, etc. allowing for widespread accessibility independent of socio-economic status or access to quality education.

- It will allow and encourage students to use their creativity, which may engage more reluctant learners, making STEM content more engaging to a more diverse group of students.

- Sketch recognition algorithms created during this dissertation may be useful for other researchers and may be used in related domains like creativity support.

- If the technology proves itself to be effective in improving learning outcomes and motivating students, it may scale up to significantly impact primary and secondary education and society at large.

3. PRIOR WORK

It's important to consider related work in other intelligent tutoring systems, sketch-based intelligent tutoring systems, sketch-based tutorial-generation tools, and sketch-based games. We learn lessons from these works and take note of where they are successful and not successful.

3.1 Intelligent Tutoring Systems

Intelligent Tutoring Systems (ITS) [29] are educational systems designed to replicate the one-on-one tutoring experience. While they are a promising approach to delivering education efficiently and even remotely, they can be difficult to evaluate [30].

Nevertheless, they have been shown to be effective and in certain cases even more effective than human instructors [31, 32]. A meta-analysis conducted by Kulik and Fletcher found that students who had access to ITSs outperformed students from conventional classes in 46 out of 50 controlled evaluations [33]. Some systems have become successful commercial products, such as *Duolingo*. One study found that spending an average of 34 hours using *Duolingo* could yield the same learning outcomes as an entire semester of Spanish at the university level [34].

Design recommendations for ITS from Sottilare et al. emphasize the importance of the "learner model," or the cognitive and motivational states of the individual that can be modeled and utilized to guide automated tutoring. [35].

3.2 Sketch Recognition

Sketch recognition is the automated recognition of hand drawn diagrams by a computer [36], and can be considered a subset of computer vision. A lot of early work in sketch recognition was focused on automatic recognition of UML diagrams [37, 38]. Follow-up work has focused on immediately turning hand drawn diagrams in to executable code [39, 40].

Some systems have been designed to be domain-independent, allowing for sketch recognition to be used in any kind of educational context [41, 42, 43, 44]. There has also been frameworks for robust recognition and beautification of basic shapes like circles, squares, stars, lines, spirals, etc.

[45, 46, 47, 48], which can be used in a variety of different domains and contexts.

3.3 Sketch-based Intelligent Tutoring Systems

A sketch-based ITS utilizes a sketch-based user interface for users to provide hand-drawn diagrams or symbols, and can provide real-time feedback on the drawings by utilizing sketch recognition. Representative examples of sketch-based ITS applied to educational disciplines outside of drawing instruction include music [49], east asian languages [50, 51], statics problems in mechanical engineering [52, 53, 54, 55, 56], geography [57], and mathematics [58].

Naturally, there have been explorations focused on drawing and sketching instruction, with the majority having a focus on fine art styles of drawing. *iCanDraw* [59] and *EyeCanDraw* [60] provide step-by-step guidance for drawing human faces and eyes, respectively. They utilize sketch recognition to provide real-time feedback on accuracy, helping users correct drawings as they progress. Since these domains are very specific, it is questionable how much general drawing skill users can gain, and whether or not improvements can transfer to other domains.

CogSketch [61, 62, 63, 64, 65, 66, 67, 68] takes a more cognitive approach and is a platform for both understanding the mental model behind sketches [69] as well as how to improve one's problem-solving ability through real-time feedback [70].

Aalto Drawing is a recent educational tool that shows promise for "fair and encouraging" automatic accuracy feedback on a variety of basic forms [71]. While it is limited to a star rating, and has yet to be formally evaluated, it hints at the possibility of the gamification of drawing practice and its potential to motivate students.

Both Keshavabhotla et al. [72] and Kuribayashi et al. [73] developed systems for recognizing the accuracy of cubes in perspective (an important skill in design sketching).

3.4 Sketch-based Tutorial Generation Tools

Some tools have focused on the automatic generation of tutorials. Tutorials, also commonly called demos, are one of the main methods for conveying how to draw. In studios, the instructor typically devotes time in each class to a tutorial, during which the students can follow along and replicate the drawing. Fernquist et al. explored how the actions of an artist could be recorded in sketching software and then be converted in to full step-by-step tutorials with stroke-by-stroke feedback and assistance for users[74]. Hennessey et al. explored converting 3D models of products into step-by-step sketching tutorials [75]. However, the tutorials are not interactive, and provide no feedback to assist the learning experience.

3.5 Digital Drawing Tools

Numerous tools have been developed which assist users in sketching in 3-dimensions. The work of Schmidt et al. [76, 77] explores powerful interaction techniques for sketching in 3D seamlessly, but the system itself is not mean to educate its users on how to sketch in 3D. Likewise, *EverybodyLovesSketch* [78] is a novel system for sketching in 3D that has been shown to be useful for novices but is not designed to be strictly educational.

ShadowDraw is a unique tool that provides an underlay for users as they draw, assisting them with drawing recognized objects like bicycles [79]. Its educational benefit is questionable, however, since the tool does not provide any instruction or convey any principles of how to draw.

3.6 Creativity Support Tools

Creativity Support tools (CST) bridge the gap between novice and expert ability in drawing and painting by providing digitally-mediated assistance, so they are also worth considering as relevant prior work. Some incorporate a wider set of figure drawing options (i.e., The Drawing Assistant [80]), while others introduce more specialized features for assisting novices to improve their drawings (i.e., PortraitSketch [81]), or provide greater creative expression from reference models (i.e., Painting with Bob [82]). These systems for novice users are primarily assistive, rather than educational, focusing more on improving drawings of specific objects or images in fine art styles of drawing as opposed to conveying fundamental concepts of design sketching, however they do provide valuable insights into successful sketch-based user interfaces and interaction design practices which can help novices create.

3.7 Co-Creative Systems

Computational co-creative systems are systems that facilitate the co-creation of artifacts between users and computers [83, 84, 85]. Also referred to as Mixed-Initiative Creative Interfaces (MICIs) [86], such systems have been considered in a variety of domains such as public art installations [87], music improvisation [88], poetry writing [89], dancing [90], etc.

A relevant example of such a system would be the *Drawing Apprentice* from Davis et al. [91, 92]. This system responds to user strokes with strokes of its own, resulting in abstract co-created drawings. We draw inspiration from this work, however we focus more on formalized concept sketching rather than making abstract art. In other words, we focus on a more practical use case of sketching specific concepts or ideas.

Another relevant example in this domain is *ShadowDraw* from Lee et al. [79]. This system responds to user strokes by recognizing what the user is trying to draw and providing underlays of that object to assist the user. For example, a user can begin drawing a bicycle, and *ShadowDraw* will generate underlays of bicycle drawings which the user can then trace, helping the user to converge on a stronger bicycle drawing. We draw inspiration from the underlay aspect of this work in particular, however our approach is oriented more towards divergent thinking rather than convergent thinking.

3.8 Sketch-based Games

Sketch-based games are by no means a new concept. Casual classics like *Tic-Tac-Toe* and *Hangman* can be played with readily accessible pen and paper. *Pictionary* [93], a charades-inspired word guessing game that involves drawing, remains a popular board game decades after its creation.

With digital games now a prevailing medium, new sketch-based games have been created which utilize either finger swiping or the use of a digital stylus. *DrawSomething* [94], a mobile game that is very similiar to *Pictionary* is worth noting for how it promotes communication skill through sketching.

Other sketch-based games have been more oriented towards education. Many have been geared towards children, e.g., the work of Paulson et al. in encouraging learning through more tactile and kinesthetic approaches [95]. *EasySketch2* was developed to assess and improve fine motor skills ion children in order to see how prepared they are for school [96]. *TAYouKi* [97, 98] was oriented towards teaching young children to draw basic shapes. A system developed by Truong et al. [99] helps children learn new concepts through sketch-based interaction. Wasson et al. [100] showed games can be utilized to learn art history and other general pictures.

Some games have been more oriented towards research and collecting data. *PicturePhone* and *StellaSketch* are collaborative sketch-based games that utilize human computation [101] to collect and classify sketch data [102]. Google released a game called *QuickDraw* that could guess what users were drawing using machine learning, and was meant to build a massive dataset of over 50 million doodles. These games offer fun play experiences for users while also fulfilling secondary research motives.

3.9 Differentiation

Our systems are designed to be more general, focusing on design sketching fundamentals, and incorporating gameplay and gamification to motivate students. The systems are general in the sense that it conveys sketching fundamentals that can be used in any domain. The goal for many disciplines like industrial design, architecture, and engineering is often to be able to sketch from one's imagination, and produce interesting and novel forms. We utilize both gamification [103, 104] and "serious" games [105] in an effort to motivate students, since our earlier research showed how low motivation to practice sketching was a major concern for students.

In terms of how Van Lehn defines intelligent tutoring systems [106], the inner loop of *Sketchtivity* gives real-time feedback on accuracy for each and every basic primitive sketch in the form of red deviation lines, while the outer loop provides summative feedback after each task. Based on performance and the student model, the outer loop allows access to more advanced lessons which the student can then choose from.

4. SENSITIZING CONCEPTS

4.1 Educational Psychology

The design of the systems described in this dissertation is heavily influenced by concepts in educational psychology. Decades of research in how humans learn has provided a wealth of information about the optimal ways to cultivate a healthy psychology around learning.

4.1.1 Self-Efficacy

Self-efficacy can be defined as a self-assessed judgment of “how well one can execute courses of action required to deal with prospective situations” [18]. Self-efficacy in various subcategories such as creative self-efficacy can be studied, and Albert Bandura, the psychologist behind the term, wrote a guide for developing self-efficacy scales in 2006 [107].

Self-efficacy is closely tied to confidence and is an important aspect of an individual being self-motivated to learn. When an individual gains self-efficacy they are much more likely to have sustained motivation to keep practicing and improving a skill [108, 109].

4.1.2 Self-Regulated Learning

This work is also heavily influenced by the concept of self-regulated learning [110] which is related to self-efficacy. It describes an ideal form of learning in which a learner is self-motivated to learn, regularly reflects on what they’ve learned (*metacognition*) [111], and is actively engaged in taking control of what they are learning.

Educational technology is a promising way to promote self-regulated learning as it can make the learning more process more transparent [112, 113, 114, 115]. An intelligent tutoring system that tracks progress on learning various concepts can provide that information to the user, helping them to self-regulate their learning. That can come in the form of real-time feedback on what to improve upon next or through “profiles” with data visualization on performance over time.

Studies have shown that when students feel that their schoolwork is meaningful and interesting, and the classroom environment is supportive, they are more likely to engage in self-regulated learning [116].

4.1.3 Self Determination Theory

Since this work is focused on motivating individuals, it's worth mentioning some sensitizing concepts in motivation theory that are relevant. Self-Determination Theory (SDT) [117] is one of the prevailing paradigms in psychology, as it describes people's innate psychological needs for competence, autonomy, and relatedness. When learning a new skill, individuals tend to be more *extrinsically* motivated at first, looking for outside rewards such as positive feedback or monetary gain. However through the natural process of self-determination, motivations are internalized and the individual starts to become more *intrinsically* motivated to learn the skill.

4.1.4 Mastery Learning

Another concept from education theory that is influential in the design of *Sketchtivity* and the various other technology probes in this dissertation is mastery learning [118]. First proposed by Benjamin Bloom in 1968, it suggests that the optimal way to learn is to always master concepts in a step-by-step manner, only moving on to more advanced concepts when the prerequisite concepts have been mastered. By always staying in the "proximal of development" [119], individuals can master concepts at their own pace, ensuring they truly understand the material as they progress.

This contrasts with much of how education is delivered in public education in the US and abroad. Typically, courses must cover a certain set of material and students are dragged along independently of their mastery of the concepts. This tends to create a stratification between students who quickly understand the concepts and those who struggle or take more time to understand them. There is no exception in studio environments, as many students learning to draw struggle to master fundamentals, but are nevertheless forced to move on to more advanced topics to cover the material of the course.

4.2 Gamification Research

In recent years, research around games and gamification has been on the rise. Games hold much promise for engaging and motivating people, and can be utilized for far more than entertainment.

Games can be utilized for education, relaxation, creativity, productivity, etc. and researchers and technologists are just scratching the surface in terms of the potential.

4.2.1 Motivating with Games

Motivation in games has been widely studied and remains an interesting area of research since games have the potential to be very immersive and engaging. Richard Bartle [120] created a taxonomy of player types for massively-multiplayer online (MMO) games that describes their motivations. We draw inspiration from this 2-axis approach to motivations.

The work of Phillips et al. [121] explored how different in-game rewards motivate people. They found little difference between individual isolated rewards, but found that a greater *variety* of rewards has a positive impact on interest and enjoyment. This is an important insight in to the design of serious games to motivate students in educational contexts.

4.2.2 Game-based Learning

Serious games have been shown to be more effective than traditional educational methods in terms of learning and retention, [105] but not necessarily motivation. Ongoing research in game-based learning is exploring the degree to which games can motivate students more [103, 122, 104].

One approach is to utilize “stealth learning” [123, 124]. Stealth learning disguises the learning within the gameplay and game mechanics. Caution is advised when using stealth learning, however, as

4.3 Drawing Pedagogy

Drawing is a broad skill with many different approaches and styles, as well as many professional disciplines centered around those various approaches. These approaches include design sketching, representational (fine art) drawing, cartooning, drafting, etc. We touch on two approaches that heavily influence our various prototypes—design sketching and representational drawing.

4.3.1 Design Sketching Pedagogy

Design sketching is a practical form of drawing that is typically taught to industrial designers, architects, and in some cases engineers. In design sketching, a mastery learning [118] approach is utilized in which students practice basic two-dimensional and three-dimensional "primitives" before moving on to combining primitives in to actual objects and forms in perspective. One advantage of this approach is it teaches the individual how to draw in a *general* sense, eventually giving them the ability to sketch from their imagination [125, 126].

Design sketching emphasizes precision, smooth fluid line work, and speed [127]. While this is a certain style of drawing focused on rapid idea generation, it doesn't have to be a rigid limitation and individuals can still utilize more traditional fine art drawing skills in their repertoire.

4.3.2 Representational Drawing Pedagogy

Representational, or fine art drawing, is the form of drawing most people are familiar with. This is the form of drawing that is typically taught in art courses in K-12 education as well as in more specialized art programs. This form of drawing emphasizes realism and capturing details of the real world. Students will typically practice portraits, still life, landscapes, etc. and be encouraged to have accurate line work, shading, and detail.

Betty Edwards is a famous art teacher best known for her book *Drawing From the Right Side of the Brain* [128, 129]. In her workshops she emphasized techniques for helping everyday people learn to draw with particular emphasis on *perception*. Indeed, drawing as a skill is closely tied to perception, and the better one is able to perceive details of the world, the better they are at replicating those details in a drawing.

Numerous other books have focused on guiding beginners in learning fundamental drawing skills [130, 131, 132] including skills like perspective [133], which beginners tend to struggle the most with. Perspective is the focus of one of the research questions and studies in this dissertation (see Chapter 8).

This form of drawing is particularly influential for one of our prototypes *DrawMyPhoto*, which will be discussed in more detail in a later chapter.

4.3.3 A Holistic Approach

This dissertation focuses on a *holistic* approach to improving drawing ability. This means that the overall goal is not to focus on a specific style of drawing such as design sketching or representational drawing, but rather to focus on a general ability to draw along with a healthy psychology around learning to draw. Individuals need to *want* to learn to draw in order to learn to draw.

Drawing utilizes a fairly complex set of abilities including fine motor skills, perception, visual memory, and spatial reasoning [8, 134]. In order to improve one's drawing ability, one must practice and put in the necessary time to cultivate these various abilities. In order to do so, individuals must be self-motivated and have some degree of self-efficacy. We identified four key aspects to learning to draw that should be cultivated in order to see the most improvement.

Skill Most obviously, skill is one of the key aspects of learning to draw. Skill includes fine motor control, perception, understanding of perspective, and so on.

Confidence As mentioned in previous sections, self-efficacy is closely tied to confidence and self-regulating one's learning. In other words, when an individual is more confident in their ability, they are more likely to be motivated to keep practicing and improving.

Motivation Naturally, motivation and confidence are closely tied to each other. Drawing, like any other skill, requires practice, and individuals must be motivated to practice.

Creativity Lastly, when one has gained sufficient ability, creativity becomes an important aspect of drawing. Tapping in to one's own creativity is where an individual can truly reap the benefits of the ability to draw.

All of these aspects work together and should be cultivated simultaneously.

5. RESEARCH QUESTIONS

5.1 R1. Can Intelligent Real-time Feedback from Sketch Recognition Help Improve Perspective Drawing Ability?

H1. By utilizing a novel sketch recognition algorithm for detecting accuracy of rectilinear sketches as a form of scaffolding, novices should be able to improve their understanding of perspective and demonstrate that through subsequent sketches.

5.2 R2. What is the Effect of Real-time Feedback Driven by Sketch Recognition on Drawing Ability and Drawing Self-efficacy?

H2. The relevant education theory suggests this will help students internalize and self-regulate their learning [110]. We would also expect to see correlated improvements in their self-efficacy and ability.

5.3 R3. What Motivates Students to Practice Sketching? What is the Effect of Game-based Learning and Gamification on Motivation to Practice Sketching?

H3. Game-based learning should have the most powerful effect on motivation to practice [135], but the nuances of this motivation will differ from student to student. Some students will be motivated more by gamification or different styles of games [120], and it will be on a spectrum from extrinsic to intrinsic motivation. We would expect to see correlated improvements in ability and self-efficacy as well.

5.4 R4. What is the Effect of Co-created Provocative Stimuli on Creativity?

H4. By providing provocative stimuli based on what students have sketched so far, it can serve as a form of scaffolding, and students should be able to avoid design fixation and produce more novel ideas than those without this stimuli.

5.5 R5. What are the Benefits and Limitations of Using an ITS for Sketching Fundamentals in Existing Courses?

H5. We would expect the system to be useful for drawing instructors for a number of reasons such as:

- Quickly getting students up to speed on basics and fundamentals of drawing.
- Seeing their activity and performance over time so that teaching can be more personalized.
- Engaging students more and making them more excited about learning drawing.

6. IMPROVING DRAWING CONFIDENCE

6.1 ZenSketch: A Novel Sketch-based Game for Practicing Line Work

This chapter contains previously published work[1].

6.1.1 Motivation

It is not difficult to find people who not only believe they can't draw, they believe they "can't even draw a straight line." As line work is fundamental for almost all drawings, it stands as a fundamental skill for anyone pursuing improvement in their drawing ability. We believe it may also be a confidence-booster, allowing learners to have more self-efficacy [18] with respect to drawing and be more motivated to learn more advanced drawing techniques.

In disciplines such as Industrial Design, a form of drawing known as design sketching (often shortened to just *sketching*) is taught to students. A typical introduction to sketching course in an Industrial Design program may allocate an entire week or two to line work [136]. A typical exercise involves connecting dots with straight lines in an effort to improve accuracy, line quality, and speed (Figure 6.2). This exercise builds the necessary motor function for more advanced sketching techniques, and allows for the learners to sketch at a level that approaches professionals.

With tablet devices and particularly stylus-based tablets that enable digital sketching becoming more ubiquitous, there is an opportunity to translate this crucial exercise in to the digital medium. While our previous work has explored this [137, 72, 138], there has been no attempt at introducing gaming mechanics to the exercise in an effort to make the exercise more fun and engaging for learners.

We have developed a game called *ZenSketch* that translates features of lines in to engaging game mechanics (see Figure 6.1). Our goal was to make an educational game that could improve freehand sketching line work while also being entertaining and motivating to learners. The game

[1]Reprinted with permission from: *ZenSketch: A Novel Sketch-based Game for Practicing Line Work*. Blake Williford, Adil Hamid Malla, Matthew Runyon, Wayne Li, Julie Linsey, and Tracy Hammond. Extended Abstracts Publication of th eAnnual Symposium on Computer-Human Interaction in Play, pp. 591–598, ACM, 2017

Figure 6.1: The game's title screen.

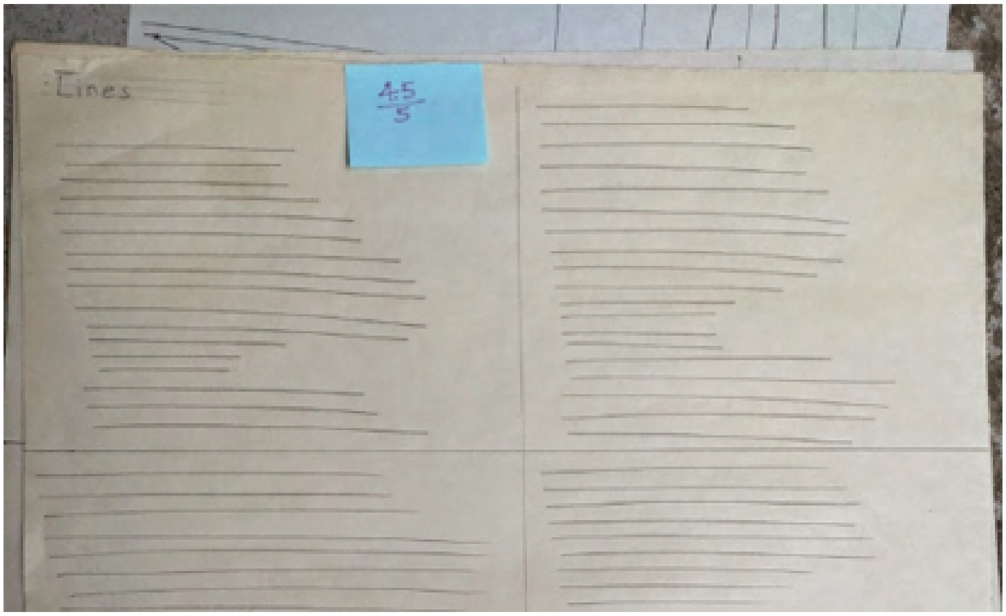

Figure 6.2: A student's homework assignment to practice lines by connecting dots. This exercise is very common and typically the first exercise in a sketching course. The student received a 4.5 out of 5 because many of the lines had an unintended arc. This is common among novices and is the result of using the wrist too much

relies on sketch recognition to extract features of the lines which include deviation from the "perfect" line (accuracy), speed, and smoothness (line quality). It then translates those features in to game play by encouraging each line to be as close to ideal as possible. The better the line work of the player, the longer they can continue playing and the higher score they can achieve.

6.1.2 Game Design

6.1.2.1 Pedagogical Basis

We have consulted with many domain experts including eight sketching instructors at three different universities, and have referenced the most widely used textbooks [125, 127] for learning sketching to determine what constitutes quality line work. In general, quality line work is *accurate*, *smooth*, and *fast*. These metrics often form the basis for grades that are given to students in an Industrial Design sketching course. They can be evaluated by an experienced instructor but are subject to human interpretation. As we will discuss later, sketch recognition can allow for objectivity in this domain and more precise determination of these metrics.

6.1.2.2 Accuracy

Every line in a sketch should have an intention behind it. The degree to which the person has achieved that intention could be described as accuracy, and it may be with respect to perspective, or with respect to an intended form or part of a form. A perfectly accurate line between two points will be the shortest path between those two points.

6.1.2.3 Smoothness

Koncelik describes a good quality line as having a clear beginning, a middle, and an end [127]. A novice may approach each line with a series of short choppy strokes which is often referred to as "chicken-scratching." This is discouraged in sketching because it results in sketches that are either unclear or too messy for further details to be added. Instead, smooth, fluid strokes are encouraged, which lead to more legible and easily interpreted sketches.

6.1.2.4 Speed

Speed is of lesser importance, however it is still encouraged in sketching because it allows for rapid ideation, and it can naturally produce more precise, smooth, and controlled lines. "Control does not come from slow, careful line production, it comes from fast, rapid movement of the drawing media on the paper surface" [127]. In this way sketching is similar to cursive writing.

It is important to note that these "rules" can be broken in certain circumstances. In fine art drawing, lines may intentionally be broken up or messy so as to render textures or suggest ambiguity. For the purposes of the game, and the exercise in general, we are focused on quality line work as described by domain experts in design sketching. This type of drawing is focused on rapid ideation, the construction of three-dimensional forms, and the clarity of those forms, therefore these rules should not be broken in general when sketching.

6.1.2.5 Game Mechanics

Translating these features in to game mechanics which encourage high values was an interesting challenge. Games with good flow have been shown to improve user performance [139] as well as keep users engaged [140]. As good game design requires immersion and clear goals [141], we decided the premise of the game was for the user to connect bridges for a fictional monk character named Maxus. By "bridging" the widely used connect the dots exercise with the digital medium, we could make the act of drawing a single line a goal within itself. This is important because when learning sketching, one must learn that every line *is* important and should have an intention behind it.

Maxus is a monk in training, and the player's goal is to help him with his training. He begins with three lives and a full balance meter. He must reach as far as he can while remaining balanced. The following game mechanics were implemented and can also be seen in the game's initial instructions screen. They are described below.

6.1.2.6 Precision

We decided to call accuracy precision in the game as it is more apt. Precision is encouraged through a game mechanic in which Maxus loses balance if the line deviates too much. He only has so much balance before losing a life. Additionally, more precise lines result in a higher score for each line.

6.1.2.7 Finesse

We decided to call smoothness finesse in the game, as sketching smooth lines does require some finesse and fine motor control.

6.1.2.8 Speed

Speed is encouraged through the collection of randomly generated power-ups that disappear after a few seconds. Zen orbs (25% chance) restore balance and add a multiplier to line scores. Hearts (5% chance) add a new life unless the player already has the max of three lives.

6.1.2.9 Stealth Assessment of Learning

Stealth assessment has been shown to help reduce user anxiety and increase performance [123]. The objective in designing the game was to at least partially hide assessment behind the game mechanics. This allows the user to be engaged with less fear of "grading" and learn better overall. We hide direct assessment behind game scores and bonuses. This encourages the users to focus more on getting a better score instead of directly focusing on their line quality. Since the users can only get better scores through better lines, we hypothesized that the line quality will increase naturally as users aim for higher scores in the game.

6.1.3 Evaluation

In addition to iterative informal testing which helped balance the game, we conducted a formal evaluation and data analysis in order to determine the effectiveness of the game in improving student's line work. This involved formally testing 14 university students. Students ranged from the undergraduate to graduate levels and had self-assessed sketching ability ranging from novice

Figure 6.3: Study participant playing the game on a Microsoft Surface Pro 3 in different device orientations.

to intermediate.

Each participant played the game on a Microsoft Surface Pro 3, and was encouraged to use the device in whichever orientation they pleased (Figure 6.3). Each participant was asked to play the game repetitively for 20 minutes, aiming for the high score each time they completed a game. We collected accuracy, smoothness, and speed data for each and every line in each game.

6.1.4 Results

Figures 6.4, 6.5, and 6.6 show performance of the 14 users for their first 100 lines (typically 5–7 games). Because the data was quite noisy, we chose to present average trends to better observe any subtle improvements.

We then performed a paired t-test comparing the average accuracy, smoothness, and speed for the first 10% and last 10% of the lines drawn by each participant to see if playing the game increased line sketching ability. From this, we were able to determine that playing *ZenSketch*

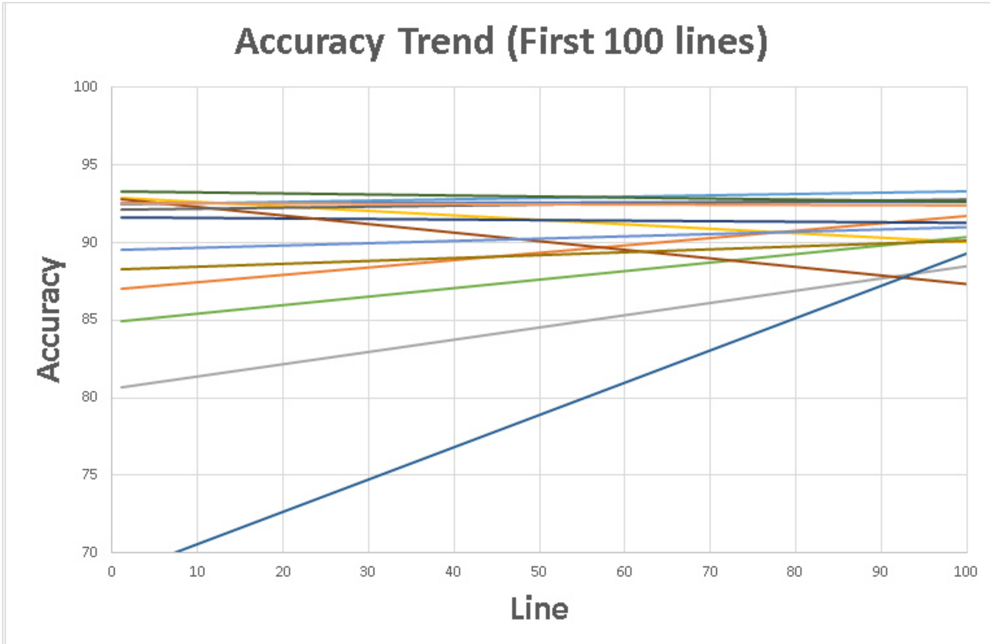

Figure 6.4: It can be seen that for most participants, accuracy trended upward or stayed stable.

statistically improved line sketching smoothness ($p < 0.02$), but not accuracy ($p < 0.15$) nor absolute speed ($p < 0.06$). When speed was standardized for each user (from 0% to 100% of their max speed), the increase in normalized speed was significant ($p < 0.03$).

The full short paper for this work was published at *CHI Play 2017* [142].

6.2 DrawMyPhoto: Automatic Generation of Drawing Tutorials from Photos

This chapter contains previously published work[2].

6.2.1 Motivation

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., [143, 144]), 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. [145]

[2]Reprinted with permission from: *DrawMyPhoto: A Novel Sketch-based Game for Practicing Line Work.* Blake Williford, Abhay Doke, Michel Pahud, Ken Hinckley, and Tracy Hammond. Proceedings of the 2019 Conference on Creativity and Cognition, 198-209

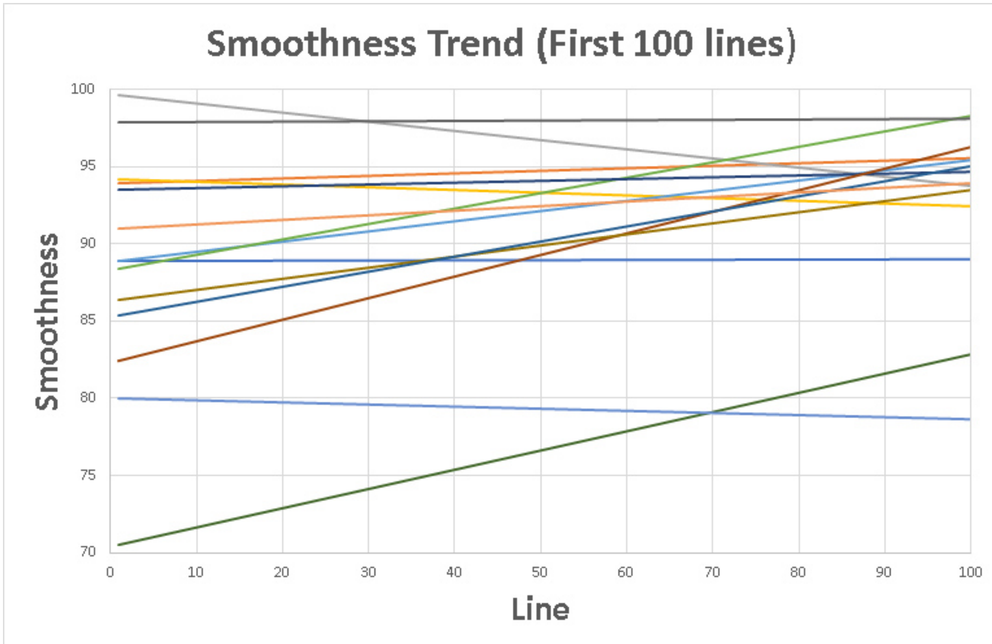

Figure 6.5: It can be seen that for most participants, smoothness also trended upward.

Figure 6.6: It can be seen that for most participants, speed trended upward or stayed stable. Interestingly, the average speed varied considerably between participants.

and the notion of "digital possessions" from Belk [146].

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 [128, 130, 125]. 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.

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.

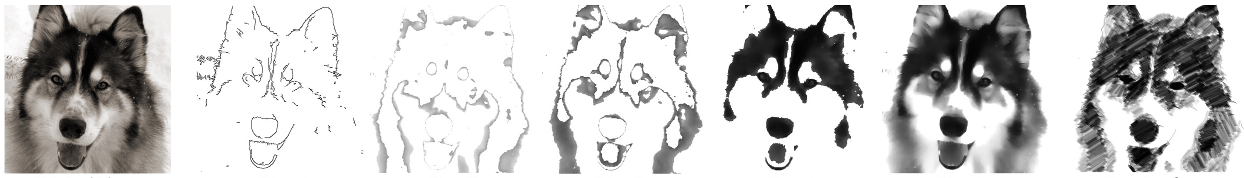

Figure 6.7: The DrawMyPhoto system generates an interactive drawing tutorial from a photograph by utilizing image processing. It generates five distinct steps and offers novel real-time feedback on pressure and tilt, allowing even novices to produce quality drawings as well as learn how stylus pressure and tilt affects shading.

6.2.2 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 [119]. This technique ensures the learner is adequately 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 6.8). 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 6.9). 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 [81, 82]. 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.

6.2.2.1 Pedagogical Basis

We consulted with a domain expert in drawing as well as referenced the leading literature in art education [128, 125, 130, 147] 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.

6.2.2.2 Image Processing

Figures 6.7 and 6.12 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 [128]). 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 [148]. 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 [149] 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 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 [150]. 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.

6.2.2.3 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 6.8). 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. [151]), 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 [128]. 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

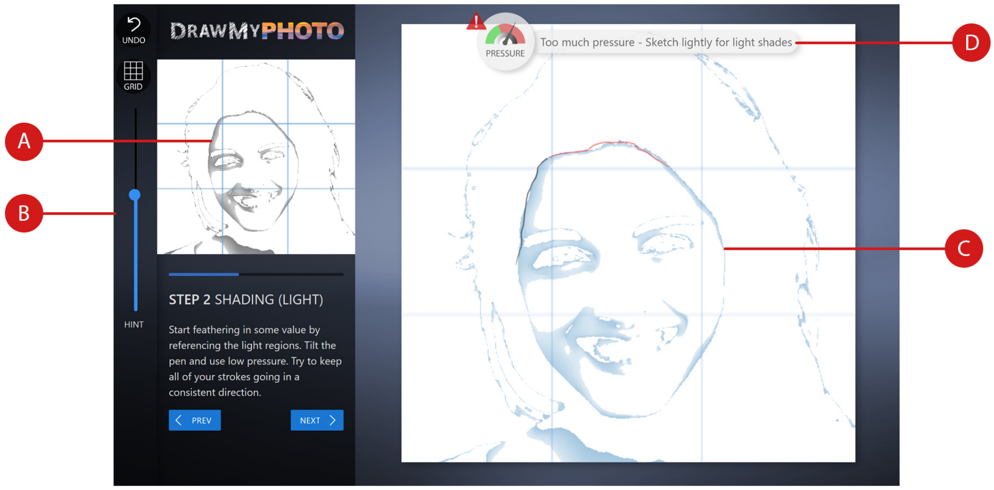

Figure 6.8: (A) is the sidebar with reference image for each step of the tutorial. It also includes some basic instructions for each step. (B) is a toolbar designed for bimanual interaction that allows user to adjust underlay feature, undo, as well as toggle the grid on and off by using the non-dominant 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.

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.

6.2.2.4 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 [152, 153], 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 6.9 and 6.10). 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

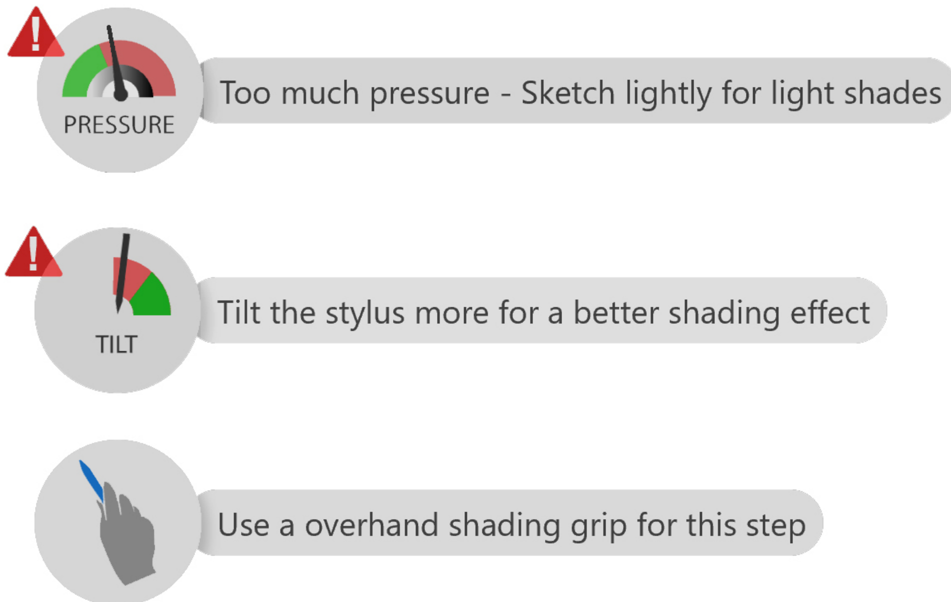

Figure 6.9: Examples of pressure and tilt feedback, along with grip suggestions. After piloting with users we found it was better to only show the feedback when necessary so that users would take notice. Constantly persistent feedback was often ignored. The "meters" and the associated ideal ranges in green were quickly grasped by most users. The values chosen allow for red-green colorblind people to still detect a difference.

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.

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 tilt$X$ and tilt$Y$ 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

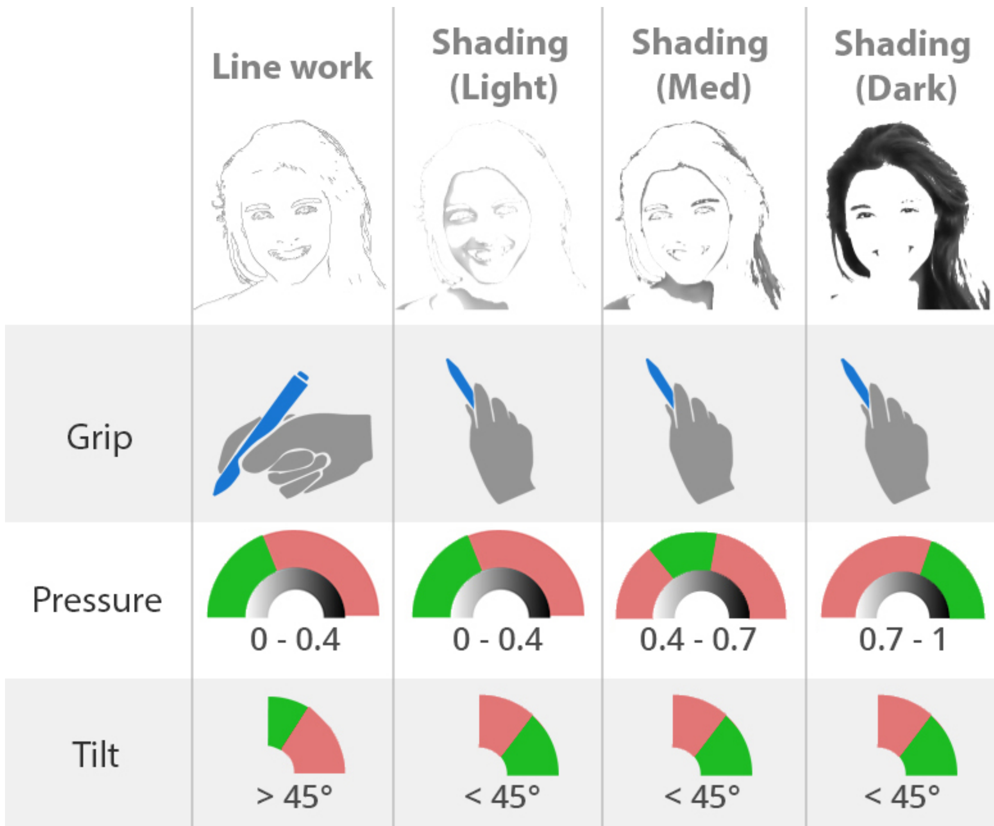

Figure 6.10: This shows how the feedback and ideal ranges change from step to step, and includes threshold values. Note that for Step 5 (Details) no feedback is given to the user because at that point they may use any range of pressure or tilt to finish the drawing, along with any grip they choose.

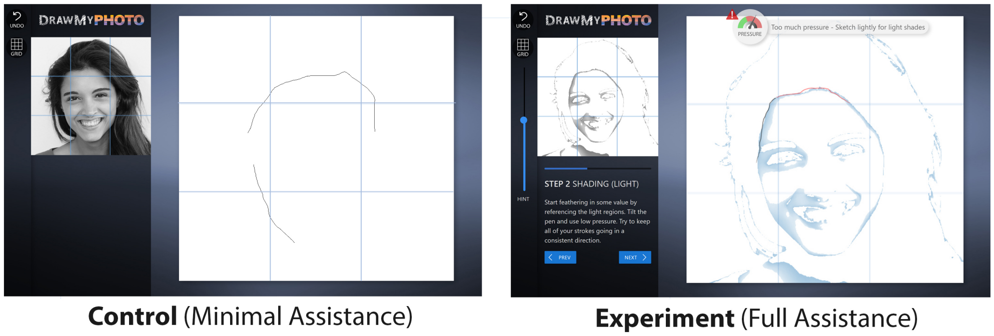

Figure 6.11: The control group experienced a minimal interface with just the grid (barely visible in this image) and the unprocessed reference photo to help. The experimental group experienced the full interface with grid, generated steps, underlay feature, and real-time feedback on pressure, tilt, and grip. We set up the study this way to avoid learning and history effects.

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 [128, 130].

6.2.3 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?*

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

6.2.3.1 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 6.11). 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 6.12) as the reference photo because many people want to learn to draw portraits [81], there is a lot of attention on portraits in the educational literature [128], and many previous research projects described earlier utilize portraits for their studies [81, 59]. This may allow for easier comparison to previous systems in the future, although that was not the main

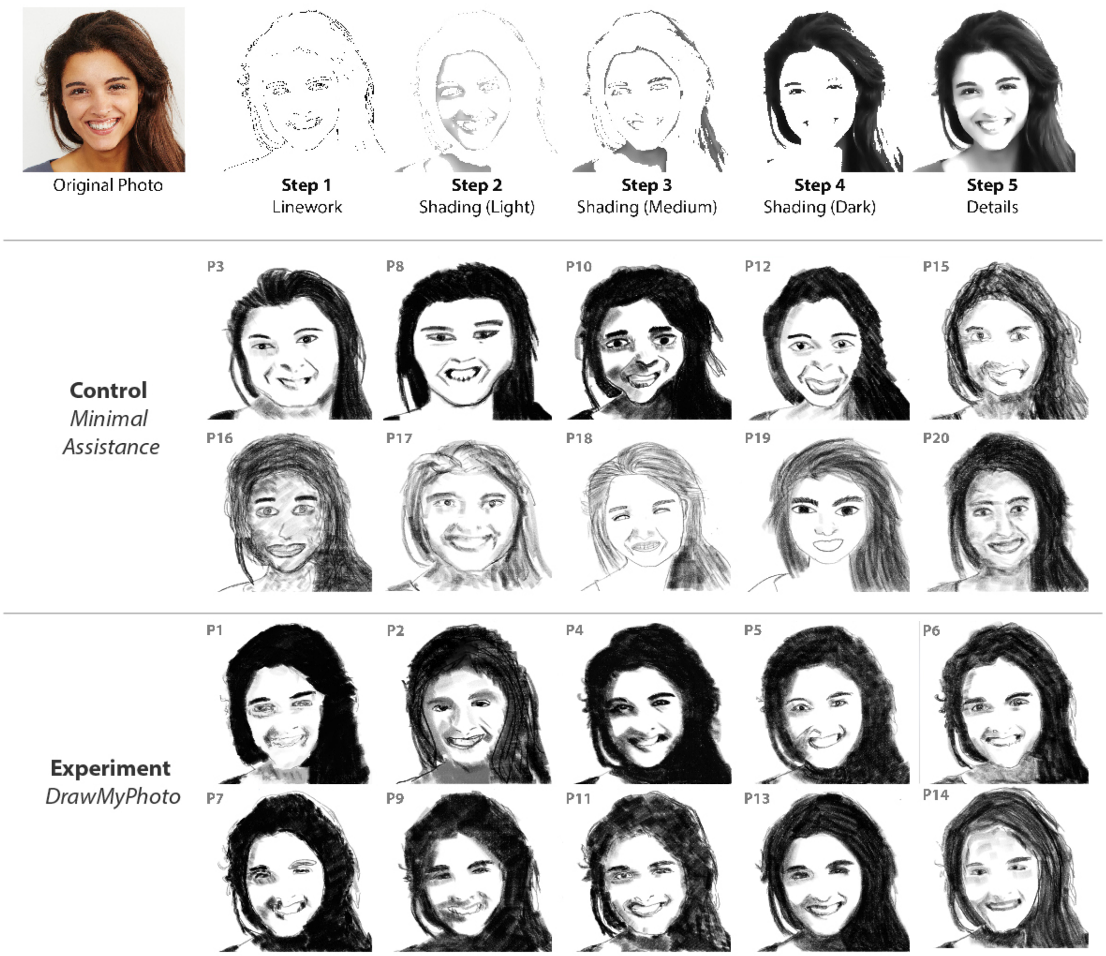

Figure 6.12: It can be readily seen that the assisted group which used the full application produced drawings that are not only more accurate, but have more consistent shading and attention to value, as well as attention to detail. This was confirmed by the expert ratings (See Figure 6.13). Note the 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

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.

6.2.3.2 Expert Ratings

We pursued expert ratings primarily to **answer R1**. We used an approach similar to Consensual Assessment Technique [154, 155], 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.

6.2.3.3 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.

6.2.3.4 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 [128], 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, average pressure, average tilt, and the total time spent to complete their drawings.

6.2.4 Results

The drawings that resulted from the study can be seen in Figure 6.12. In this section we will discuss and analyze the results and how they relate to our design goals and research questions.

6.2.4.1 Expert Ratings

The expert blind ratings are shown in Figure 6.13. 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 6.12 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.

6.2.4.2 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 6.14). We conducted a Wilcoxon Signed-Rank test between the groups and found statistical significance for both ratings of their own drawings

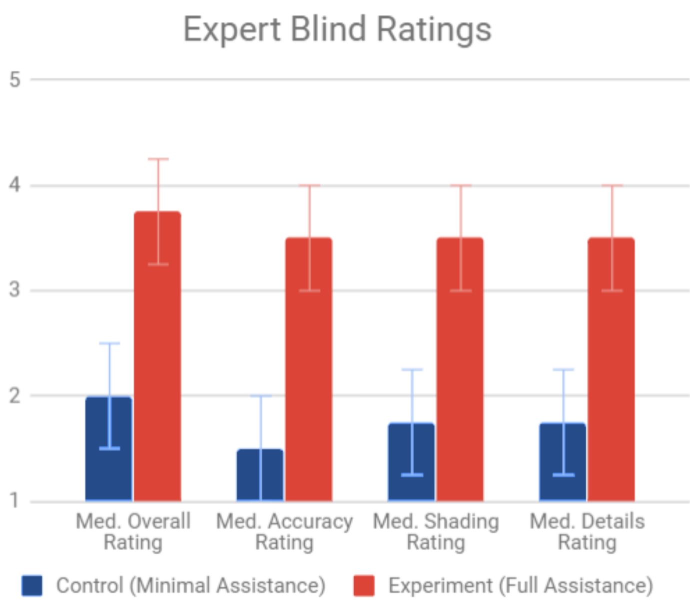

Figure 6.13: The experts gave statistically significant higher ratings to the experimental group in every category ($p < 0.003$ for overall quality, $p < 0.003$ for accuracy, $p < 0.005$ for shading, and $p < 0.006$ for details).

($p < 0.02$), 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 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 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.

Table 6.1: Log Analysis

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

"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 [18] for users. This is important with respect to motivation to practice and continuing to improve drawing skills.

6.2.4.3 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 6.1). 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 6.1. 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 approach to drawing. It is known that different

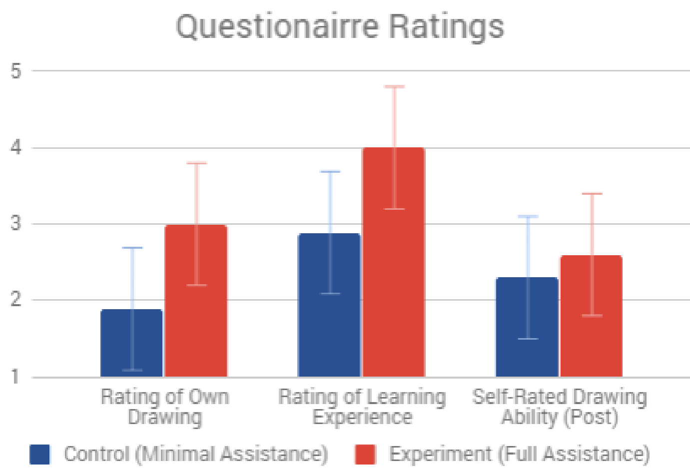

Figure 6.14: Participants in the fully assisted group rated their drawings higher, considered the learning experience better, and rated their drawing ability slightly higher than the control group.

people develop different, and sometimes very unique grips and approaches to handwriting and drawing [156]. 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.

6.2.5 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.

6.2.6 Discussion and Future Work

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

6.2.6.1 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" [157]. 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.

6.2.6.2 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 *percep-*

tion [128, 130]. We believe it would be important to fade the scaffolding away in the application over time [158], 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.

6.2.6.3 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. [159] 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., [160, 161]) 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.

6.2.6.4 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.

The full paper for this work was published at *Creativity and Cognition 2019* [162].

7. IMPROVING MOTIVATION TO PRACTICE DRAWING[1]

7.1 A Framework for Motivating Sketching Practice with Sketch-based Gameplay

This chapter contains previously published work.

7.1.1 Motivation

Sketching has traditionally been considered a valuable skill for professions such as industrial design, architecture, and engineering as it has been shown to benefit problem solving ability [1] and improve communication skills [4]. It is also increasingly valued in the multidisciplinary field of human-computer interaction (HCI) as a means to rapidly generate ideas for user experiences [14, 16, 15]. However, the skill relies on many different fine motor skills and cognitive skills working together [134, 163] and it can take an individual years of practice to make improvements.

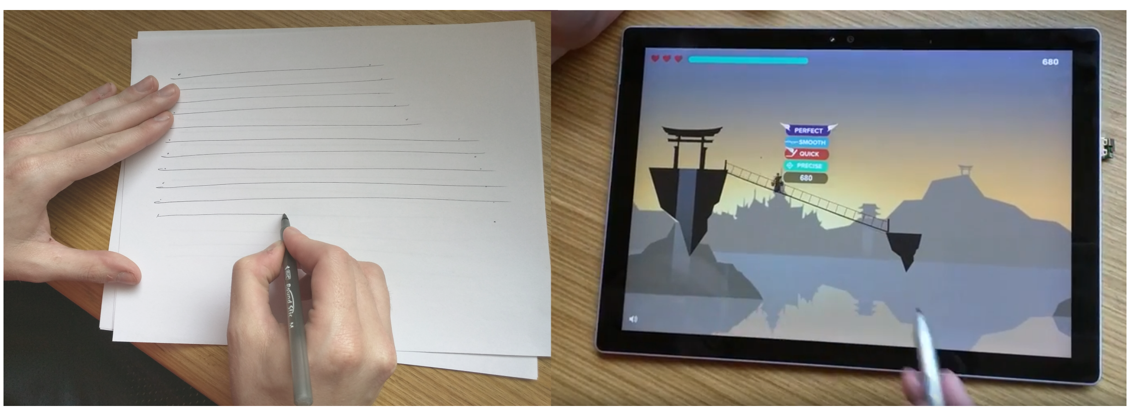

Figure 7.1: Sketch-based gameplay can help motivate practice of sketching fundamentals. We developed a theoretical framework for motivating individuals of varying skill levels, and explored two different gameplay approaches to a popular connect-the-dots exercise for practicing line work. The left image shows the traditional exercise with pen and paper, while the right image shows one of our gameplay approaches *ZenSketch*.

It is not uncommon to find people who lack confidence in their sketching ability, and a common

[1]Reprinted with permission from: *A Framework for Motivating Sketching Practice with Sketch-based Gameplay*. Blake Williford, Matthew Runyon, Josh Cherian, Wayne Li, Julie Linsey, and Tracy Hammond. Proceedings of the Annual Symposium on Computer-Human Interaction in Play 2019

remark one might hear is "I can't even draw a straight line!" This form of low self-efficacy [18] and learned helplessness is common in many people and often begins in childhood [21] as many children abandon artistic and creative activities. This low self-efficacy naturally makes it difficult to be motivated to practice. Even students studying industrial design and learning to sketch at a professional level struggle with motivation and can view practicing fundamentals as a "chore" [72].

Gamification, game-based learning, and serious games are promising areas of research which may help to motivate students in educational contexts. Serious games have shown to be more effective than traditional educational methods [105], but not necessarily more motivating. On-going research in these areas is exploring the degree to which games can motivate people and in what manner [103, 122, 104, 164]. There has even been exploration in to a comprehensive design framework for gamification in order to push the emerging field towards "inspiring and empowering" gameplay versus the more sinister "manipulative and obsessive" gameplay [165].

Our goal with this research was driven by two main research questions:

R1—*What motivates individuals to practice sketching and how do those motivations change with skill level?*

R2—*How are our current approaches to sketch-based gameplay motivating students?*

To answer R1, we conducted a grounded theory study with eighteen participants, including four professionals. To answer R2, we conducted a semester-long implementation of two previously designed approaches to motivating line work practice with sketch-based gameplay in a university sketching course as well as a high school art course.

Our contributions in this work include:

- An analysis of how motivations of individuals with different sketching skill levels change and evolve. This analysis may be useful for other researchers, educators, and technologists involved in design education and/or art education.

- A framework for how various approaches to sketch-based gameplay can target these motivations and encourage sketching practice for individuals with different skill levels.

- The results of a semester-long case study of two sketch-based gameplay implementations in a university sketching course and a high school art course. We gathered interesting insights about how students engaged with our approaches and how they were motivated.

7.1.1.1 Motivation Study

In an effort to understand what motivates students to practice their sketching ability, we initially conducted a grounded theory study [166] with eighteen participants, including fourteen undergraduate and graduate industrial design students and four established industrial design professionals. Seven of the students were male, while seven were female. Figure 7.2 is a table that provides more detail on the participants.

Semi-structured interviews were conducted with these individuals with a primary focus on understanding what motivates them to practice sketching, and what factors hinder their motivation or discourage them. Interviews lasted approximately 20–30 minutes. Theoretical sampling [167] was used to ensure a variety of skill levels were represented as well as both males and females. The constant comparative method [168] was used to continuously update codes and categories as interviews were conducted.

Participants were asked to self-identify as being either **Novice**, **Intermediate**, or **Advanced** at sketching. Descriptions of these categories are below:

Novice—*I am still learning fundamentals and have much to improve on. I am not very confident in my sketching ability.*

Intermediate—*I have mastered some fundamentals, but am still improving and learning. I am somewhat confident in my sketching ability.*

Advanced—*I have mastered fundamentals and advanced methods and am very experienced and confident in my sketching ability.*

Ultimately, four categories emerged which encompass what motivates students to practice their sketching ability: *Achievement*, *Competition*, *Communication*, and *Creativity*. A high level overview of the categories can be seen in Figure 7.3. More detail about these categories is described in the following sections.

STUDENTS			
CODE	GENDER	AGE	SKILL LEVEL
S1	Female	25	Novice
S2	Male	24	Advanced
S3	Female	24	Intermediate
S4	Male	25	Intermediate
S5	Male	25	Novice
S6	Male	29	Intermediate
S7	Male	21	Intermediate
S8	Female	25	Novice
S9	Female	28	Novice
S10	Female	23	Novice
S11	Male	26	Novice
S12	Female	21	Intermediate
S13	Female	25	Advanced
S14	Male	22	Novice

PROFESSIONALS			
CODE	GENDER	AGE	SKILL LEVEL
P1	Male	41	Intermediate
P2	Male	42	Advanced
P3	Male	24	Intermediate
P4	Male	25	Advanced

Figure 7.2: A Table showing participants interviewed during our grounded theory study.

7.1.1.2 Achievement

Particularly prominent among the most novice students interviewed, a sense of achievement was a strong motivator for practicing sketching. This was the most extrinsic and ego-driven motivator. Some students wanted to just "feel" better at sketching, while others were looking for external praise from their instructor or peers. Some were simply driven to avoid negative criticism from their instructor or peers.

"I've never felt too motivated because I never get any praise from my teacher."—S10

Figure 7.3: The four categories that emerged from our grounded theory analysis, along with high-level codes that contributed to the categories.

> *"I'm afraid of my instructor criticizing me. I want to do well."—S1*

7.1.1.3 Competition

Some students expressed a desire to be the best in the class, or to have their sketches highlighted during critiques. Others simply wanted to avoid being the *worst* in the class. Such motivations fall under the category of competition. There was also a consensus that if there was going to be competition, everyone should be around the same level.

> *"Competition can definitely motivate me, but only if they are around my skill level."—S3*

> *"Some friendly competition and seeing my friend's sketches was always a motivator for me when I was a student."—P1*

7.1.1.4 Communication

Some students were focused on the practical reasons for learning sketching such as collaborating and communicating ideas. These students tended to be more intermediate or advanced and were focused on their careers and finding jobs. From that perspective, the students were motivated by producing higher quality drawings that communicated ideas better and could be better understood by their peers.

"If you have an idea and you can't put it on paper, it doesn't count."—S6

"I'm always inspired to practice by seeing work from the pros like Syd Mead and Joe Johnston."—P1

7.1.1.5 Creativity

Some of the most advanced students and the professionals we interviewed were motivated by novelty and a desire to make the most creative ideas with sketching. The students who were most motivated by this were also approaching graduation and were focused on their careers. Since sketching and the rapid generation of ideas is a key part of the discipline of industrial design, it's no surprise that these individuals were motivated by creativity.

Interestingly, we found that the most advanced individuals also rejected competition as a motivator, with some even suggesting it can be a negative. This suggests a shift towards intrinsic mastery-driven motivation versus extrinsic ego-driven motivation. This is an ideal evolution of motivation from the standpoint of SDT [117].

"Competition is not what sketching should be about."—P4

"I've been in a position where others are discouraged because of my sketching ability, so I think it should be more about personal mastery."—S2

7.1.2 Motivation Framework

A fundamental insight of the grounded theory study was the realization that individual's motivations evolve based on skill level and confidence. Figure 7.4 shows an approximation of how individual's motivations change with skill level in a radar chart inspired by the gameplay motivation work of Richard Bartle [120], among others. We produced the charts by coding instances of each category for each participant and averaging them. While not a rigorous approach, it provides a general sense of these changes in motivation.

This insight influences a key component of our motivation framework—To employ different motivators based on an individual's self-efficacy and skill level, be it self-evaluated or determined

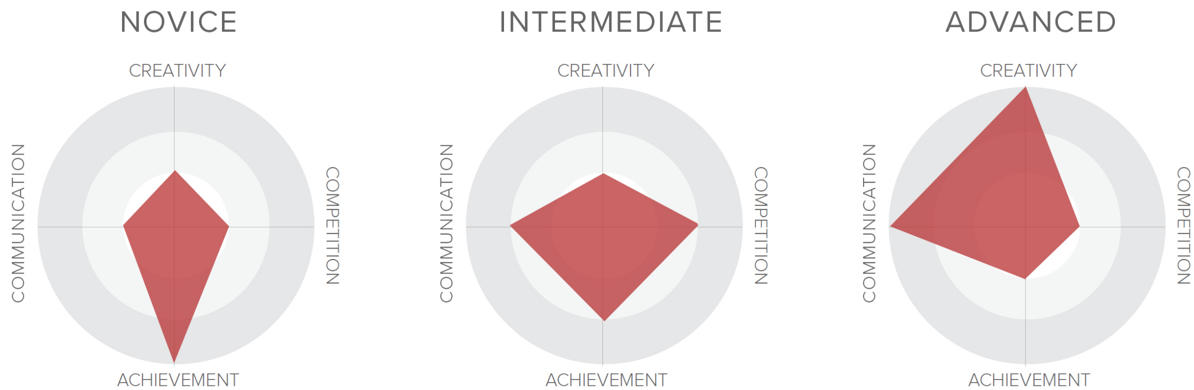

Figure 7.4: These radar charts depict an approximation of how motivations evolve as one's skill level increases in sketching. Naturally, not all individuals follow this exact evolution, but it was the general trend we found. Novices are primarily in need of a sense of achievement. As they become intermediate, their motivations can expand to competition and communication. Eventually, as they master sketching and become professionals, they see sketching as a useful *tool*, and are more motivated by creativity and communication, while being less interested in ego-driven motivations like achievement and competition. This follows self-determination theory.

by an intelligent system based on performance. The following sections explain this approach for novice, intermediate, and advanced individuals.

7.1.2.1 Novice

It is important for many novices to have a sense of *achievement*, providing a boost in self-efficacy and motivation that can carry over in to more time spent practicing. Achievement in the context of sketch-based gameplay can include high scores, RPG-influenced stat and skill improvements, improvements in performance metrics, and completed artifacts that evoke a sense of pride. Previous work has shown that when individuals invest time and effort in to creating something, it holds more value to them [145]. We believe it is best for novices to compete with themselves initially, rather than competing with peers, since many individuals expressed being discouraged when compared to peers.

Recommended motivators for Novice skill level:

- Single-player gameplay

- High scores
- Achievements / badges
- RPG-inspired leveling / stats
- Assistance in producing quality artifacts

7.1.2.2 Intermediate

At the intermediate level, motivators can expand to include *competition* and *communication*. At this point, individuals may have developed enough sketching self-efficacy to no longer be discouraged by comparison to peers.

Friendly competition can bring out strong motivation in individuals, particularly those who already identify as being competitive. Some individuals may not identify as being competitive at all, so we believe competition should always be an *option* rather than a requirement. Competition in the context of sketch-based gameplay can include leaderboards and thematic sketching competitions.

Communication can also be a strong motivator at this stage because individuals may be starting to understand that sketching is a useful and practical "tool." *Pictionary* and *DrawSomething* are examples of how gameplay can foster communication skills and possibly motivate someone to improve their sketching ability, although the games are not explicitly designed for that purpose. Nevertheless, the more clearly an individual can communicate an idea, the better they can perform in such games.

We recommended introducing these motivators for Intermediate skill level:

- Multiplayer gameplay
- Themed sketching competitions
- Collaborative gameplay
- Leaderboards

- Communication game mechanics

7.1.2.3 Advanced

As individuals become more advanced at sketching, they begin to tap in to their creativity and be highly motivated by producing truly novel ideas and forms. At this point it can be advantageous to use *creativity* as a motivator. Creativity in the context of sketch-based gameplay can include idea generation game mechanics, puzzles that require divergent thinking, etc.

It's also important to consider that at an advanced level, individuals can inspire others with their sketches and those that we interviewed expressed they were motivated to do so. It may be beneficial to bring in *community* aspects, and give these individuals the opportunity to both share their sketches and teach others concepts.

We recommended introducing these motivators for the Advanced skill level:

- Idea generation game mechanics
- Opportunities to share sketches
- Opportunities to teach others

7.1.2.4 Limitations

We recognize that the proposed four-component motivation model and Figure 7.4 are based on an exploratory qualitative analysis with students from one university industrial design program along with some associated professionals, and should be validated in future work using quantitative methodologies. We believe that the model does provide a *general* sense of the motivations of individuals with different skill levels learning sketching, however it could be improved and expanded upon with more quantitative methods such as as motivation scales [169] and more rigorous studies.

7.1.3 Case Study

We conducted a study with students who were all novice to intermediate skill levels with respect to sketching, and explored two of the core motivators of our motivation framework: *Achievement* and *competition*. We focused mostly on understanding how these motivators influenced the students, and how much they encouraged more practice. We also wanted to compare two different approaches to sketch-based gameplay for basic line work—A gamified line lesson and a serious game for practicing line work called *ZenSketch*. Both are web-based prototypes that can be used on any device with an internet connection, and are described in more detail below as well as in the companion video.

7.1.3.1 Gamified Line Lessons

The gamified line lesson is based on a popular exercise in design sketching pedagogy which involves connecting dots to improve freehand line work [125, 127, 170]. This exercise promotes accurate and fluid freehand line work and is common warm-up and practice exercise in the first few weeks of a design sketching course.

We "gamified" the exercise by providing real-time feedback in the form of deviation lines in red. Feedback is important in gamification and serious games, particularly to help players self-regulate their learning and achieve the desired learning outcomes [171, 172]. We also include a 5-star score after 8 lines have been sketched and provide feedback on average precision, smoothness, and speed, all important aspects of design sketching [125, 127, 170]. See Figure 7.5 for more detail on the gamified line lessons.

7.1.3.2 ZenSketch

ZenSketch is an approach to the same connect-the-dots exercise influenced by game-based learning and serious games [142]. The feedback and game mechanics are influenced by stealth learning since it has been shown to help reduce user anxiety and increase performance [123]. Rather than sketching "lines," the player is actually focused on sketching bridges for the character. The feedback comes in the form of bonuses and individual line scores, which communicate what the player is doing correctly (see Figure 7.6). The summative feedback is in the form of a gameover screen which shows average performance in precision, smoothness, and speed as well as providing some tips based on these values e.g., "Sketch faster to have smoother lines!"

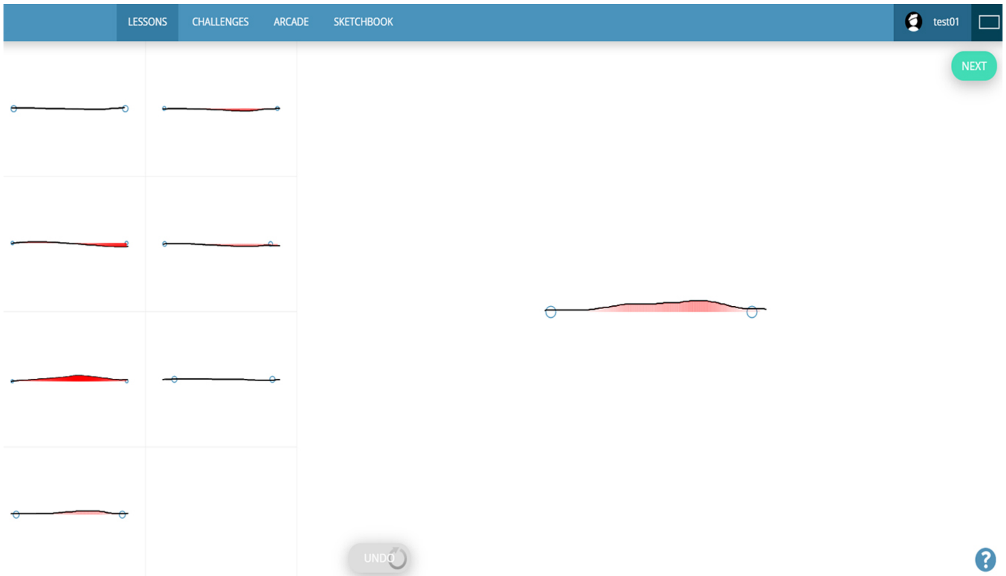

Figure 7.5: Screenshot of the gamified line lesson. The lesson provides some basic feedback on accuracy and a summative feedback screen, but is otherwise just a digital version of the classic connect-the-dots exercise.

ZenSketch also adds a leaderboard component (see Figure 7.9) to motivate individuals who are competitive. Utilizing high scores and leaderboards as a motivator in games has been utilized for decades [173], particularly in arcade games where players compete for "bragging rights."

7.1.3.3 Methodology

We implemented the gamified line lessons and *ZenSketch* in two courses for a semester-long study in Fall 2017. One course was an art course at a Giddings High School in Texas, and the other was *ID1418*, an introductory Industrial Design sketching course at Georgia Institute of Technology. All participants had access to both the game and the line lessons in a system called *Sketchtivity* [138]. They had freedom to use one or the other approach, or both.

There were a total of 150 students that engaged with the systems, including 58 students from the Texas high school, and 92 students from the Georgia Tech sketching course.

We allowed the instructors to implement *Sketchtivity* in their classes as they deemed suitable so as to not interfere with their teaching methodology. This encouraged more naturalistic gaming and resulted in data which could be described as *organic*. The students had access to the game and

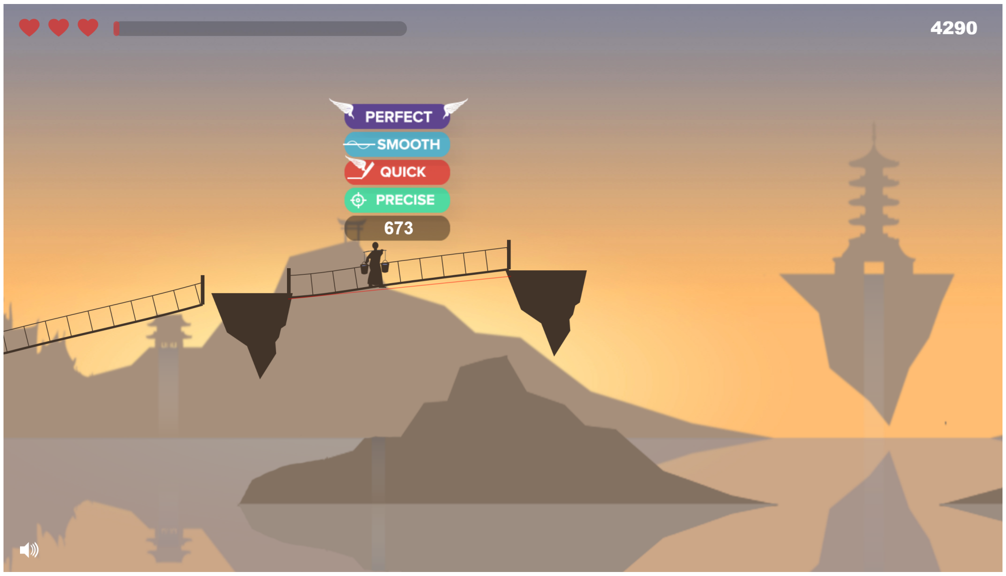

Figure 7.6: Screenshot of *ZenSketch*. The game provides real-time line feedback and bonuses during gameplay. The game also provides summative feedback and a high score, allowing for more competitive gameplay.

line lessons and could play them at their own leisure instead of being strictly required to play it at certain times, although the university students were encouraged by their instructors to practice periodically. The importance of this methodology was to be able to make observations on how well the approaches motivated the students.

In an effort to understand which approach students seemed to prefer, we logged lines practiced, time stamps, games played, scores, and metrics for each line including precision, smoothness, and speed data. We also conducted semi-structured interviews with three Georgia Tech students who played *ZenSketch* extensively for more specific insights on the experiences they had while playing the game over the semester. It is important to note that we did not look at overall improvement trends because this data was very asymmetrical, however we did investigate the performance of the students who played the game extensively.

7.1.4 Results

In general, both approaches were engaging and motivating to students with 72,842 lines practiced overall and the average student practicing 243 lines ($\pm$ 425.12) over the semester.

Collectively, nearly twice as many lines were practiced with the gamified line lessons (47,842) versus *ZenSketch* (25,002). This is a statistically significant difference (**p = 0.0001**) when comparing data from the two systems with a paired t-test.

The students who were highly engaged by *ZenSketch* were a minority. The students that played the game 10 or more times comprised only 26 of the 150 students, but accounted for nearly 64% of the 25,002 lines practiced in *ZenSketch*. In later sections we further investigate three of these participants.

The following results are organized in themes we found for the semester study based on analysis of the data. We did not focus on comparing activity between the two courses since the data was asymmetrical (both prototypes were used much more in the Georgia Tech sketching course where it was a requirement to practice for at least 15 minutes each week).

7.1.4.1 Overall Engagement

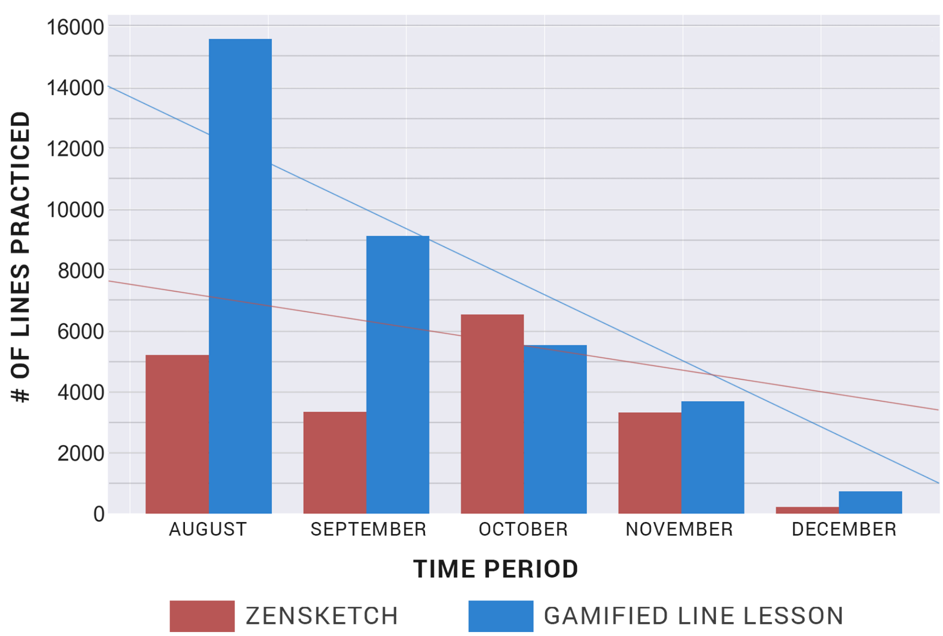

Figure 7.7: A histogram showing lines practiced over time for both *ZenSketch* and the gamified line lessons. There was a natural drop-off over time in both cases, but in October there was more lines practiced with *ZenSketch*. This happens to coincide with when the high school students used the system the most.

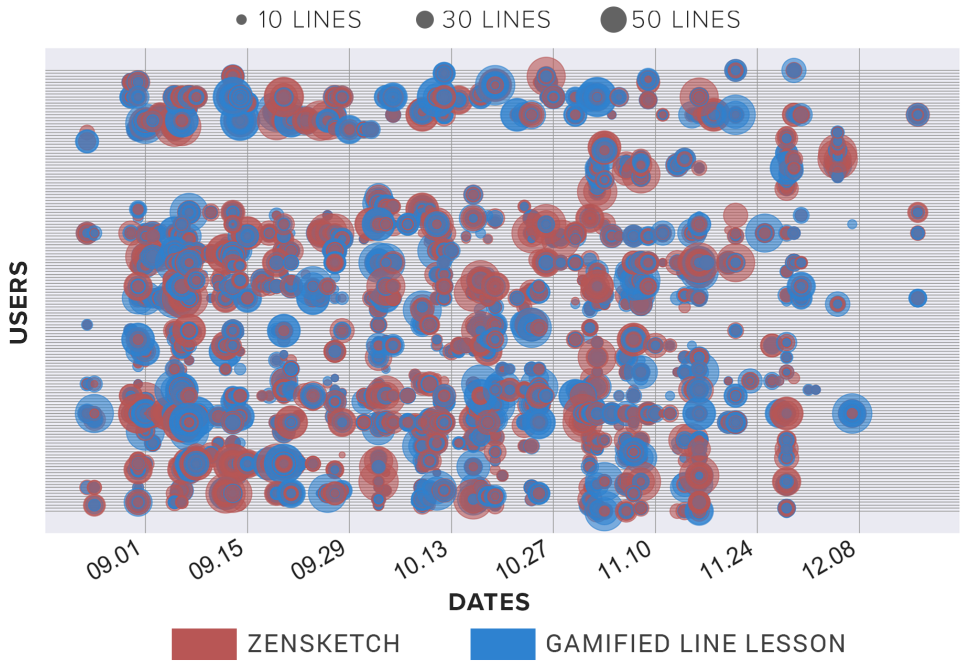

Figure 7.8: Game plays of all 150 users over the semester study. Each horizontal line in the chart corresponds to a user. If multiple lines were practiced in a day (sessions) we encoded that in to bubble size. It can be observed that players tended to practice lines in "bursts," likely in an effort to attain a better star rating (achievement) or a better rank on the leaderboard (competition).

As one might expect, there was clearly initial excitement, but some natural loss of interest over the course of the semester. Figure 7.7 shows lines practiced from when they created an account and had access to the system.

Figure 7.8 shows lines practiced over the course of the semester in a bubble chart. It can be observed that participants tended to practice lines in sessions or "bursts." This is likely due to motivations to improve on their score and/or the competitive aspect of the game which is discussed in the following sections.

7.1.4.2 Competition as a Positive

We found that the high score leaderboard component of the game (Figure 7.9) was a strong factor in encouraging repeated plays, particularly for participants who identified as being competitive.

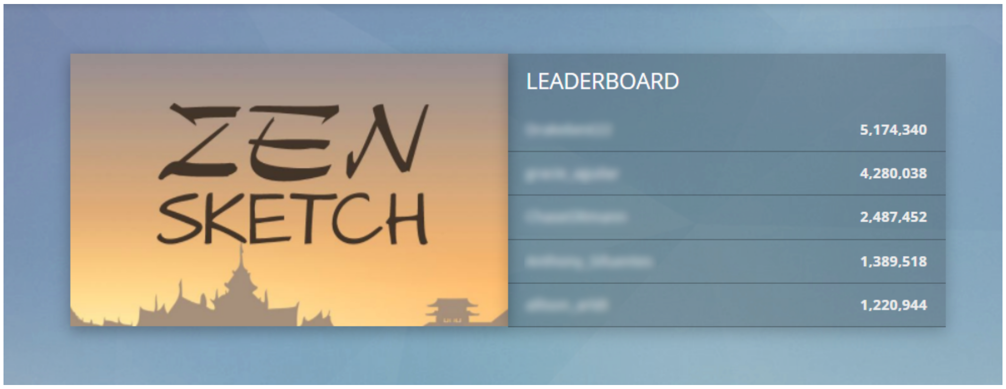

Figure 7.9: The leaderboard was found to be a highly motivating factor for some students, including participant CS1 who played the game 58 times in an effort to achieve the top score.

Participant CS1 was an undergraduate student at Georgia Tech studying industrial design who was noteworthy for having played the game the most of any student—58 game plays in a period less than two months, practicing 1,711 lines, also the most of any student. We found that she was primarily motivated to play the game as much as she did because she found it "addictive" and was aiming for the high score in her class of 92 students.

> *"I found it was a pretty addictive game." "There was one particular score at the top. They had a top score and I was in 2nd place or something. I was wanting to try to beat it. I could never beat it!"—CS1*

Participant CS2 was another undergraduate student at Georgia Tech studying industrial design who played the game 14 times in her own effort to get the top score.

> *"...also having the scoreboard at the end where you can see oh this person in my class got like a million points, it gives you a goal—I have to beat them. I got more in to it when I was trying to beat other people's scores."—CS2*

CS1 found the game to be much more appealing than the traditional connect-the-dots exercise common in the first few weeks of a design sketching course. She also found it more appealing than the gamified line lessons.

"I was doing a lot of the exercises but honestly was just getting kind of bored with it... And the game was just a lot of fun."—CS1

CS2 felt similarly, noting that the game made her feel more accomplished.

"Having the game definitely makes motivation easier. It makes you feel like you're actually accomplishing something instead of just like doing practice for nothing."—CS2

7.1.4.3 Competition as a Negative

One participant, an undergraduate Mechanical Engineering student from Georgia Tech whom we will identify as O1, admitted to using a laptop with a track pad instead of a stylus to achieve the highest score of all participants in both studies. The study was relatively informal and encouraged the use of any number of devices, therefore we could not account for this. We had been suspicious that this participant utilized a bug in the game to achieve her score, but she openly admitted to using her track pad in the interview.

This data was thrown out, however we believed it was important to mention as an example of how extrinsic motivators such as competing with others for a high score can sometimes lead to unintended and unwanted behavior in serious games. O1 was so motivated to achieve the highest score in her class that she bypassed the intention of the game and did not cultivate improvements in her line work through using a stylus.

7.1.4.4 Self-Perceived Improvement and Actual Improvement

Both CS1 and CS2 believed that they improved their line work by playing the game over the course of the semester.

"It could still be a little tedious at times since it was just lines. But I did feel like I was a lot more confident at drawing straight lines after that."—CS2

We analyzed their data and found that CS1 kept precision and smoothness high while trending better in speed over time (Figure 7.10). CS2 improved in precision, but sacrificed smoothness and

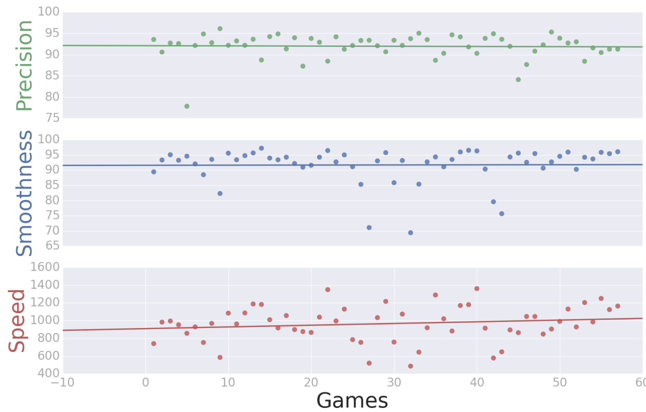

Figure 7.10: CS1 performance over 58 *ZenSketch* games

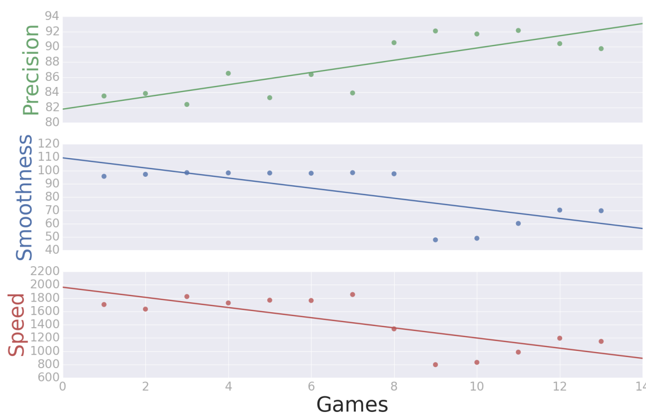

Figure 7.11: CS2 performance over 14 *ZenSketch* games

speed (Figure 7.11). This is consistent with a common “trade-off” that is made between precision and speed in motor control tasks in general [174, 175] and specifically in sketching [176]. People tend to sacrifice some precision in favor of speed and rapid idea generation when sketching.

This also suggests that significant *ZenSketch* gameplay does not necessarily correlate to a correspondingly significant improvement in line work. This may mean that only a moderate amount of gameplay can result in the desired improvement in line work that the game is designed for and further plays may have diminishing returns.

7.1.4.5 Translation to Paper and Traditional Media

When asked if she felt like playing *ZenSketch* on a tablet translated to better line work on paper, CS1 responded that she believed it did.

> *"Oh yeah, yeah it definitely did—At least for lines."—CS1*

Other work in this area has shown it to be a very promising approach [96, 95, 177]. Independent of drawing on a tablet or paper, the user is cultivating the same fine motor skills, namely holding a precision instrument and moving it with increased confidence, precision, and fluidity.

7.1.5 Discussion and Future Work

This section includes lessons we've learned from this work, and general principles that could be useful for other researchers, educators, and technologists exploring how serious games, game-based learning, and gamification can motivate students.

7.1.5.1 Diving Deeper into Motivation

Because of the intentionally open nature of the implementation and the organic use of our systems, it is difficult to more deeply understand student engagement and motivation from this data. Some things worth noting that influence these results:

- A greater number of lines can be practiced in a shorter time frame with the gamified line lessons. This may contribute to the overall trend that more lines were practiced with the line lessons.

- The university students were required to practice at least 15 minutes a week as part of a participation grade during the entire semester, whereas the high school students had more freedom to use or not use the system over a shorter period. This may have influenced them to use the tools quite differently.

- Students may have had limited access to quality devices that allow for stylus input, influencing how they engaged with the system. This has been a concern with previous studies [72].

A more controlled study in the future, and motivation scales [169] along with more interviews could determine more nuances of how these approaches are engaging and motivating students.

7.1.5.2 Consistency with Our Motivation Framework

From the lens of our motivation framework and SDT [117], the results of the case study make sense. Most of the participants were novices with little to no exposure to sketching, so many of them may have been deterred by the competitive aspect of *ZenSketch*. While the participants we highlighted loved the game and were very motivated by it, they were a minority among the participants from the university course. They did self-identify as being fairly confident and intermediate at sketching, so this is consistent with them being open to the competitive aspects of the game and its leaderboard, where comparison to peers is inevitable.

This does not mean the competitive aspects of *ZenSketch* were a poor design choice, rather that they only engaged a certain subset of people who were likely competitive by nature and/or more confident in their sketching skill. This is why we suggest that competition be used as an *option* for gameplay, but not as a requirement. Variety is key, and some individuals may never be drawn to competition regardless of their skill level.

7.1.5.3 Utilizing Variety in Motivators

It is well known that motivations differ greatly from student to student. Even within a medium like video games, there is a spectrum of preferences from individual to individual. Based on the results of our study, as well as recent research on rewards in games [121], we recommend researchers and game designers consider as much variety as possible when creating gamified or game-based learning solutions to motivate students. This could include adaptive systems that give the ability for students to self-select what motivates them, or select whether or not they are competitive. It could also include games that offer many different motivators within the same experience such as high scores, collectible items, cooperation with others, character growth and development, etc. It's worth noting that there is no "one-size fits all" approach and that it is highly dependent on the domain and context.

7.1.5.4 Advancing Sketching Education

In order to reap the many benefits of sketching, students must practice, and educators must consider all tools available to them to motivate their students. We believe this motivation framework can be one such tool.

We want to emphasize that this framework does not have to be limited to *digital* gameplay. Indeed, gameplay has been utilized in educational contexts long before the digital era. As was discussed earlier, sketch-based gameplay can include traditional pen-and-paper games like *Tic-tac-toe* and *Hangman*. Even modern educational approaches such as "sketch aerobics" utilize traditional media. The most important objective it to use sketch-based gameplay to motivate individuals of different skill levels to practice sketching, regardless of the particular medium used. Granted, digital media does not allow for certain advantages such as real-time feedback, more immersive experiences, and opportunities for remote gameplay.

We also want to emphasize that although this framework is oriented towards sketching in industrial design, many of the same principles are likely to hold true in other drawing contexts including fine art drawing [128], sketching in HCI [16], concept sketching [125], cartooning, etc. In all of these domains, students will struggle with motivation from time to time. Sketch-based gameplay can be a way to bring joy in to the learning experience while also sparking motivation. We encourage educators to experiment and researchers to share their findings.

8. IMPROVING DRAWING ABILITY[1]

8.1 Recognizing Perspective Accuracy: An Algorithm and Intelligent Interface to Assist Novices

This chapter contains previously published work.

8.1.1 Motivation

Drawing is valuable skill to learn for a wide variety of reasons. It can benefit internal cognition and problem solving ability [1]. It can improve peripheral skills like writing, brainstorming, and visual communication [4]. It can even promote better overall academic achievement in some students [6]. In this work we focus on a practical form of drawing commonly known as design sketching or just *sketching*, and specifically on perspective sketching.

Sketching in perspective is a core skill that artists, designers, and engineers should learn in order to depict objects in a manner that closely resembles how reality is seen by the human eye. It tends to be a skill that novices struggle grasping however, and as a result there is much focus on it in the educational literature [125, 127], as well as in curricula that emphasize sketching skills [178, 179]. Our discussions with domain experts with teaching experience also confirmed

[1]Reprinted with permission from: *Recognizing Perspective Accuracy: An Intelligent Interface for Assisting Novices*. Blake Williford, Matthew Runyon, Josh Cherian, Wayne Li, Julie Linsey, and Tracy Hammond. Proceedings of Intelligent User Interfaces Conference 2020

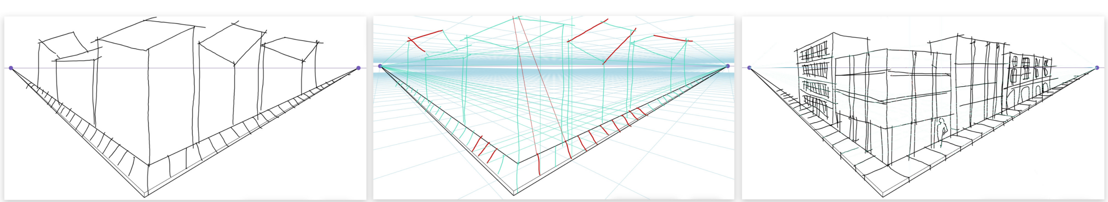

Figure 8.1: Our algorithm and intelligent interface can provide real-time assistance to novices trying to sketch in perspective. In the example above, incorrect convergence is happening in the roofs of the buildings above the horizon, and in the sidewalk lines (both common novice mistakes). Awareness of the mistakes during the sketch can help novices learn perspective and produce more accurate sketches.

that perspective tends to be the skill that beginners struggle with the most.

Our goal with this research was to develop an intelligent user interface that could help novices grasp the concepts rectilinear perspective more readily with their own digital sketches and in real-time. We accomplished this via a perspective accuracy recognition algorithm and technology probes which explore different ways of giving real-time feedback to users. The algorithm can classify perspective strokes versus non-perspective strokes, as well as classify which perspective strokes are accurate and which are not. The algorithm does this on a stroke-by-stroke basis. Our contributions include:

- A general-purpose rectilinear perspective accuracy sketch recognition algorithm with strong performance and novel sketch recognition features.
- The promising results of a user study and exploration of providing intelligent real-time feedback to novices using the algorithm. On average, participants could improve their perspective accuracy by a statistically significant amount by using the intelligent feedback, and this improvement carried over when the feedback was removed.
- Insights about the optimal way to utilize the different forms of feedback for faded scaffolding to best help novices learn perspective sketching skills.

8.1.2 Pedagogical Influence

In an effort to fully understand this problem space, we consulted two sketching experts with more than ten years of experience teaching sketching at the university level. We also referenced some of the leading literature in drawing education [128, 125, 127].

Our discussions with experts revealed that one of the biggest obstacles for novices is learning how to sketch in *perspective*, particularly two-point perspective. In one-point perspective, their is only convergence to one vanishing point, so it is much easier to grasp, however in two-point perspective, there are multiple convergences. This often confuses novices as they have to make decisions about where lines should converge to properly depict the object in perspective.

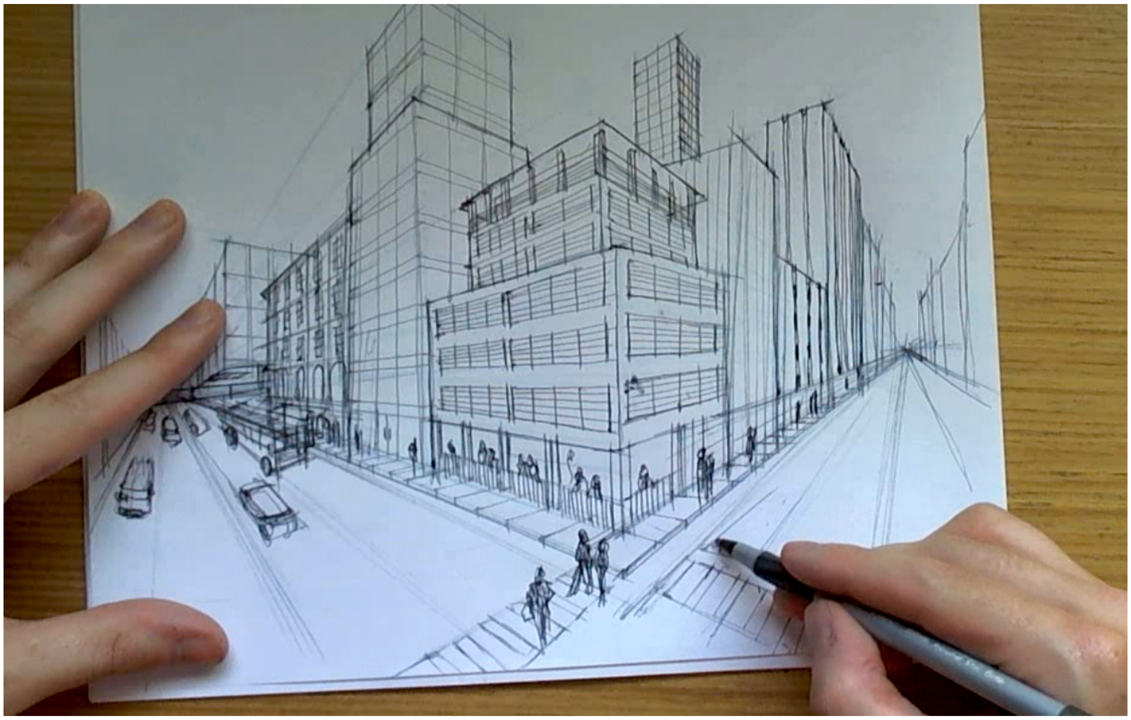

Figure 8.2: A freehand two-point perspective sketch of a city corner generated by an expert during an educational demonstration.

Typically, perspective is taught through demonstrations, where the instructor produces a perspective drawing in front of the students and relays concepts such as vanishing points, the horizon, convergence, and foreshortening. A very common exercise for both understanding and practicing perspective is to think of city streets and corners. One-point perspective is analogous to a city street, in which all lines parallel to the street are converging at the single vanishing point. Two-point perspective is analogous to a city corner, in which the same convergence is happening but for two streets perpendicular to each other. One of the other reasons why city streets and corners are a great way to practice perspective is because buildings are generally *rectilinear* forms, meaning they are cuboid shapes which are far easier to sketch in perspective than curvilinear shapes. Figure 8.2 shows a two-point perspective drawing of a city street.

These insights drive our focus on two-point perspective for both the algorithm, and for the user study, however it's important to note that it is not difficult to adjust the algorithm and intelligent interface to assess one-point perspective, or even three-point perspective.

8.1.3 Algorithm

There are up to four steps in the perspective recognition algorithm: pre-processing, perspective identification, convergence identification, and perspective accuracy.

Pre-processing helps remove noise from the sketch data. A stroke is represented by a series of x, y, and time values for each sampled point from pen down to pen up in a format based on SketchML [180]. Since a user can draw multiple segments in a single stroke, each stroke must be segmented into substrokes. The corners of a stroke are found via an algorithm called ShortStraw [181], and these corners are used to segment the original stroke into its substroke components. After this pre-processing step, we can extract features for every substroke in the sketch.

The extracted features are then used to determine if a substroke should be considered a perspective stroke or a non-perspective stroke. If the substroke is deemed to be perspective, then we determine which point the stroke is converging towards: left (VP1), right (VP2), or vertical. Finally, we determine if the substroke was drawn accurately converging towards the point from the previous step.

8.1.3.1 Features

We chose 26 features that we believed would be highly relevant for the algorithm. In general, these features are in one of two categories: *gestural* features and *contextual* features.

Our gestural features are extracted from the stroke itself and are influenced primarily by Rubine's seminal work in gesture recognition [182]. These features were primarily aimed at classifying whether or not a stroke was a perspective stroke or a non-perspective stroke.

The contextual features are more novel and explore relationships between the stroke and the provided perspective grid. For example, we believed the angle between the stroke's angle and the "perfect" convergence to each vanishing point as well as perfect vertical and horizontal lines would be highly relevant for both convergence classification and accuracy classification. Therefore, most of the contextual features are different ways of measuring that angle including cosine, sin, and degrees. Table 8.1 reveals all of the features we explored and which subset of features were most relevant for at least one stage of the algorithm (in bold).

8.1.3.2 Feature Selection

We performed a best-first forward subset selection to determine the most important set of features for each stage of the algorithm. These features were then used in training the random forest classifier used in this section.

For determining if a stroke was a perspective or non-perspective stroke, the most important features were the length of the bounding box diagonal, angle of bounding box diagonal, distance between first and last point, sum of squared values of angles, line similarity ratio, and substroke siblings (3, 4, 5, 11, 12, 13 in Table 8.1). These features make sense because the sum of squared angle values and line similarity ratio can determine the roundness of a substroke which helps classify non-perspective strokes. In addition, the substroke siblings feature was intended to classify non-perspective strokes that occur from embellishments drawn with a single stroke such as a bush that would result in multiple substroke siblings.

For determining which point the stroke was converging towards if it was perspective, the most important features were the angle of the bounding box diagonal, cosine between first and last point, sine between first and last point, cosine made with vanishing points 1 and 2, and the angle made with vanishing points 1 and 2 (4, 6, 7, 15, 16, 23, 24 in Table 8.1). These features all deal with the direction the stroke is point as a whole as well as the relative difference between an ideal line to each of the respective points that the stroke may converge.

For determining the accuracy of the perspective line, the most important features were the angle of the bounding box diagonal, cosine between first and last point, sine between first and last point, and the cosine, sine, and angle with respect to the vanishing point determined in the previous step (4, 6, 7, 15, 16, 17, 19, 20, 23, 24 in Table 8.1). These features closely correlate to similarity with a perfect line drawn to the vanishing point and thus act as thresholds for accuracy within the classifier.

Table 8.1: Features Used to Classify Perspective Strokes, their Convergence, and their Accuracy

Gestural Feature	Novel?
1. Cos of Initial Angle	No [182]
2. Sin of Initial Angle	No [182]
3. Length of Bounding Box Diagonal	No [182]
4. Angle of Bounding Box Diagonal	No [182]
5. Dist Between First and Last Point	No [182]
6. Cos of Angle Between First and Last Point	No [182]
7. Sin of Angle Between First and Last Point	No [182]
8. Total Stroke Length	No [182]
9. Total Angle Traversed	No [182]
10. Sum of Absolute Values of Angles	No [182]
11. Sum of Squared Values of Angles	No [182]
12. Line Similarity Ratio (#8 / #5)	No [183]
13. # of Substroke Siblings	Yes
Contextual Feature	Novel?
14. Above Horizon?	Yes
15. Cos to VP1	Yes
16. Cos to VP2	Yes
17. Cos to Vertical	Yes
18. Cos to Horizontal	Yes
19. Sin to VP1	Yes
20. Sin to VP2	Yes
21. Sin to Vertical	Yes
22. Sin to Horizontal	Yes
23. Degrees to VP1	Yes
24. Degrees to VP2	Yes
25. Degrees to Vertical	Yes
26. Degrees to Horizontal	Yes

Table 8.2: Classification Performance

Classification	Precision	Recall	F-Score
Perspective Stroke?	0.892	0.892	0.892
Convergence	0.957	0.958	0.957
Convergence Accuracy	0.884	0.885	0.883

8.1.3.3 Performance

In order to evaluate the performance of the algorithm we compiled a set of 22 digital perspective drawings from seven individuals with varying skill levels. These individuals were given the city corner template as shown in Figure 8.4 but with no assistance. The drawings ranged from having very accurate perspective to having very poor perspective and many mistakes. This provided a rich and varied set of data for training.

One of the authors who is a sketching domain expert manually labeled 2026 strokes across the 22 drawings. The drawings were labeled in three different ways:

- Is it a perspective stroke or not? (Binary)
- If so, where is it converging? (VP1, VP2, Vertical)
- Is it accurate or not? (Binary)

After labeling the data we trialed a variety of classifying algorithms and found random forest [184] to perform best. The performance can be seen in Table 8.2.

While this performance could improve, it's important to note that this domain is highly subjective and perfection is not realistic. In particular, the scrutiny with which a perspective stroke can be considered "accurate" does not have a strong answer. A computer can demand a perfectly converging line, however a threshold this demanding would be frustrating to a human sketching freehand. A line directly below a vanishing point could be considered converging to the vanishing point or a vertical line, depending on context.

Since the training data was labeled by a human domain expert, the binary accuracy classification of this algorithm can be considered "good enough" to the human eye. In other words, the algorithm is forgiving if a stroke is reasonably accurate enough. We believe this is very appropriate for freehand sketching, in which some minor inaccuracies are to be expected. We considered also use a continuous floating point value instead of a binary value for perspective accuracy, however that would demand a more arbitrary threshold. The manually trained binary value serves as that threshold.

8.1.3.4 Flexibility

It is important to describe the flexibility of the algorithm and the ways it can be generalized. Because of the way the features are extracted contextually, vanishing points can move anywhere in the 2D euclidean space of the sketching canvas and this will be embedded in the features. This can allow the user to adjust vanishing points in the interface itself with no detriment to algorithm performance. Vanishing points can also be added or removed, allowing for 1-point and 3-point perspective accuracy to be recognized. This flexibility allows the algorithm to be used for almost any rectilinear perspective drawing. The algorithm can also be useful when there are curvilinear forms in perspective such as cylinders, spheres, or ovoids. In these cases, the algorithm can help the individual to establish the rectilinear "scaffolding" for those forms. For example, cylinders are based in a cube or cuboid scaffolding. If the cube or cuboid is in accurate perspective, the cylinder will be in accurate perspective (see Figure 8.3).

8.1.4 User Interface

We designed a web-based interface which can be accessed with any device that has an internet connection. We provide the user with a basic set of drawing tools and some scaffolding in the form of a city corner sidewalk and horizon (see Figure 8.4). This establishes the perspective of the sketch for the user. For purposes of the study, we intentionally do not provide more power to the user such as the ability to adjust the perspective grid or a wide variety of drawing options. These features can be added in the future.

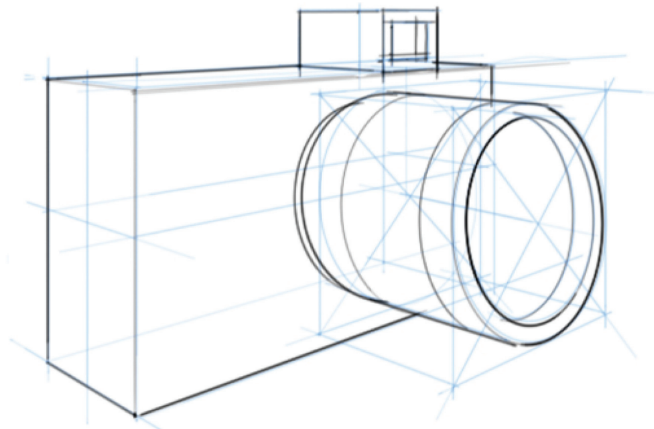

Figure 8.3: An example of a sketch with a curvilinear form (cylinder). The algorithm is still helpful for these forms because they are typically constructed within a rectilinear scaffolding (see blue construction lines). The more accurate the scaffolding, the more accurate the curvilinear form will be.

We provide the horizon at two-thirds the height of the screen of the device, and the vanishing points a short distance from the edges of the screen so that they can be seen by the user. Depending on the device resolution, this can translate to different X and Y coordinates which are captured during feature extraction. While this "forces" the perspective and makes it quite dramatic, we believed this was the best approach for novices. Our discussions with sketching instructors informed us that forcing the perspective this way and keeping the vanishing points visible can be very helpful for novices.

In an effort to answer one of our research questions and understand the best way to provide feedback to the user, we designed three technology probes using this interface; an unassisted probe, a real-time feedback probe, and a summative feedback probe. The details of each are described below.

8.1.4.1 Unassisted Probe

The unassisted probe was designed to assess baseline performance of participants and does not provide any feedback on the accuracy of the perspective. It is little more than a sketching canvas. However, the algorithm is still running in the background, allowing us to gain perspective

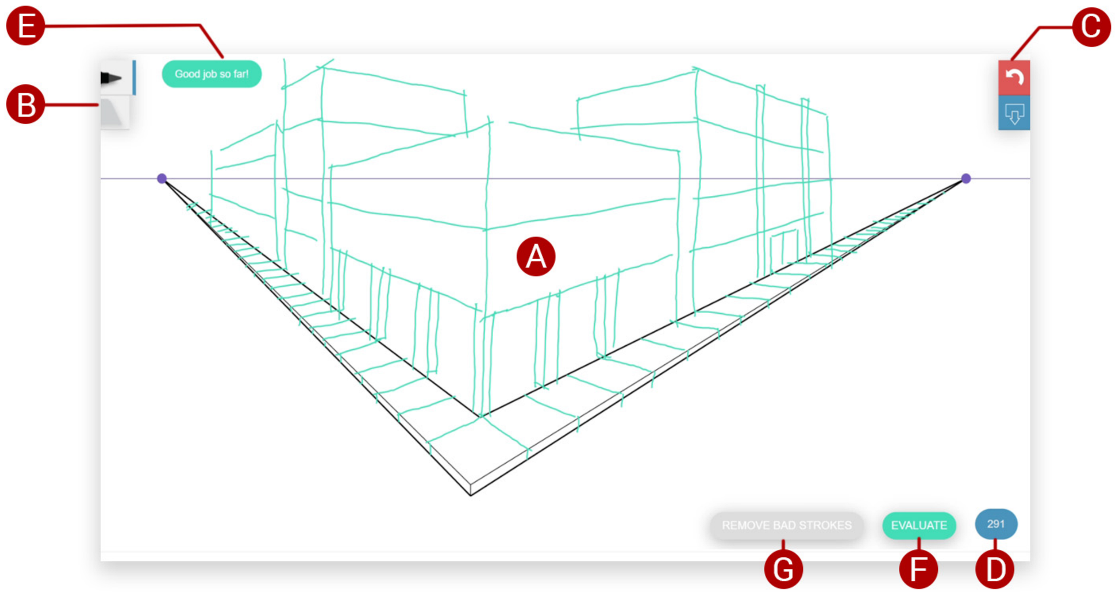

Figure 8.4: (A) The sketching canvas, which provides two vanishing points and a sidewalk. (B) Pen and Eraser. (C) Undo and Save sketch. (D) 5-minute timer. (E) *Real-time probe only*—Text feedback on most recent stroke. (F)(G) *Summative probe only*—An evaluate button and "remove bad strokes" button for removing incorrect strokes.

accuracy data. As will be discussed in the methodology, the control group uses a slight variation of this probe that adds a static perspective grid in the background (without any intelligent feedback).

8.1.4.2 Real-time Feedback Probe

The real-time feedback probe provides feedback after each and every stroke sketched by the user. This comes in the form of turning the stroke green or red for three seconds, along with all of the previously sketched strokes. If a stroke turns red, the user would know that it is an inaccurate perspective stroke, and can either erase or undo that stroke. Additionally, text feedback either provides positive encouragement "Great job so far!" or constructive feedback "This stroke appears to converge to a vanishing point but it is inaccurate." If the stroke is labeled incorrect, the perspective grid also appears for three seconds (see Figure 8.5).

It's important to note that for users who may be red-green color blind, we use a light green color that has noticeable contrast between the dark red color used for incorrect strokes.

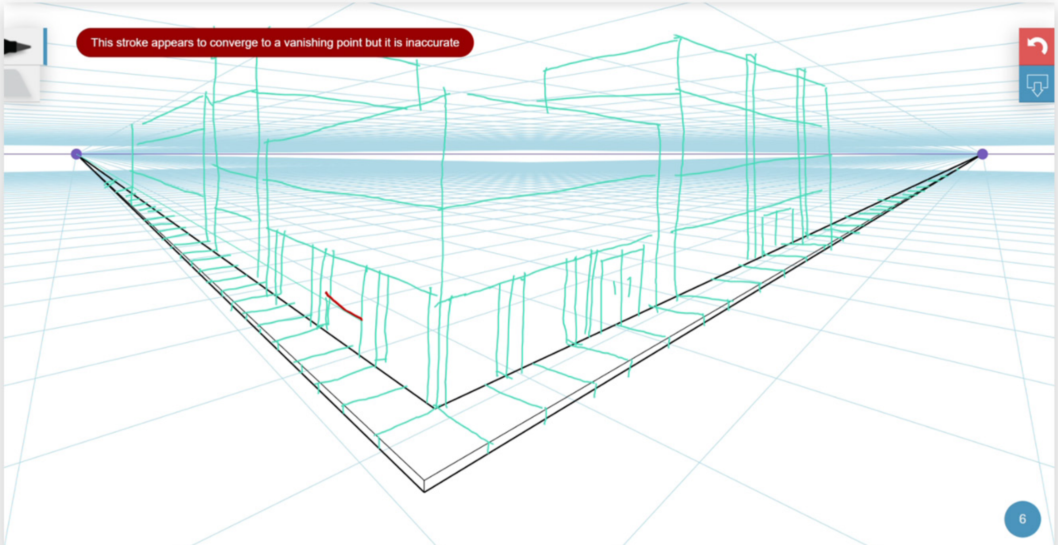

Figure 8.5: The perspective grid is a form of scaffolding in a light blue color commonly associated with construction lines in design sketching. It appears briefly when an incorrect perspective stroke is sketched in the *real-time feedback* probe, and is always present when the evaluation mode is toggled in the *summative feedback* probe.

8.1.4.3 Summative Feedback Probe

The summative feedback probe places the power with the user and gives them the ability to evaluate the whole sketch at any time. For example, the user can sketch for a few minutes, then press the "Evaluate" button to see if any strokes were found to be inaccurate by the system. This variation also gives the user the useful ability to remove all inaccurate strokes in the sketch in one fell swoop. The evaluation mode also provides the perspective grid (see Figure 8.5) and the user can leave it on for as long as they like while they grasp what could be wrong with their sketch.

It's worth noting that while these probes could provide automatic correction to perspective strokes such as what is available in some commercial tools like Sketchbook Pro, we believed there would be more educational value in letting users make mistakes. This is supported by the educational literature [185]. When individuals are free to make their own mistakes, and then are corrected, it makes them more aware of the mistakes. This awareness can translate into learning more readily.

8.1.5 Methodology

We conducted a user study to answer our two main research questions associated with this work:

R1—*Can novices improve their sketching perspective accuracy from exposure to intelligent feedback driven by the algorithm? Can it transfer to subsequent sketches?*

R2—*What are the advantages and disadvantages of real-time versus summative feedback in this context, and which form of feedback is most useful to novices?*

We conducted a between-subjects user study with 40 participants. Participants were undergraduate and graduate students from a variety of majors who were recruited primarily through convenience sampling. The group included 26 males and 14 females. The participants were predominantly novices at sketching with an average self-rated drawing ability of **2.50 (**$\pm$ **0.88)** out of 5, and an average self-rated ability to draw specifically in perspective of **2.17 (**$\pm$ **0.87)** out of 5.

After the initial questions on drawing ability, participants were asked to make four perspective sketches of city street corners. Each sketch was given a time limit of five minutes and participants were encouraged to be creative with the sketches, but draw them in 2-point perspective to the best of their ability. In order to ensure that any improvements in perspective accuracy were truly caused by the intelligent real-time feedback and not just practicing with a grid, half of our participants (20) were in a control condition in which they still produced four drawings, but with just a static perspective grid (see Figure 8.6).

8.1.5.1 Sketch 1 - Pre Test

Initially, all participants were asked to sketch with the unassisted probe, which served as a pre-test and baseline of sketching performance.

8.1.5.2 Sketches 2 + 3

Participants in the experimental group were exposed to the two feedback probes. For each of those participants, we alternated between exposing one feedback probe first versus the other to

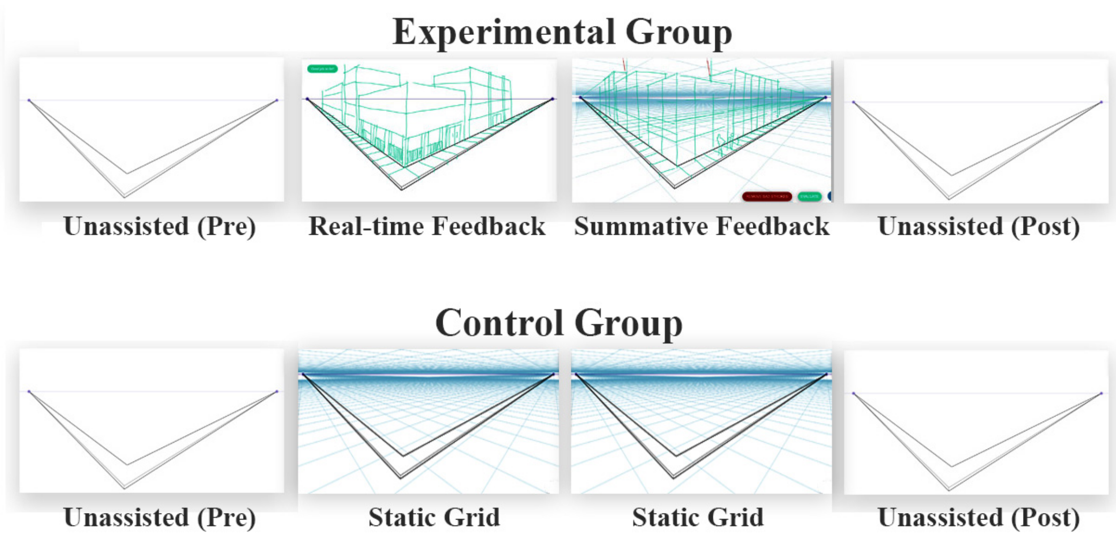

Figure 8.6: The study involved all participants producing four sketches. The first and last sketch were unassisted and served as pre and post tests. The second and third sketches differed for the two groups. The experimental group (n = 20) used the feedback probes while the control group (n = 20) were only assisted with a static perspective grid.

avoid learning effects from one specific probe being more pronounced.

Meanwhile participants in the control group continued to use the unassisted probe, albeit with a static perspective grid in the background added. A perspective grid is something that students can practice drawing in perspective with readily, without using any technology. However, this may not be enough for novices to truly grasp the rules of perspective and we hypothesized that practice with the grid would not lead to any significant improvement in perspective accuracy afterwards.

8.1.5.3 Sketch 4 - Post Test

Lastly, all participants were asked to sketch with the unassisted probe one more time, which served as a post-test.

8.1.5.4 Data Analysis

For each sketch, we logged the number of strokes drawn, perspective strokes classified, accurate perspective strokes classified, and any corrections that were made (including “undo,” “erase,” and the unique “remove bad strokes” for the summative feedback probe). To mitigate imperfection in the algorithm, a sketching expert reviewed the sketches and the logs to ensure the numbers were as accurate as possible.

Between each sketch, participants answered survey questions centered around self-rating their sketch, and self-rating their confidence during the sketch. For the experimental group specifically we asked what they thought about the intelligent feedback provided by the feedback probes. Additionally, think aloud protocol [186] was used to elicit rich qualitative data from those participants about the intelligent feedback during the sketches.

8.1.6 Results

During the course of the study the 40 participants produced 160 perspective sketches. The average sketch in the experimental group contained around 87 strokes (± 34.59) of which an average of 44.45% (± 16.04%) were classified as perspective strokes. Individuals varied quite a bit with regards to how much non-perspective detail they chose to include in their sketches, with some focusing mostly on the main rectilinear buildings in perspective and others choosing to add people, signage, cars, etc. Unless otherwise noted, the following analysis is focused only on the perspective strokes generated by participants and how accurate they are. All p-values reported are from paired t-tests.

8.1.6.1 R1 Improvements in Perspective Accuracy

Not surprisingly, the real-time feedback provoked participants in the experimental group to make nearly twice as many corrections on average, a statistically significant increase ($p < 0.001$). This resulted in sketches that were much more accurate, also a statistically significant increase ($p < 0.0001$). Since the real-time feedback immediately showed participants any perspective error they made, it was easy for them to erase or undo that stroke and make a correct one. 17 of the 20 participants produced 100% accurate perspective sketches with this probe.

The summative feedback resulted in similar statistically significant increases in corrections and accuracy ($p < 0.05$ and $p < 0.0001$, respectively), though slightly less dramatic. Participants used this feedback in varying ways, and they did not always make all corrections suggested by

Table 8.3: Average Performance In Control Group

Condition	Perspective Acc.	Corrections
Unassisted (Pre)	56.91% (± 15.87%)	7.55 (± 11.51)
Static Grid	72.74% (± 17.51%)	8.00 (± 11.93)
Static Grid	63.70% (± 16.69%)	9.50 (± 13.13)
Unassisted (Post)	56.64% (± 15.30%)	11.59 (± 13.38)

Table 8.4: Average Performance In Experimental Group

Condition	Perspective Acc.	Corrections
Unassisted (Pre)	64.43% (± 18.12%)	11.95 (± 11.79)
Real-time	99.27% (± 1.80%)	22.45 (± 6.86)
Summative	93.36% (± 11.31%)	19.15 (± 10.53)
Unassisted (Post)	78.94% (± 11.58%)	17.85 (± 12.58)

the feedback. Since the feedback in this probe was suggesting multiple corrections, they would sometimes forego the "remove bad strokes" feature and only correct certain mistakes. 11 of the 20 participants produced 100% accurate perspective sketches with this probe (see Table 8.4).

What is perhaps most promising is that when the feedback was taken away and the participants returned to the unassisted probe, there was still a statistically significant increase in corrections ($p < 0.05$), and improvement in perspective accuracy ($p < 0.0005$) compared to the initial unassisted sketch. This suggests that there was a mild learning effect from exposure to the intelligent feedback for many participants, with some transfer of knowledge happening in the final unassisted sketch. 16 of the 20 participants in the experimental group improved their perspective accuracy in the final sketch (see Figure 8.9 for two examples). This is supported by the qualitative data which included participants sharing what they thought they learned about perspective.

> *"I understood that basically all vertical strokes need to go straight, perpendicular to the horizon, while horizontal strokes need to be drawn at angles with perspective from your drawing to the horizon. Also you can't see the tops of objects above your*

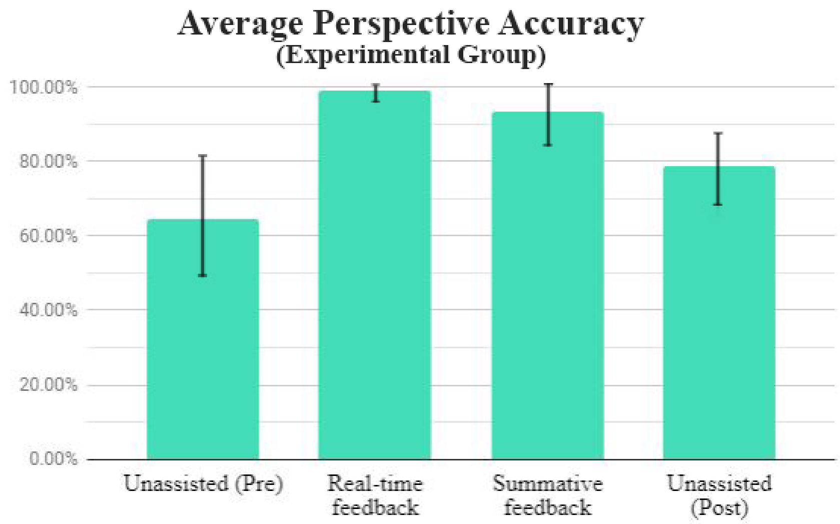

Figure 8.7: Accuracy improved with both feedback conditions, particularly real-time feedback. More importantly, when feedback was removed in the final sketch, participants retained some improvement in perspective accuracy on average ($p < 0.0005$), suggesting a learning effect and transfer of knowledge. Error bars are 1 standard deviation.

perspective to the horizon."—P2

"It really helped in understanding how the perspective changes depending on how far away from the vanishing point I get."—P8

Another result that is very promising is the improvement in self-rated sketching confidence in the experimental group (see Table 8.4). After each sketch participants were asked to rate their sketching confidence during the sketch on a scale from 1 to 5 (5 being very confident). The average rating after the initial unassisted sketch was 2.10 (± 0.79) out of 5, and the average rating after the final unassisted sketch was 3.45 (± 0.51) out of 5, a statistically significant increase ($p < 0.0001$). This suggests the tool could be very useful for improving sketching self-efficacy [18], an important factor when learning sketching and staying motivated to practice.

Meanwhile, participants in the control group showed signs of improvement in their perspective accuracy on average by using the static grid ($p < 0.0005$ for second sketch, statistically significant; $p = 0.072$ for third sketch, not statistically significant). However, they did not show any improvement at all in the subsequent unassisted final sketch ($p = 0.36$). This can be readily seen

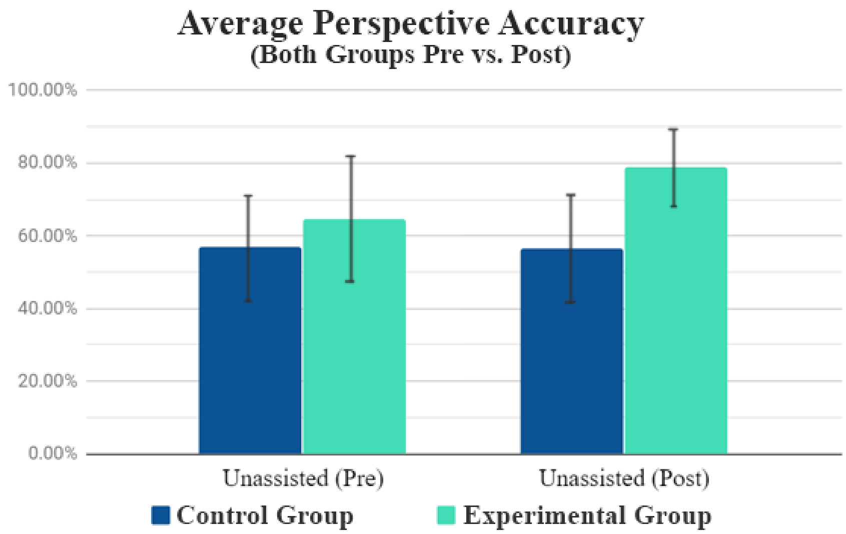

Figure 8.8: When comparing the pre and post sketch performance between the two groups, it can be readily seen that participants in the control group showed no improvement in perspective accuracy on average ($p = 0.36$), suggesting no learning effect. Meanwhile participants who were exposed to the real-time feedback improved by a statistically significant amount ($p < 0.0005$). Error bars are 1 standard deviation.

in Figure 8.8. This suggests that practice with a perspective grid alone does not help individuals fundamentally improve their understanding of perspective, but the intelligent real-time feedback provided by our system can have that effect.

8.1.6.2 R2 Both Forms of Feedback Considered Useful

A content analysis of feedback given from both surveys and think aloud protocol revealed insights on how useful participants in the experimental group found the intelligent feedback provided by the probes.

In general, participants found both approaches to feedback useful, but for different reasons. When asked how useful they found the feedback on a scale from 1 to 5 (5 being very useful), participants rated the real-time feedback 4.25 (± 0.64) on average, and the summative feedback 4.20 (± 0.77) on average. The following sections reveal more nuances of how the participants felt about the different forms of feedback.

8.1.6.3 Advantages and Disadvantages of Real-time Feedback

Many participants appreciated that the real-time feedback allowed them to maintain a constant flow in their sketching, correcting individual mistakes as they progressed.

> *"It's easier to stay on task with what you are drawing. The feedback rate is consistent, you're getting feedback after every step."—P2*

> *"Since the feedback is real-time I am less worried about making mistakes and having to check them later. This improves my workflow significantly."—P14*

However, some participants found that the feedback could be distracting, since it appears after every single stroke. Some got concerned when they were making many mistakes in a row, and not fully grasping why the lines were inaccurate.

> *"The feedback is somewhat distracting while sketching."—P12*

> *"Its easy to get flustered and want to start over completely."—P13*

8.1.6.4 Advantages and Disadvantages of Summative Feedback

Many participants appreciated that the summative feedback gave them the power to choose when to get help.

> *"More 'freedom' to draw without getting immediately notified at if you do something wrong"—P3*

> *"Receiving Feedback when I asked for it allowed me to work at my own pace without having to stop and retrace every line"—P5*

> *"I got to choose when I wanted to receive feedback, meaning I could just check when I wasn't confident in a line"—P7*

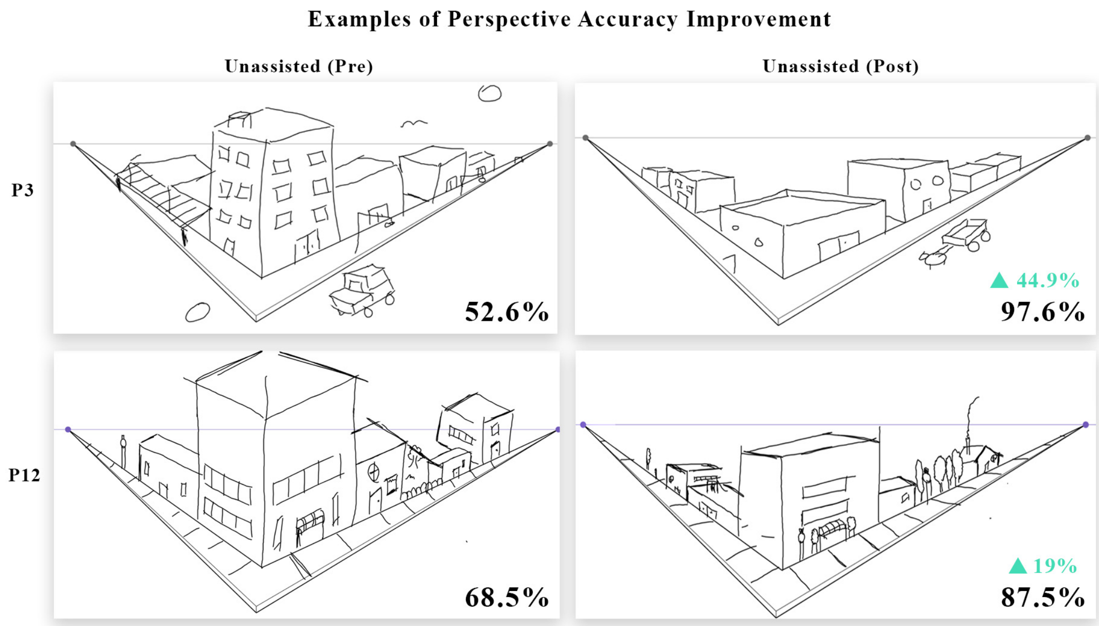

Figure 8.9: These participants improved noticeably from the initial unassisted sketch to the final unassisted sketch. Note that in the initial sketches both participants make mistakes with the roofs of the buildings (you should not be able to see the roof of a building above the horizon). In their final sketches they have properly depicted the buildings and their roofs.

However, some participants did not like the idea of making many mistakes only to later find out that much of their drawing was incorrect. In a few cases, this was an overwhelming amount of mistakes that was difficult to correct all at once.

> *"If you had a lot of bad lines, then the program would suggest removing a large portion of the image. Sometimes causing you to change what you are drawing completely."—P1*

> *"You might end up drawing the whole thing out of perspective and not check until it's too late."—P3*

There was also a consensus among many participants that the real-time feedback was more appropriate for absolute beginners, and the summative feedback was better for when one is more intermediate to advanced.

"This type of feedback would be for those who already had some understanding of drawing in perspective and made minor mistakes. I like being able to draw larger items and then check how 'good' it was at one time."—P1

"This type of feedback seems good for intermediate level."—P9

8.1.7 Discussion and Future Work

We gained insights from this work which are worth discussing and sharing with other researchers and technologists working in the areas of creativity-support, intelligent tutoring systems, and sketch recognition.

8.1.7.1 Fading Scaffolding and Changing Feedback Based on Skill Level

Faded scaffolding [158] is a technique used in educational contexts to ease beginners in to a subject area while keeping them in their proximal zone of development [187]. It involves giving a beginner extra support until they improve their ability and confidence, at which point the support can begin to fade away. Adding, then removing training wheels on a bicycle is a classic example of faded scaffolding.

We believe the insights gained from the user study reveal an ideal way to utilize faded scaffolding in the context of helping novices learn perspective sketching with intelligent software. Initially, a persistent perspective grid can help novices to essentially trace perspective gridlines, and make accurate perspective sketches immediately. This would be used in combination with real-time feedback for every stroke to catch any mistakes from the individual immediately. Then, the persistent perspective grid can be removed, and only shown when mistakes are made. When the individual has reached a more intermediate to advanced level, the real-time feedback can be replaced with summative feedback, bringing them the power to receive help only when desired, and giving them more freedom to sketch without scaffolding. Lastly, when the individual has shown mastery of perspective, all forms of feedback can be removed and the individual can have full freedom to sketch anything they desire on any medium (digital or traditional).

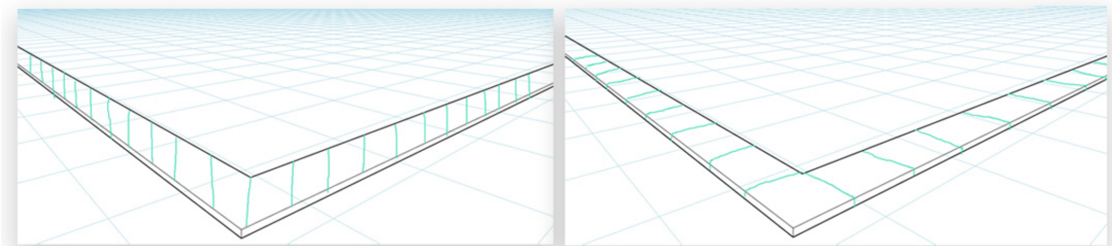

Figure 8.10: (A) shows a common novice mistake which is to miss convergence in sidewalks. The algorithm assumes these are "correct" vertical strokes because it has no context of knowing that it is meant to depict a sidewalk. (B) shows the correct convergence.

8.1.7.2 Algorithm Limitations

It's important to emphasize that this algorithm is limited to rectilinear perspective accuracy recognition. Recognizing perspective accuracy of *curvilinear* forms is significantly more complex, and can be subjective even for human sketching experts. Such recognition would likely have to involve more robust and nuanced perspective stroke classification, since the current algorithm does not classify curves as being perspective strokes. It would also need many more contextual features which describe the relationship between strokes in the sketch and what they are trying to depict. The systems built by Xie et al. in automatically completing sketches and paintings [188, 189] use contextual features and could be a source of inspired improvement.

Edge cases can also detriment the potential educational benefit of such an algorithm. For example, a common novice mistake was to draw sidewalk lines as a series of vertical strokes, even though they should converge to the vanishing points (see Figure 8.10). More contextual features could help with these edge cases and it may be possible to subdivide provided templates into sections where specific perspective convergences are expected.

8.1.7.3 Supplementing Existing Pedagogy

For reasons mentioned in the previous section, we believe this user interface would be best used in combination with traditional pedagogy. In a traditional studio environment where sketching is typically taught, a qualified instructor may perform demonstrations, show videos, and thoroughly explain the concepts behind perspective. An interface such as what we've presented in this work can allow the students to practice, ensuring they at least grasp the concept of rectilinear two-point perspective before moving on to other advanced topics. The real-time feedback provided by the tool could allow students to self-regulate their learning [190].

Additionally, the instructor could offload the task of grading to such a tool, significantly reducing the time needed for grading. They can then observe the performance of individual students, and better personalize their instruction for the students that need it most.

9. IMPROVING DRAWING CREATIVITY

9.1 Ambiguous Stimuli: Exploring Creativity Support for Form Ideation

9.1.1 Motivation

Ambiguity, by definition, is the quality of being open to more than one interpretation. Since multiple interpretations by a single individual or across several individuals can lead to multiple solutions [191], ambiguity has long been considered a useful tool in the creative process for designers, artists, and other creative professionals. Renowned concept artist and art educator Scott Robertson regularly uses and encourages ambiguous underlay drawings in his approach to concept art [125]. Gaver et al. support ambiguity as be a wonderful resource for design [192].

Sketching is a domain where ambiguity can be utilized readily. In addition to facilitating one's internal thought processes [1] and problem-solving ability [6], sketching can be considered a "dialectic" thinking exercise [193]. Sketches form a dialogue between the brain and the external artifact that is constantly evolving and changing. When the artifact is ambiguous, it holds potential for many different interpretations and new directions.

In creative professions like concept art, animation, and game design, creativity and original ideas are highly valued and sought after. Our curiosity in this research centered around the role technology can play in promoting ambiguity and thus creativity in the creative process while sketching. We primarily wanted to understand the effects of *ambiguous stimuli* on novice and intermediate creators who would likely benefit the most from assistance from such technology.

We recruited 20 university students from a Visualization major at ****, most of whom are aspiring concept artists, animators, and game designers, and conducted a study that involved sketching twelve alien creature concepts. For half of the concepts the students sketched in a normal manner on a digital tablet, and for the other half they were exposed to a technology probe that generates ambiguous stimuli underlays based on the user's strokes.

We make the following contributions in this work:

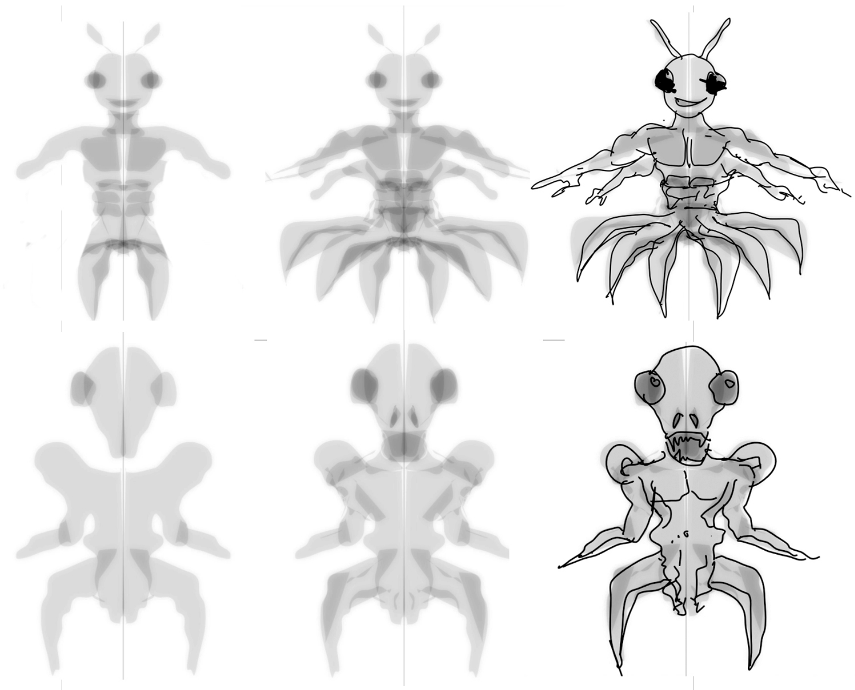

Figure 9.1: Examples of ambiguous stimuli co-created by the user and the system and the resulting creature concepts after further development.

- A discovery that using ambiguous stimuli in the beginning of the ideation process resulted in less fixation on humanoid forms and provoked more unusual forms. Other effects included cultivating holistic thinking, stimulating exploration, enjoyment, resistance to a loss of control and freedom, and a consensus that the stimuli are most useful in the earliest stages of the creative process.
- Implications for further research in utilizing ambiguous stimuli to augment human creativity and provoke more exploration, along with design implications for related co-creative systems and creativity support tools.

9.1.2 Technology Probes

We built two web-based drawing technology probes, each with a basic set of drawing tools and limited options. We were not focused on providing the users with power and features; rather, we wanted to ensure they could quickly sketch ideas and be focused on those ideas rather than the user

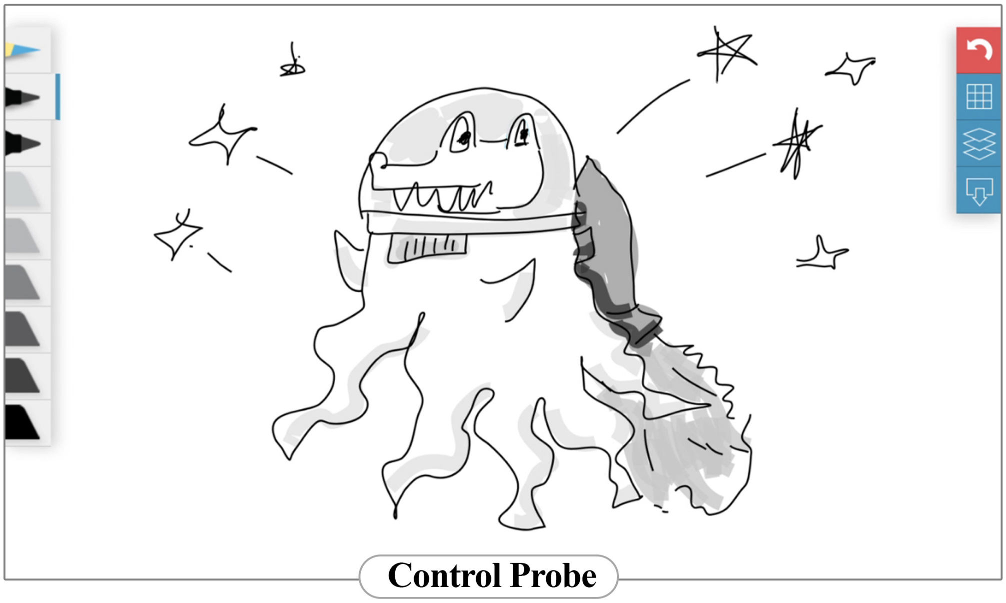

Figure 9.2: The control probe is a basic sketching interface with options of three stroke thicknesses (including a "blue pencil" for underlays) and 6 gray value "markers" for adding value to a sketch. The probe was intentionally designed to be simple and mirror the basic tools commonly used by concept artists in analog sketching.

interface (see Figure 9.2).

One of the tools has the additional capability of allowing the user to co-create ambiguous stimuli with the system (see Figure 9.3). It does this by converting the user's stroke input into shapes of semi-transparent light gray value and mirroring those shapes. The probe allows this allows for a simple series of strokes to yield fairly complex, layered, and symmetrical ambiguous stimuli which can then serve as an underlay for normal line work. While symmetry could be optional in a more feature-rich co-creative system, we did not make turning off symmetry possible for the study. The symmetry can lead to unexpected and stimulating forms which we considered an important aspect of our exploration on ambiguous stimuli.

9.1.3 Methodology

We conducted a study primarily aimed at answering our research question: *How do these ambiguous stimuli affect an individual's creativity and form ideation?* We chose a within-subjects experiment design where all participants used the ambiguous stimuli probe as well as a basic

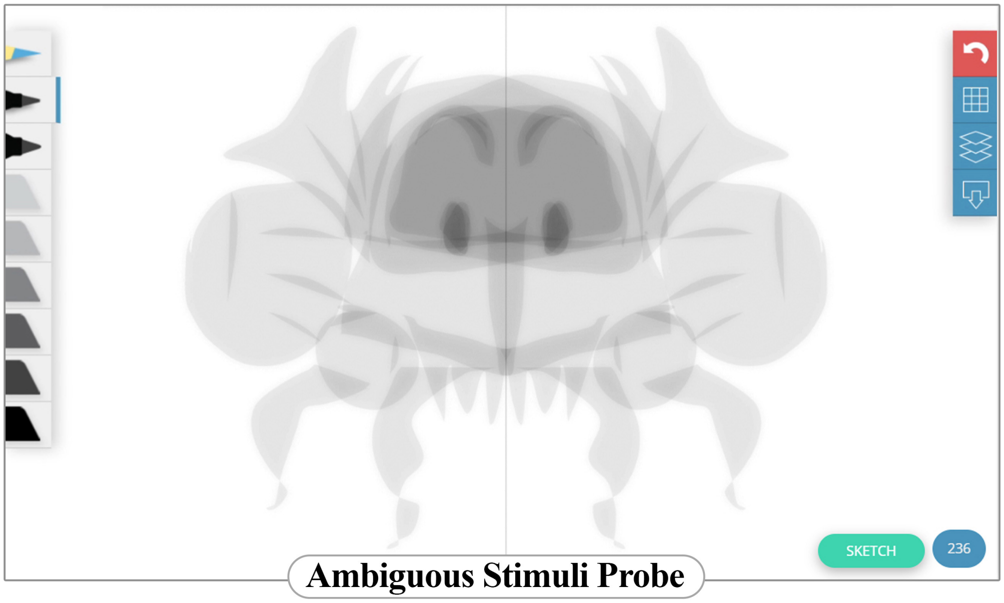

Figure 9.3: The ambiguous stimuli probe adds an additional capability of producing ambiguous stimuli underlays. For purposes of our study design, the system switches to a normal sketching mode (line work) automatically after two minutes.

sketching interface as a control condition. For each participant, we alternated between which condition they were exposed to first, as we were also interested in learning if ambiguous stimuli could have lingering effects on their ideation process.

9.1.3.1 Recruitment

We recruited 20 university students for the study, most of whom were pursuing a degree in Visualization. These students are working towards careers in concept art, animation, game design, and other fields related to entertainment. We chose from this pool because we wanted participants who had a stake in creativity and who could potentially benefit the most from creativity support for form ideation. We expected more rich sketch data and qualitative data from these participants.

The participants included 10 males and 10 females at the undergraduate and graduate level ranging from 19 to 32 years old. In general, participants ranged from novice to intermediate with respect to their confidence in their drawing ability and creativity. Their average self-rating for drawing ability was 3.55 ($\pm$ 1) out of 5 while their average self-rating for creativity with respect

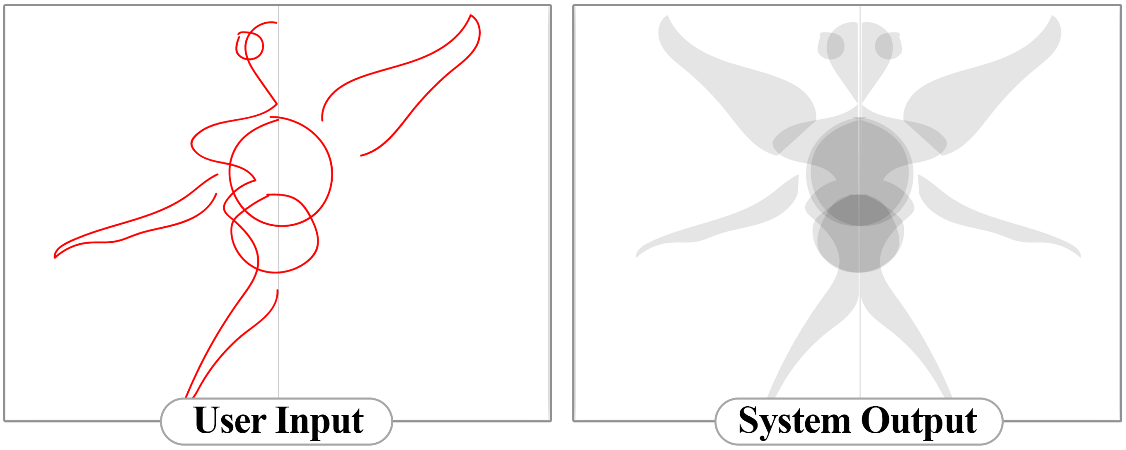

Figure 9.4: Example of how the ambiguous stimuli probe converts user input strokes (left) to an abstract underlay. Each stroke is turned in to a closed shape via the shortest path possible, then duplicated and mirrored. Each shape was black with 10% opacity, allowing for layering to occur, and more nuanced and detailed forms to emerge.

to drawing was 3.85 (± 0.59) out of 5.

9.1.3.2 Protocol

After answering some demographic questions and self-assessing their ability, participants were prompted with the following task:

An animation studio has recruited you to design an alien creature which will star in their upcoming animated film. Your task is to generate 12 early-stage concepts for this creature.

Each participant was asked to draw 12 creatures, each having a time limit of four minutes. Drawing inspiration from Shah's *C-Sketch* [191] study, we chose a four minute time limit because it gave students enough time to fully flesh out a concept. For six of the creatures, the participants used the control probe (basic sketching interface), and for the remaining six they used the ambiguous stimuli probe (experimental condition). For the ambiguous stimuli probe, the four minutes was split between two minutes of creating the ambiguous stimuli underlay and two minutes of normal sketching. This was timed automatically by the system.

We alternated between participants being exposed to the control condition first or the experimental condition first as we were interested in observing if that had any lingering effects on their

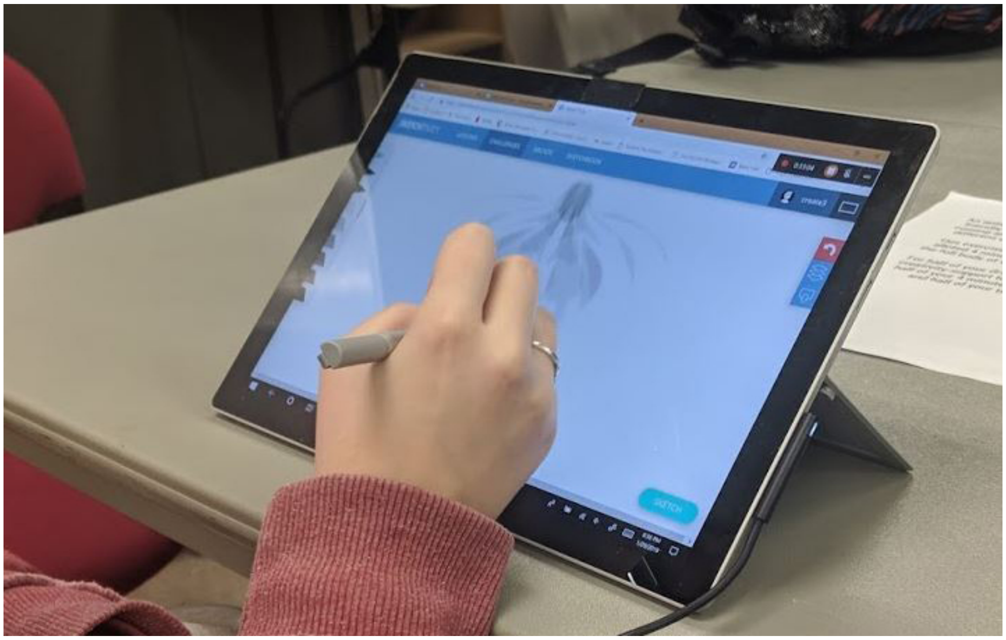

Figure 9.5: A participant sketching a creature using the ambiguous stimuli probe on a Surface Pro device. Participants were encouraged to sketch in a naturalistic manner and orient the device in a way that was comfortable for them.

creativity or fixation. Participants were encouraged to think aloud [186] and discuss their creative process as they sketched, including narratives around their creatures, yielding rich qualitative data.

9.1.3.3 Analysis

For the sketch data, we saved static PNG images of each creature sketch as well as time-stamped vector stroke data in JSON format to a MySQL database. A total of 240 sketches were collected. We analyzed the sketch data by randomizing the images and coding them visually [194, 195] based on themes and patterns seen in the images. We also took inspiration from previous creativity research in which sketches were coded visually, such as Ward's studies on structured imagination [196]. Our visual analysis resulted in six types of creatures, four "novelty factors", and the emergence of four key findings.

We used grounded theory [166] and the constant comparative method [168] to analyze the qualitative data. In the initial coding phase, we analyzed 362 codes from transcripts and notes gained from the studies. After focused coding and theoretical coding, six categories emerged

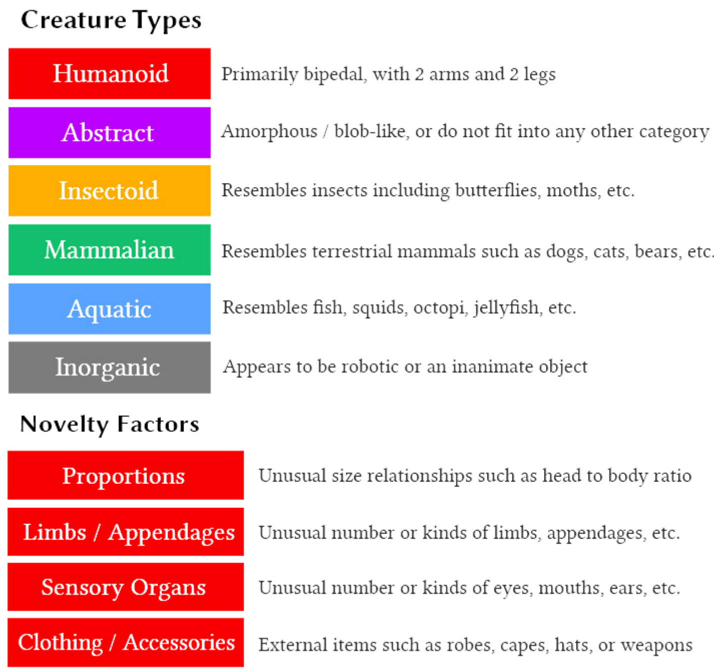

Figure 9.6: Categories that emerged from the visual analysis.

which are described in the results. Theoretical saturation was reached when data from the last several participants did not add any new categories or dimensions to the existing categories.

9.1.4 Visual Analysis Results

Our visual analysis resulted in the categories described in Figure 9.6. Note that creature types alone was not enough to sufficiently measure the creativity or novelty of the creature concepts, so novelty factors was an additional metric we introduced to quantify unique features of the creatures that the participants explored.

9.1.4.1 Overall Fixation on Humanoid Forms

43.8% of all creatures sketched were categorized as humanoid, suggesting a fixation on humanoid forms across most participants. This is consistent with similar prior research and is very reminiscent of Ward's "structured imagination" work in which most creatures had human-like traits like limbs and sensory organs [196]. Participants may have been attempting to anthropomorphize their creatures and make them relatable due to the prompt. This can be observed in Figures 9.7 and

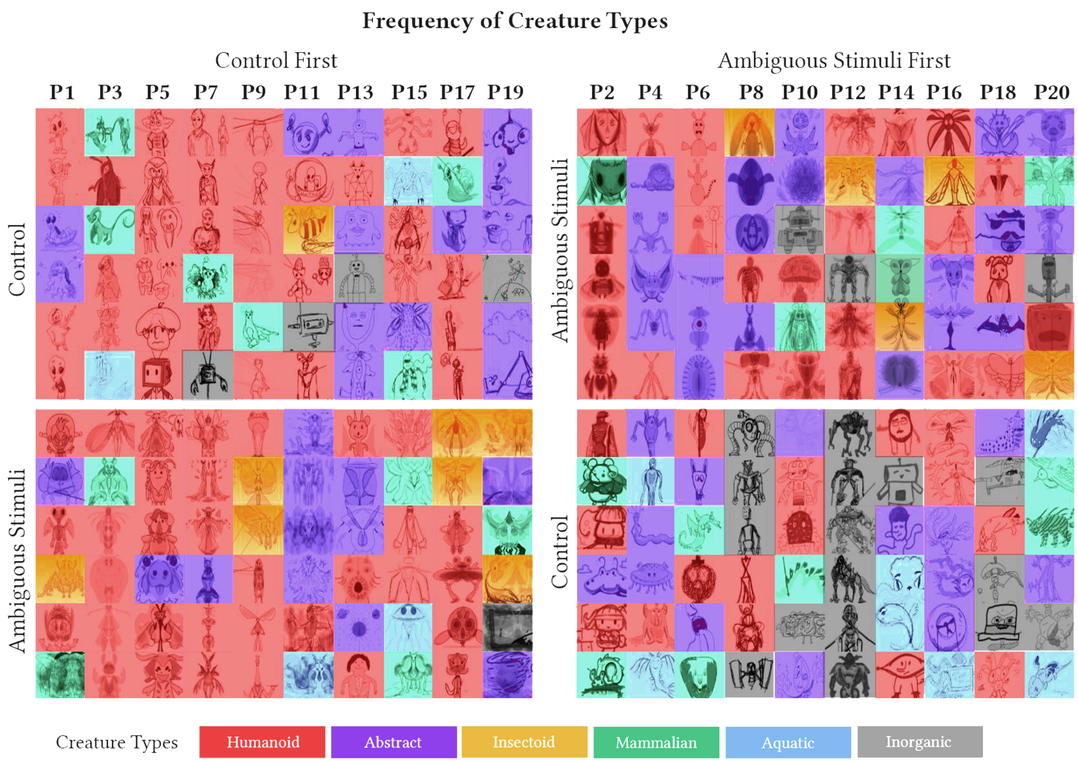

Figure 9.7: A heatmap visualization of creature types across all participants and both conditions. Note that participants who were exposed to the ambiguous stimuli condition first produced significantly less humanoid forms and produced more abstract and inorganic forms.

9.8.

9.1.4.2 Ambiguous Stimuli First Caused Less Overall Fixation on Humanoid Forms

The participants who were exposed to the ambiguous stimuli first had had 41% less humanoid forms, 42% more abstract forms, and nearly four times as many inorganic forms, e.g., robots and inanimate objects (see Figure 9.9). We acknowledge that one participant P12 was clearly fixated on robots. A Pearson's chi-squared test revealed that the effect of beginning with the ambiguous stimuli first was statistically significant ($p < 0.01$) in influencing different creature varieties.

9.1.4.3 Ambiguous Stimuli Provoked Particular Forms More Than Others

When isolating the sketches from the ambiguous stimuli condition, we found that particular forms were provoked. For example, there was significantly more (12 to 1) insectoid forms. Mean-

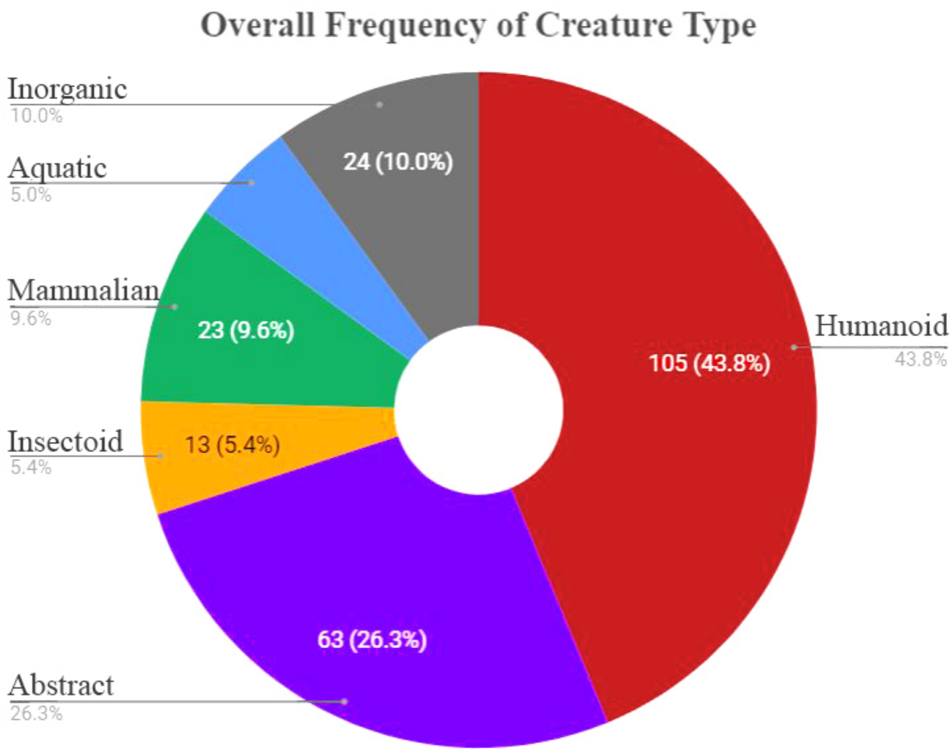

Figure 9.8: The overall frequency of creature types independently of condition. Most participants were fixated on humanoid forms.

while, the control condition resulted in much more inorganic forms such as robots and inanimate objects, as well as more aquatic forms like octopi and squids. This finding is very consistent with the qualitative data and we will touch on it again in that section. A Pearson's chi-squared test revealed that there was statistical significance ($p < 0.01$) in the effect of ambiguous stimuli on creature varieties.

9.1.4.4 Ambiguous Stimuli Don't Impede Exploration of Novelties

The stimuli resulted in the exploration of slightly more novelties including proportions, limbs, and sensory organs, but less of clothing and accessories. A Pearson's chi-squared test revealed that there was statistical significance ($p < 0.02$) in the effect of ambiguous stimuli on exploring more novel features.

We recognize that the coding of novelties is highly subjective, so the key finding we are

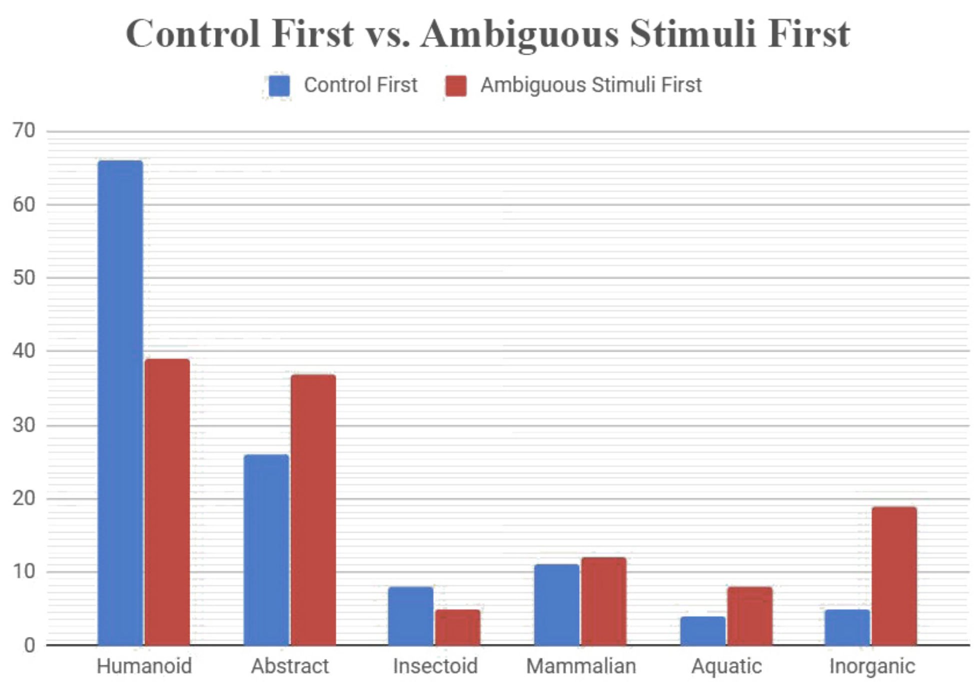

Figure 9.9: Participants who were exposed to the ambiguous stimuli condition first produced significantly less humanoid forms and instead produced more abstract or inorganic forms.

claiming is that ambiguous stimuli *do not impede* an individual's creativity or inclination for exploring novel features. The average participant explored around two novelties (1.98 ± 0.82) per creature concept regardless of condition.

9.1.5 Qualitative Analysis Results

9.1.5.1 Holistic Thinking Cultivated

The ambiguous stimuli seemed to cultivate many participants' focus on the "big-picture" and overall form rather than being caught up in details. This was likely due to the combination of symmetry and value allowing for a full form to emerge very quickly. Some participants expressed that they normally begin by drawing silhouettes, suggesting that the ambiguous stimuli meshed well with their creative process.

> *"I do like how it made me think in generalities"—P9*

> *"It was forcing me to figure out the shape of the character before the details"—P4*

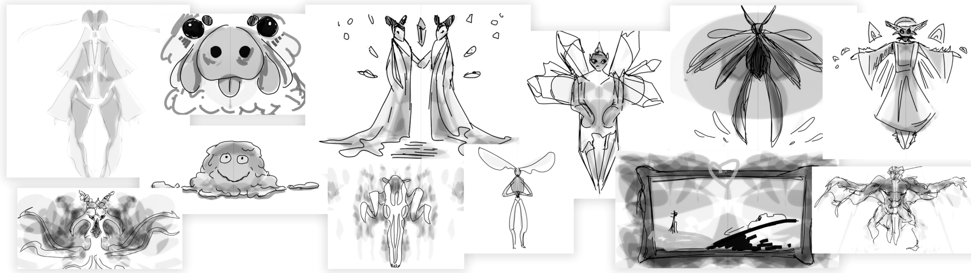

Figure 9.10: Examples of creature sketches generated by participants with the ambiguous stimuli probe.

> *“You can get stuck in the details so it prevents you from that”—P3*

9.1.5.2 Exploration Was Stimulated

Many of the participants found that the stimuli cultivated more exploration and led to less fixation on what they normally are accustomed to drawing. Some participants felt that it loosened them up and expressed that it caused them to approach the sketching with less of a preconceived idea.

This is consistent with the finding that the participants who were exposed to the ambiguous stimuli first had less fixation on humanoid forms despite the overall trend of humanoid forms across both conditions. It is also consistent with the finding that the ambiguous stimuli condition resulted in different forms and slightly more exploration of novelties.

> *“It’s freeing me up a bit”—P3*

> *“The really expressive shapes keep your mind going. It makes you think of something else”—P4*

> *“After you get used to using it, it gives you an interesting way to think of ideas”—P10*

9.1.5.3 Interesting and Enjoyable

Most of the participants found it interesting and enjoyable to generate the ambiguous stimuli. Most of them had never used symmetry tools or sketched in such a way before. Some found it relaxing while others found it humorous or pleasantly surprising.

"It's weird, but not in a bad way"—P9

"It feels very satisfying. It's super fun to make creature on this"—P2

"This is cool, I've never tried drawing like this before"—P7

9.1.5.4 Particular Forms Provoked

Many of the participants expressed that the stimuli pushed them towards certain specific forms. The most common themes were "insects" and "bugs," including "butterflies" and "moths." Some participants also mentioned clothing including "robes" and "tuxedos." A few participants expressed that the stimuli encouraged "organic" forms as opposed to more angular and inorganic forms. This was not seen as a negative by most participants, although some did feel constrained by it.

This is very consistent with the finding from our visual analysis discussed earlier. The ambiguous stimuli did result in significantly more insectoid forms and significantly less inorganic forms such as robots and inanimate objects. This may have been influenced by the particular way we designed the ambiguous stimuli probe.

"It makes me think of insect shapes"—P3

"It's definitely easier to think of organic forms"—P5

"I felt constrained by bugs"—P9

9.1.5.5 Loss of Control and Freedom

Some participants were initially jarred by the loss of control in their creative process. The ambiguous stimuli forced symmetry and areas of values which could sometimes be surprising and not what was originally intended. Many expressed that they prefer asymmetry when sketching for a variety of reasons.

"This is not the typical way I draw"—P11

"I definitely feel like the symmetry is constraining me"—P20

"Sometimes you need asymmetry for a character to stand out"—P1

9.1.5.6 Benefit to Creative Process

Still, most participants acknowledged they could see the benefit of giving up some power and freedom for what the ambiguous stimuli allow for. Participants acknowledged their usefulness in speed, forcing reconsideration, and adding beneficial "chaos" to the creative process. There was also a nearly universal consensus from participants that the ambiguous stimuli would be useful in the earliest stages of ideation but not in any other part of their creative process. This includes "thumbnailing" and the initial effort of generating a variety of ideas from which to later narrow down the best ones. This is consistent with our intention for ambiguous stimuli.

> *"This would help with production speed. I feel like I can make a lot of silhouettes this way"—P2*

> *"It kind of forces chaos in to your design work flow"—P12*

> *"I'd use it only for thumbnailing and early concepts for characters"—P9*

9.1.6 Discussion

We believe the results show preliminary evidence that ambiguous stimuli can offer a variety of benefits for creatives in their ideation process, however we acknowledge the limitations of the study. This work opens up new avenues of research to further investigate and holds implications for how co-creative systems can be designed to augment creativity.

9.1.6.1 Limitations

We acknowledge that this study focused only on university students within one major at one university. The study is also focused on one particular subject (creatures). The claims made on the effects of creativity are therefore preliminary, and should be validated and expanded upon in further research. It would be interesting to investigate other subjects (such as vehicles or environments), other participants (professional concept artists and illustrators, novices with low creative self-efficacy), and other technology probes that can facilitate ambiguous stimuli co-creation.

9.1.6.2 Furthering Creativity Research

The results showed that being exposed to the ambiguous stimuli first resulted in less fixation on humanoid forms in subsequent sketches, however the participants were overall still largely fixated on these forms. This is consistent with previous creativity research which suggests that humans rely on categorical knowledge when ideating and that creativity is little different than normal cognition [197, 196].

An interesting further inquiry is how can these stimuli further push people towards less expected results and help them overcome the limitations of "structured imagination." Adding more chaos such as no symmetry, randomization, altering user strokes more drastically, and giving the system more influence could do so. If the current system can provoke particular forms such as insects, an adjustment to the system could provoke other forms. This could be parametric and controlled by the user. More influence from the system could add constraints, however it has been repeatedly shown that specificity and constraints can benefit creativity [198, 199].

9.1.6.3 Balance Between User, System, and Control

These results lead to an interesting question: How much control should the user and the system each have when generating ambiguous stimuli? Some participants resisted the forced symmetry aspect of the ambiguous stimuli. Allowing them the flexibility to turn this feature off and on is simple to implement and would give them more control and freedom, however it would push the system more towards a standard sketching program. Adding adjustable parameters may give users more flexibility and power and make the ambiguous stimuli more useful and provocative. We purposely avoided making the system in this study too sophisticated and feature-rich, as we wanted to isolate the ambiguous stimuli as a variable.

Perhaps allowing the system to automatically adapt to a users sketching style or preferences may be a promising future direction [200]. These "adaptable interfaces" are a promising way of using machine learning to automatically personalize the experience to the user. Building on the "enactive" model of creativity from Davis et al. [83], the interactions should be designed "like a conversation where each party tries to make sense of contributions and respond appropriately given the history of interaction." As artificial intelligence will allow for increasingly sophisticated models of users, this should theoretically allow for that "conversation" to be richer, pushing ideas farther and making the co-creative process more fluid.

9.1.6.4 Benefits at Various Skill Levels

We believe ambiguous stimuli can be helpful to individuals with a variety of skill levels. We chose to conduct the study with students in creative majors because we sought feedback from people who had a stake in creativity and would benefit the most from co-creative tools. The participant pool was more intermediate in terms of drawing ability on average. However, we believe there is also potential for ambiguous stimuli to help absolute novices overcome a phenomenon described by many drawing instructors and their students as "Fear of the Blank Page" [198]. By quickly gaining momentum through generating ambiguous stimuli, novices could potentially improve their creative self-efficacy [201].

Meanwhile, creative professionals with advanced skill level may also benefit from ambiguous stimuli. Under pressure and time constraints to generate novel concepts in industry jobs, professional concept artists and illustrators demand tools that can stimulate their creativity and provoke them to think of their next big idea.

10. IMPROVING THE STUDIO EXPERIENCE[1]

10.1 Exploring the Potential of an Intelligent Tutoring System for Sketching Fundamentals

This chapter contains previously published work.

10.1.1 Motivation

Sketching is a powerful skill that can be used to explore ideas, communicate ideas, and work through problems. As a form of drawing, it has been shown to benefit student's internal thought processes [1], boost self-confidence [7], promote general academic achievement and problem-solving [6], improve peripheral skills such as writing and brainstorming [4], and even improve spatial recognition skills [8].

However, sketching can be a difficult skill to learn that requires many hours of practice. Sketching is often taught in dedicated "studio" courses at the high school and university level [136]. Sometimes it is also integrated as a component in engineering courses, with several weeks dedicated to it [178]. Students often struggle with low self-efficacy [18] with respect to sketching, which in turn effects their motivation to practice [19]. The inevitable comparisons to others which is part of studio environments doesn't help, as students can easily get discouraged by more skilled peers [20].

Sketching teachers have their own struggles as class sizes can become quite large, reducing their ability to personalize their instruction for individual students who need feedback and making grading overwhelming [17]. This reduces their overall efficacy and can lead to poor learning outcomes. There is a critical need to reconsider how we are teaching sketching to students. Since sketching is a skill that requires many hours of practice, students need to be engaged, confident, and motivated to practice. Teachers need to be effective and able to provide personalized feedback.

Intelligent Tutoring Systems (ITS) [29] are an interesting approach to educating students with-

[1]Reprinted with permission from: *Exploring the Potential of an Intelligent Tutoring System for Sketching Fundamentals*. Blake Williford, Matthew Runyon, Wayne Li, Julie Linsey, and Tracy Hammond. Proceedings of the annual ACM Conference on Computer-Human Interaction (CHI) 2020

out human instructors. They are designed to replicate the benefits of one-to-one personalized tutoring in contexts where it is either difficult or impossible. While challenging to evaluate [30], they have been shown to be effective in various domains and in certain cases even more effective than human instructors [31, 32].

We built an ITS for sketching fundamentals called *Sketchtivity*. We designed the system not to replace human instructors but rather to augment the educational experience of learning sketching with technology. The system includes lessons on the basic "primitives" of design sketching including forms like ellipses, planes, cubes, and cylinders. To convey these concepts the system provides real-time feedback to students driven by sketch recognition, a form of artificial intelligence focused on recognizing hand-drawn sketches.

A previous study showed that the majority of students who used the system throughout the semester could improve their accuracy and speed with regards to sketching primitives [72]. However, very little qualitative data was gained on the actual experiences of students and teachers. This motivated us to dive deeper in to their subjective experiences in order to understand how the software was affecting them. We deployed the system in to six existing courses at the high school and university level during the 2017-2018 school year. We then conducted a grounded theory [166] study which involved six semi-structured interviews with the teachers who implemented the software, as well as a focus group with six students from a course where the tool was used extensively. Our primary research questions in this study were:

R1—*What are the advantages and disadvantages of the system for teachers?*

R2—*What are the advantages and disadvantages of the system for students?*

Ten categories emerged (five for teachers and five for students) which are described below:

Teachers

Offloads grading work

Validates insights about students

Cultivates exploration and community

Easy to deploy

Limited access to quality devices

Students

Real-time feedback insightful but limited

Gameplay motivates in different ways

Considered a strong warm-up tool

Digital practice seen as an advantage

Not engaged by creative challenges

These findings describe the potential of an ITS for sketching fundamentals to improve the educational experience for both teachers and students.

It's worthwhile to discuss the design of the system itself to provide better context to the results and analysis. First, we discuss the pedagogical influence on the system, including the challenges we discovered from interviewing two experienced sketching instructors. We then detail the various features and components of the system which we believe accomplishes the goal of providing students a strong learning experience independent of their instructor.

10.1.2 Pedagogical Basis

The system is heavily influenced by design sketching pedagogy [125, 127, 202]. Design sketching, or conceptual sketching, is a form of drawing that is commonly taught to industrial designers, concept artists, architects, and engineers. It is heavily focused on rapid ideation and the development of concepts rather than producing finished works of art. There is a large focus on perspective sketching and sketching in 3-dimensions as well. The way in which industrial designers are taught is heavily influenced by mastery learning [118] such that beginners ideally master basics of 2-dimensional and 3-dimensional primitives like ellipses, cubes, cylinders, before moving on to more advanced topics like combining primitives and sketching actual objects.

In addition to referencing the leading educational literature [125, 127, 202], we interviewed two domain experts whom we will label as E1 and E2 for anonymization purposes. E1 and E2 both have more than ten years of experience teaching sketching at the university level. We interviewed them to find out challenges they faced when teaching sketching. Our main research questions centered around how our system could help with these challenges and most benefit their instruction practices. A content analysis revealed the following principles which we used as a guide for designing the ITS:

10.1.2.1 Every Student is Different

E1 emphasized that every student is different and has different motivations, and for this reason there is no "one-size fits all" approach. He mentioned flexibility as an important factor for the system to have, with different gameplay motivators for different students. As an example, E2 utilized friendly competition in his classes and mentioned that it could be a great motivator for many students. He also mentioned that personal style was important. He emphasized that software should not discourage students from developing their own personal style and approach to sketching.

10.1.2.2 Time is Limited

A strong theme from the interviews that is also supported by relevant literature [17], is that the time instructors have for teaching can be very limited. From this standpoint, it is difficult to personalize instruction to each student and ensure each student is advancing at a reasonable pace. This tends to result in a stratification of student skill levels in which many students lag behind others. A negative externality of this stratification is that the students who lag behind tend to get discouraged and have lower drawing self-efficacy than their more advanced peers.

> *"I really wish I had the time to give every student more feedback”—E1*

10.1.2.3 Providing Personalized Feedback is Difficult

Also related to the aforementioned time constraints, providing personalized feedback to each student can be difficult. Instructors may have classes as large as 60 or more students and may or may not have teaching assistants to assist with grading. While some real-time feedback can be given to a small set of students during class time, much of the class is reserved for lecturing and conveying new concepts. The remaining feedback has to come in the form of grading assignments which can take weeks to complete. These weeks without feedback are missed opportunities for the students to correct themselves, cultivate good habits, and improve their overall ability quickly.

10.1.2.4 Encouraging Practice is Difficult

Both instructors emphasized the crucial importance of practice and doing many hours of it. E1 mentioned that he would often assign very large quantities of sketches for practice knowing that many students would not even complete the quantity assigned.

> *"Putting in the mileage and practice is key. It doesn't happen overnight"—E1*

10.1.3 System Features

The following sections describe various features of the ITS which were designed to support both instructors and students and their various needs.

10.1.3.1 Lessons

Sketchtivity's lessons follow a mastery learning approach [118] that closely mirrors how perspective sketching is taught traditionally in industrial design curricula. Users can begin with the basics including lines and basic 2D primitive shapes like circles and squares. They can then move on to learning perspective and 3D primitives such as cubes, cylinders, and spheres. Each lesson allows the user to practice eight examples of the given shape (see Figure 10.1) and provides feedback on the accuracy, line quality (smoothness), and speed of the sketches submitted (see Figure 10.2). Some of the lessons, such as the cube lesson, employ faded scaffolding [158] to keep users in their proximal zone of development [187] and support them as they learn the fundamentals of sketching.

10.1.3.2 Creative Challenges

The creative challenges give users an opportunity to flex their creativity in exercises that are more open-ended. Templates are provided which serve as scaffolding and a starting point from which they can sketch anything from their imagination. 2-dimensional exercises include subjects like creatures, cars, spaceships, while 3-dimensional exercises include city streets (see Figure 10.3), a very common exercise when learning perspective. In the creative challenges, users are

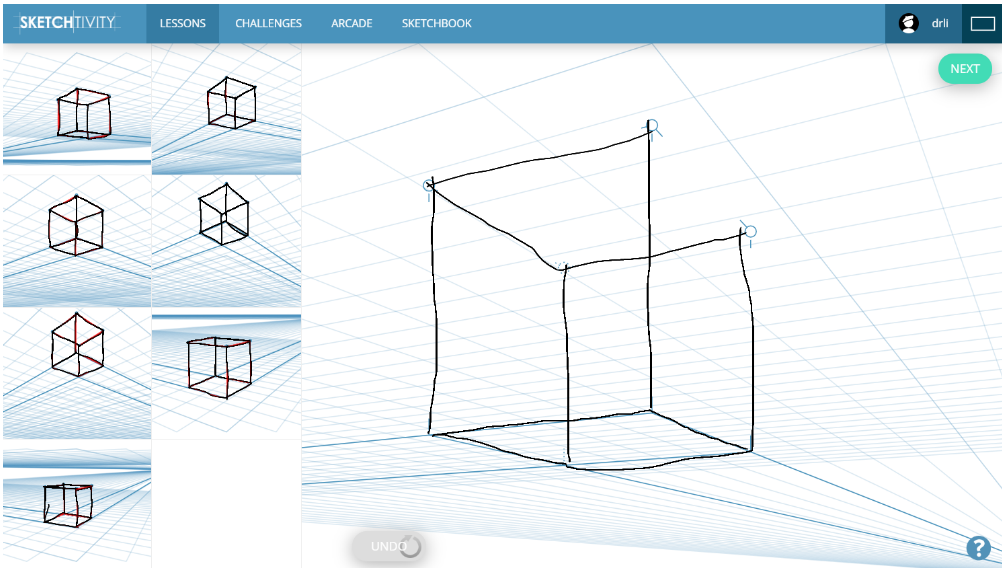

Figure 10.1: The interactive lesson for cubes. Students can practice cubes in various perspective orientations and receive real-time feedback on accuracy.

given more power and flexibility with regards to stroke width, and value in the form of gray "markers".

10.1.3.3 Arcade

The arcade contains a sketch-based game called *ZenSketch* [142] which allows users to practice basic line work (see Figure 10.4). The game utilizes the concept of "stealth learning" [123] to make practicing line work more fun and feel like less of a chore. The feedback given is similar to the line lesson, however *ZenSketch* also includes gameplay elements such as a high score and leaderboard, which are designed to motivate students who enjoy competition.

10.1.3.4 Profile

The profile gives users an opportunity to view progress and stats regarding their sketching skills (Figure 10.5). This includes "levels" for accuracy, smoothness, and speed based on their performance in the lessons. It also includes *ZenSketch* high scores and stats on how many of each basic shape have been practiced. The profile is designed to give students a broad sense of their progress and improvement and help them self-regulate their learning [203].

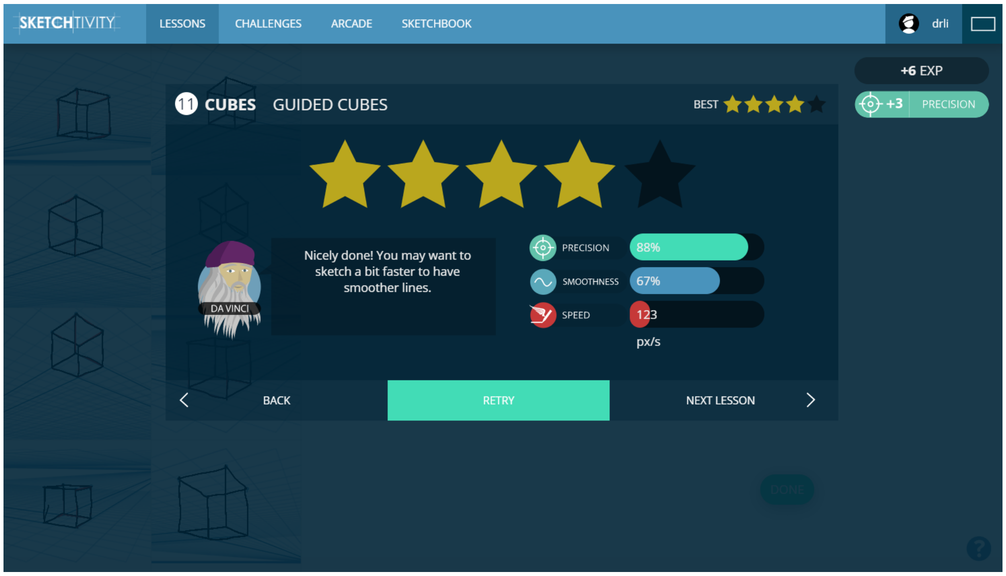

Figure 10.2: Summative feedback after completing a lesson. Students receive feedback on average accuracy, line quality (smoothness), and speed, along with text feedback on what they can work on to achieve the highest score.

The profile also provides teachers the ability to view student engagement— specifically the number of each basic form students have practiced (Figure 10.6). This can be helpful to gauge engagement with the system and for grading purposes.

10.1.4 Methodology

Sketchtivity was deployed in three different courses at three different Texas high schools, and at Georgia Institute of Technology in two different courses (one of the courses had two sections). Figure 10.8 shows the instructors we interviewed and details on which course they taught and what semesters they implemented the software.

It's important to emphasize that teachers were given flexibility and were encouraged to use the software in a way that was appropriate for their course. This encouraged naturalistic and realistic use of the software. For example, T2 covered sketching for three weeks in her *Introduction to Engineering Design* course, so she had her students use the software during just those three weeks. Meanwhile T5 and T6 were teaching a course entirely focused on sketching, so they asked their students to use the software every week for the entire school year.

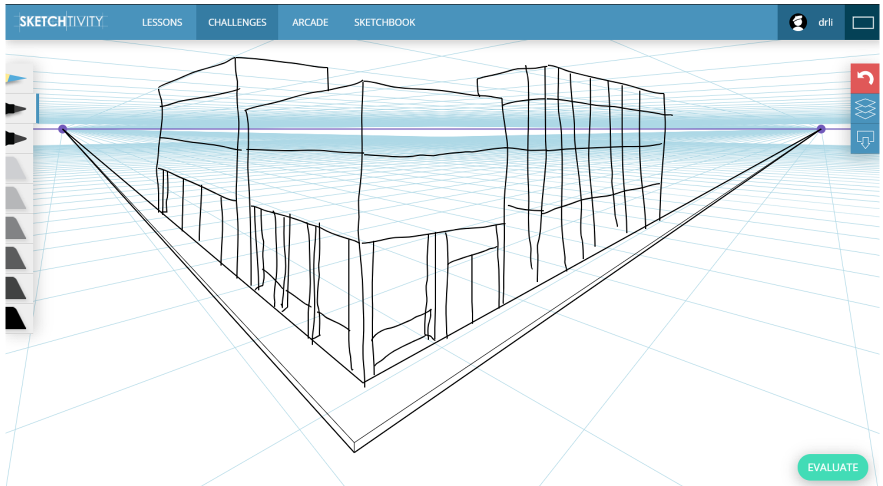

Figure 10.3: An example of one of the creative challenges the system offers. The city street allows a student to practice perspective and demonstrate understanding of 1-point and 2-point perspective.

Upon completion of their deployment, semi-structured interviews were conducted with each of the teachers to gain data on their experiences using the tool in their class. The interviews lasted about 30 minutes on average. Interview questions focused on how they were using the system in their classroom and their perspective on its advantages and disadvantages.

To understand the student perspective, we conducted semi-structured interviews with students from the university *Introduction to Sketching* course (section taught by T5 and T6) who used the system extensively. On average, students in these sections practiced over 1000 sketches of basic primitives during the school year. We did not interview high school students because they were a more difficult to access population (under 18), and because they did not use the system as extensively as the university students (see Figure 10.9).

Grounded theory [166] was used to analyze the data, and the constant comparative method [168] was utilized to improve our interview questions and dive deeper into the experiences of the teachers and their students. In alignment with Charmaz, we viewed the data from a social constructivist [204] viewpoint.

A strong degree of theoretical saturation was reached when the final instructor and student interviews did not add any new codes or categories. Ultimately, ten categories emerged from the

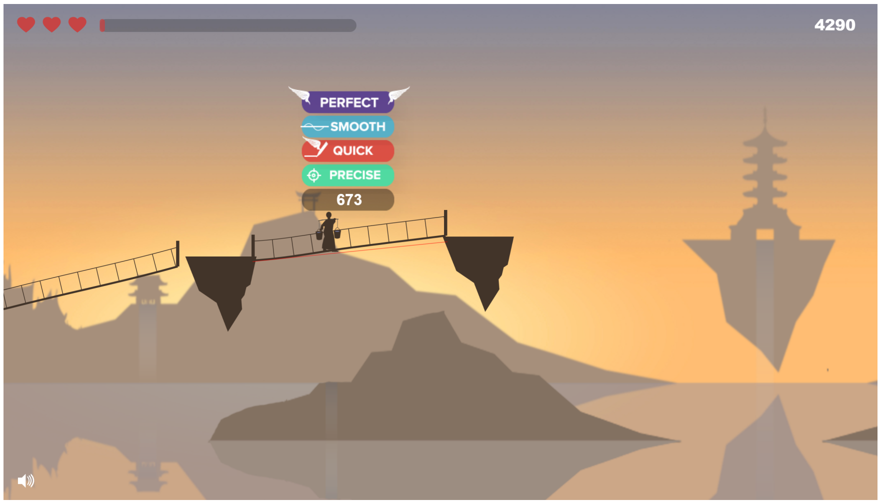

Figure 10.4: ZenSketch, a game for practicing basic line work.

data, including five categories from teacher experiences, and five from student experiences. These are discussed in more detail in the following section.

10.1.5 Results (Teacher Experiences)

10.1.5.1 Offloads Grading Work

One of the things that instructors spoke most enthusiastically about was the advantage of the platform in allowing them to offload grading work, making grading much easier and faster. This was especially useful for T5 and T6, who deployed the system for the entire school year and made using it part of a participation grade. The instructor profile allowed them to quickly and easily view individual student engagement at the end of the semester and grade accordingly. T5 especially appreciated that he didn't have to look through overwhelming quantities of basic forms and grade them.

> *"Yeah, so the profile feature is very useful to know how much time students are spending on it."—T4*

> *"Data is instantaneous and I don't have to do anything. From a time standpoint there's a pain factor for analyzing 500 sheets of ellipses. But hey it can do that, there's no reason for me as a professor go through 500 sheets of that. That's really menial."—T5*

Figure 10.5: The profile view for students, showing "levels" for accuracy, smoothness, and speed based on their performance in the lessons.

10.1.5.2 Validates Insights about Students

Some of the teachers found that the performance data validated insights they had about students. They believed it could help them personalize their instruction.

> *"It kind of validates some things that I've already learned about the students. So one was general effort. It correlates back to attempts. There'd be a group of students who are more or less in the middle. If you just look at say total number of attempts, you'll see that it kind of tiers out. The middle section, better section, the best, the slightly worse, the worst. You have kids who didn't even bother. They were either being neglectful or they just totally forgot. But I liked that there was this correlation between the amount of effort they put in to Sketchtivity and the grade they would get in other assignments."—T6*

10.1.5.3 Cultivates Exploration and Community

The teachers mentioned how their students explored aspects of the platform and discovered its features organically. In this way the learning experience did not feel prescriptive but rather new and exciting. The high school teachers (T1, T2, T3) in particular appreciated this aspect.

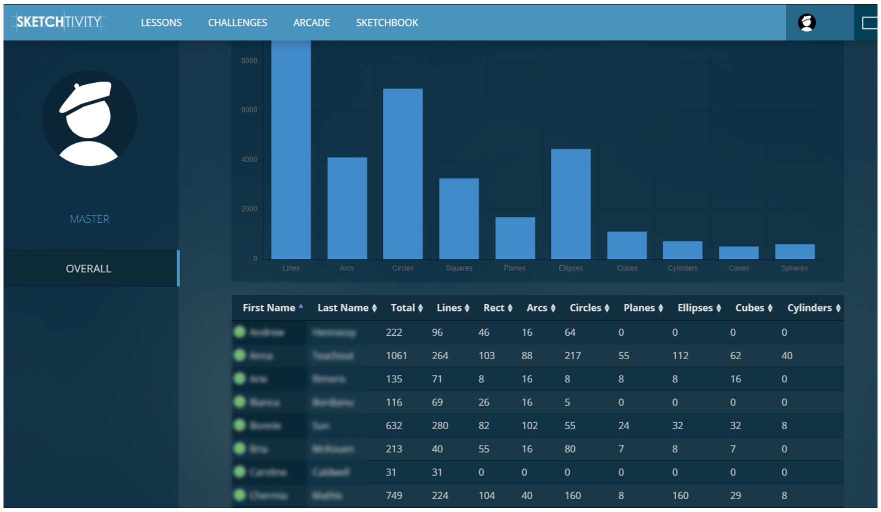

Figure 10.6: The profile view for instructors, showing total lessons completed by the class along with more detailed information on each student.

Many students also shared their scores and performance with their peers, cultivating friendly competition and a community around sketching practice. T5 appreciated this, as he had found in previous years of teaching sketching that community really helps to engage students and keep them motivated to practice.

> *"They did some exploration, and got through the basic front stuff. And then they discovered the game, and the game became the all-consuming thing."—T1*

> *"Some students got in to it because they saw their section as a team. I saw also some of them thinking it was a competition of their section against another section. So I did notice hearing conversations between students over the course of the semester. 'Hey let's draw together.' 'Hey did you get your accuracy up?' So four or five times over the semester I noticed this kind of thing."—T5*

10.1.5.4 Easy to Deploy

In general, the teachers found *Sketchtivity* easy to implement into their curriculum. Teachers were provided a course code that allows students to register themselves. The system was accessible via the web and the teachers appreciate how simple it was to get up and running. Because of the

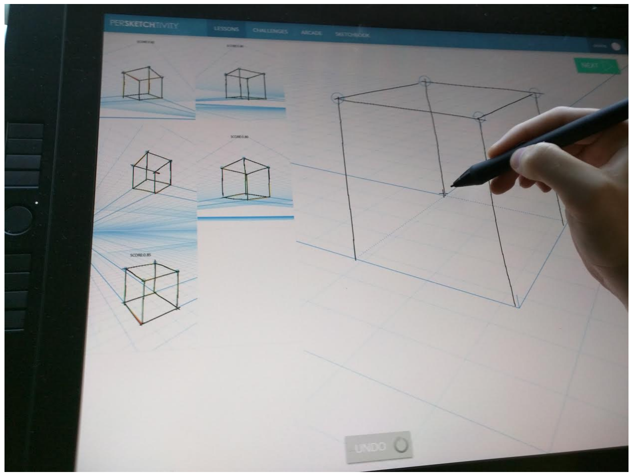

Figure 10.7: A university student using the system to practice cubes in perspective on a Wacom Cintiq device.

Teacher	Level	Course	Semester	Students	Sketches
T1	High School	*Principles of Engineering*	Fall 2017	55	3,879
T2	High School	*Introduction to Engineering Design*	Fall 2017	16	2,692
T3	High School	*Fundamentals of Art*	Fall 2017	79	24,405
T4	University	*Introduction to Engineering Design*	Fall 2017	53	5,572
T5	University	*Introduction to Sketching*	Fall 2017, Spring 2018	34	40,381
T6	University	*Introduction to Sketching*	Fall 2017, Spring 2018	31	39,884

Figure 10.8: Teachers interviewed and details on which course they taught and what semesters they implemented the software. Sketches is the total number of primitive shapes practiced (ellipses, cubes, etc.) in the course. Note the substantial amount of sketches practiced from the Introduction to Sketching course sections (T5 and T6). This course ran the entire school year and was entirely focused on sketching.

naturalistic deployments of the system, teachers could use the system as little or as much as they wanted in their curriculum.

> *"Signing up was easy."—T2*

10.1.5.5 Limited Access to Quality Devices

In all courses there was not a 1:1 ratio of quality sketching devices to student. Georgia Tech teachers mentioned use of "drawing labs" with Wacom Cintiqs, where up to 20 students could use

Student	Gender	Level	Course	Semester	Sketches	ZenSketch plays
S1	M	University	*Introduction to Sketching*	Fall 2017, Spring 2018	1283	11
S2	F	University	*Introduction to Sketching*	Fall 2017, Spring 2018	1807	24
S3	M	University	*Introduction to Sketching*	Fall 2017, Spring 2018	2060	6
S4	M	University	*Introduction to Sketching*	Fall 2017, Spring 2018	805	10
S5	F	University	*Introduction to Sketching*	Fall 2017, Spring 2018	3997	0
S6	M	University	*Introduction to Sketching*	Fall 2017, Spring 2018	2678	6
S7	F	University	*Introduction to Sketching*	Fall 2017, Spring 2018	371	58
S8	F	University	*Introduction to Sketching*	Fall 2017, Spring 2018	346	29
S9	F	University	*Introduction to Sketching*	Fall 2017, Spring 2018	2596	14

Figure 10.9: Students interviewed along with how many sketches they practiced and how many times they played ZenSketch.

the software at any time. High school teachers used a limited supply of chromebooks or iPads. In many cases, students were either left out completely or had to take turns using the devices.

> *"Section was split, because we only had 20 monitors. Half the class got to use the software. That half went to the industrial design department to use those devices."—T4*

> *"In general, they have to be highly reliable systems. They have to work whenever students feel like using them. Students get frustrated by buggy systems, and students throw up their hands."—T4*

10.1.6 Results (Student Experiences)

10.1.6.1 Real-time Feedback Insightful but Limited

Students appreciated the real-time feedback provided by the system on accuracy, smoothness, and speed but expressed that it was limited and desired more nuanced feedback. They also recognized that the feedback should be taken somewhat lightly. For example, students liked that the system could measure the speed of their sketches, but knew that speed wasn't necessarily the most important factor in sketching.

"I like seeing the error."—S2

"I like how it records experience and progress."—S1

"It does help, but then the sketches that are better aren't necessarily faster."—S4

10.1.6.2 Gameplay Motivates in Different Ways

Students felt motivated to improve their score by both the gamified lessons and the game *ZenSketch*. However, some students who identified as not being competitive were more motivated by the gamified lessons. This speaks to individual differences between students and the value of offering multiple ways of motivating students through sketch-based gameplay. S1, S4, and S7 all self-identified as being competitive and were very motivated by *ZenSketch* due to its high score feature.

"I know [S4] and I are super competitive. We like being able to look at the smoothness rating, I think it's fun."—S1

"I like it because I can give it a few. You can keep trying to be as precise and smooth as possible, and I can just do it."—S4

"There was one particular score at the top. They had a top score and I was in 2nd place or something. I was wanting to try to beat it. I could never beat it!"—S7

Meanwhile S5 admitted that she was not very competitive and was more drawn to the lessons where she could compete with herself and only improve her own score.

"I'm not [competitive]. I'm not very good at it so I just do the lessons because the game doesn't really help."—S5

10.1.6.3 Considered a Strong Warm-up Tool

Many of the students thought that the system was great for "warming-up" since it provides a way of practicing all the basic primitives.

It's worth noting that even professionals warm up with basic primitives and it is encouraged in the educational literature [125, 127]. Warming up allows professionals to loosen up their sketching arm and prepare for a sketching session.

> *"I could see it being a really great warm-up activity. Even in your professional career. Just crank out a bunch of lines. Even if you had an endless mode on it where you just crank out a bunch of straight lines in a row. Or a bunch of circles in a row."—S1*

10.1.6.4 Digital Practice Seen as an Advantage

Students appreciated that the platform allowed them to practice without as many construction lines and that they didn't have to go through lots of paper. This is consistent with previous studies, where students expressed that they liked this aspect of digital practice and the scaffolding the system provided [72].

> *"The part I like best is that you don't have to set up a drawing template. You don't have to make the lines and then draw ellipses inside of them. It just sets it up for you. You'd spend half the time just drawing a template on the paper."—S6*

> *"Yeah it's better than worrying about having paper and bringing paper around."—S1*

10.1.6.5 Not Engaged by Creative Challenges

Students were not engaged by the creative challenges. They expressed that if they wanted to do a full drawing they would rather use a different tool (e.g., Sketchbook Pro, Procreate) or traditional pen and paper. This is understandable considering the limitations of a web-based sketching tool to replicate the look and feel of sketches that professional commercial software can provide.

They also expressed that the challenges were too open-ended. In creativity research specificity has been shown to aid creative process [197].

"I have no motivation to do anything that isn't a lesson or a game"—S1

"I feel like I know what I need to practice, and it's more than just ellipses. I don't want to spend 10 minutes on one sketch"—S2

"I feel like actual drawings is for your sketchbook where you can have the feel of pen and paper"—S4

10.1.7 Discussion

This study was grounded in the actual experiences of students and teachers. It showed us that *Sketchtivity* is providing the kind of experience it was designed for and, despite some limitations, can provide many benefits for both instructors and students. The following sections summarize our findings and discuss future research directions such as ways to better quantify student motivation, confidence, sketching ability, etc. which can lead to an improved learner model [35].

10.1.7.1 Limitations of the Study

We acknowledge that there was some variance among the classes that deployed the tool. This was the intention, as we wanted instructors to implement the software in a way that was *naturalistic* and supported their preexisting curriculum and instruction practices. In this way, it allowed us as researchers to better understand the viability of a tool like *Sketchtivity* to help improve the educational experience of learning sketching in these various settings.

Future studies could dive deeper in the nuances of how the tool is effecting student confidence and motivation. However, to do so, we believe the studies need to be longitudinal and potentially last years making it a difficult undertaking. Sketching skills can take years for noticeable improvements to be made, and the changes in motivation and confidence are very nuanced and subjective. For measuring motivation, scales such as the MSLQ can be used, however they are not focused specifically on sketching [169]. For measuring confidence, scales such as Bandura's self-efficacy scale [18], or a self-efficacy scale designed specifically for sketching [107] could be utilized. We did not use scales in this study because the ITS was not a fully isolated variable. All of the teachers utilized a mix of the ITS and traditional sketching practice in their curricula, so any improvements in confidence and motivation could be attributed to just practice in general, not using the ITS specifically.

10.1.7.2 Benefits of an ITS for Sketching Fundamentals

The teachers primarily appreciated the ITS for its capability to automatically “grade” student sketches and track engagement, while also fostering a community around sketching practice. Meanwhile the students found it to be a great warm-up tool and more motivating way to practice.

10.1.7.3 Automatic Grading

The teachers were most enthusiastic about the automatic feedback the system offers. This is particularly important in dedicated drawing or sketching courses where several weeks may be devoted to fundamentals and primitive forms. Potentially hundreds of these basic forms must be graded or at least checked by the teacher and or teaching assistants.

10.1.7.4 Fostering Community

Although the current system does not offer a message board, private messaging, or other ways for students to interact with each other remotely, the teachers found that it still fostered a community around sketching practice in the classroom itself. They noted overhearing students sharing their scores and accomplishments while using the system and that it motivated some “friendly competition” among some of the students. We believe this is another critical benefit an ITS for sketching fundamentals can offer. In the future more community features can be implemented such that students can interact with each other remotely, paving the way towards MOOCs (Massive Open Online Courses) that can utilize the ITS. Such features could also allow for the more advanced students to better teach and motivate each other.

10.1.7.5 Warming Up

Even though many students remarked that they would rather use other tools for more complex and time-consuming sketches, they loved “warming up” with the ITS. Indeed, this may be one of the core benefits for students as sketching primitives is an important part of sketching in much the same way stretching is an important part of playing sports. In both cases, the relevant muscles are engaged, preparing the body for action. In the case of sketching, warming up can help the individual more readily access the "flow" state [205] that is highly associated with deep focus and creativity. An interesting avenue of future research would be to explore the effects of warming up with the system on subsequent sketches in terms of quality and creativity.

10.1.7.6 Motivating Practice

It was clear from the interviews that both competitive and non-competitive students could find benefit in the tool, since the non-competitive students could avoid competitive features and "leaderboards" and just compete with themselves for their own best scores. We believe this is an important aspect of utilizing gaming and gameplay in education in general. Special care should be taken to ensure that non-competitive students aren't forced to compare their work with each other unless they are ready to. Otherwise, these students who are often more modest can be discouraged. Also, a wide variety of rewards should be used since motivations among individuals can vary greatly [121].

10.1.7.7 Limitations of an ITS for Sketching Fundamentals

It is important to acknowledge the limitations of our system and the the future challenges in making an ITS for sketching viable for widespread use.

10.1.7.8 Technological Limitations

Providing useful feedback on sketches is a difficult task to automate. Even if it is depicting a real object, a sketch is an abstraction of the real world. Even state-of-the-art computer vision classification and segmentation techniques may fail as sketches become increasingly complex. As was mentioned earlier, the students appreciated the real-time feedback on accuracy and line smoothness, but were not convinced that it was the only feedback they needed to improve. Indeed, there are many nuances to sketching well that human instructors can communicate readily, but intelligent systems are not currently capable of providing.

However, the technology is always improving. Future data-driven approaches may allow for more nuanced feedback on details like shading [162], and eventually feedback on more sophisticated and complex perspective sketches that combine many primitives together. There have also been attempts at finding a holistic "score" which encompasses how a student is improving overall [176]. As systems can build more sophisticated learner models, the feedback they can provide will also become increasingly powerful and useful.

10.1.7.9 Challenges of Digital Sketching

There is evidence that skills gained with digital sketching transfer to analog sketching readily as an individual is essentially training the same fine motor skills [206]. However, that does not change the fact that digital sketching has different affordances and "feel." Sketches on paper can be pinned up on a wall and shared easily. They have a tangible quality that is ideal for ideation and the creative process. Students expressed that while they saw advantage in digital practice for basic forms, they did not want to invest time in more complex sketches using the ITS. For more complex sketches they wanted to use either more fully-featured digital tools or traditional paper. An important design decision for the future is to determine the overall scope of the ITS and if it should go beyond just fundamentals. At a certain point, students will want to use media they prefer and are used to using.

There are also ongoing challenges with devices. As was mentioned in the results, several of the teachers who implemented the software mentioned challenges with getting enough devices for their students as well as ensuring they were of sufficient quality. While tablets have become increasingly ubiquitous, tablets that also allow for stylus input and quality sketching experiences have only started to become widely available and affordable for everyday consumers in the last few years. As these devices continue to become available and the price point lowers, an ITS for sketching will become much more viable for a wider variety of classrooms.

10.1.7.10 Challenges of Different Settings

In this study, we deployed the software in various domains (high school and university), as well as different curricula (art, industrial design, mechanical engineering). Each domain comes with unique needs and challenges and so does specific curricula. With the exception of the *Introduction to Sketching* course which was completely focused on sketching (T5, T6), the other courses included digital sketching as a minor component. In the case of our deployments, it was up to each teacher how much time they wanted to devote to teaching sketching skills in their curricula, and their deployment of the ITS was based on that decision.

We have yet to fully understand how much use of the ITS leads to the ideal benefits in improving sketching skills and motivating practice of sketching skills. This makes it challenging to make specific recommendations for deploying in the various domains. Presently, we believe the approach used by the *Introduction to Sketching* course is adequate. T5 and T6 asked their students to use the system at least 30 minutes per week over the semester as part of a "participation" grade. The students found this a reasonable request, and a previous study showed they could make measurable improvements in accuracy and speed over the course of the semester with that same approach [72]. Using the system more than that might even have diminishing returns, as another prior study has shown improvements in the accuracy of basic forms can be made in just 15 minutes [142].

Independently of how little or how much the system is utilized, and in what domain, we believe teachers can immediately extract the benefits of a tool like *Sketchtivity* as it can be deployed with very little burden on them while making it easy to grade and gauge class participation.

11. SUMMARY

11.1 Answers to Research Questions

11.1.1 R1. Can Intelligent Real-time Feedback from Sketch Recognition Help Improve Perspective Drawing Ability?

A1. Yes. The perspective accuracy recognition algorithm described in chapter 8 was developed to answer this research question. We found that the 20 participants in the experimental group who received real-time feedback driven by the algorithm on their perspective sketches improved the accuracy of their sketches by a statistically significant amount ($p < 0.0005$) on average afterwards. Meanwhile the 20 participants in the control condition who did not receive real-time feedback and merely sketched with a perspective grid did not improve by a statistically significant amount ($p = 0.34$) on average. This implies the real-time feedback driven by sketch recognition helped them understand how 2-point perspective works and that their was a mild learning effect and transfer of knowledge.

Additionally, the comments from the participants confirmed that the feedback was helpful to them in understanding perspective. The comments also helped us to better understand the best way to use the real-time feedback to guide learners. Based on our findings, we believe the best way to utilize the algorithm is to pair it with traditional teaching methods.

11.1.2 R2. What is the Effect of Real-time Feedback Driven by Sketch Recognition on Drawing Ability and Drawing Self-efficacy?

A2. Through various studies we found that real-time feedback driven by sketch recognition can help individuals make measurable and statistically significant improvements to accuracy, line quality (smoothness), and speed of basic primitive forms [72, 142].

We also found through qualitative data and self-reported ratings of confidence that the feedback could be beneficial for self-efficacy. The real-time feedback on shading in *DrawMyPhoto* led to very high quality drawings that pleasantly surprised many participants [162] and their self-rated confidence on those drawings was much higher than those in the control group ($p < 0.02$).

As for *Sketchtivity*, the students interviewed said that the feedback along with the scores and profile helped them with confidence, suggesting improvement in self-efficacy. Further evaluations may determine more nuances of how the system is effecting their self-efficacy.

11.1.3 R3. What Motivates Students to Practice Sketching? What is The Effect of Game-based Learning and Gamification on Motivation to Practice Sketching?

A3. As was discussed in Chapter 7, both the game-based learning approach (*ZenSketch*) and the gamification approach were motivating to students and helped them improve their basic line work. However, only a minority of participants who identified as being competitive played *ZenSketch* many times. We discovered that the game-based learning approach is not necessarily better than gamification in this context because it can lead to individuals trying to "game" the system and miss the point of the game—which is to help them learn something. We also found that the game could have diminishing returns. Initially individuals will improve their basic line work in a measurable way, but after many repeated plays they will no longer keep improving.

Further research would involve applying gamification and game-based learning to other drawing fundamentals.

11.1.4 R4. What is the Effect of Co-created Provocative Stimuli on Creativity?

A4. As was discussed in Chapter 9, we discovered that using ambiguous stimuli in the beginning of the ideation process resulted in less fixation on humanoid forms and provoked more unusual forms. Other effects included cultivating holistic thinking, stimulating exploration, and enjoyment. There was also a consensus that the stimuli are most useful in the earliest stages of the creative process.

Based on this finding we believe provocative stimuli could be utilized in a variety of ways to help individuals generate more creative drawings and be less fixated on obvious solutions. In particular, novices can benefit from provocative stimuli to help them overcome the widely known "fear of the blank page" concern [198].

11.1.5 R5. What are the Benefits and Limitations of Using an ITS for Sketching Fundamentals in Existing Courses?

A5. This was answered through the deployment and grounded theory study conducted during the 2017–2018 school year described in Chapter 10. The categories that emerged for teachers and students respectively were:

Teachers

Offloads grading work

Validates insights about students

Cultivates exploration and community

Easy to deploy

Limited access to quality devices

Students

Real-time feedback insightful but limited

Gameplay motivates in different ways

Considered a strong warm-up tool

Digital practice seen as an advantage

Not engaged by creative challenges

11.2 List of Contributions

- A novel sketch-based game called ZenSketch and a study which showed participants could improve basic line work in a measurable way in just 15 minutes.

- A description of the DrawMyPhoto 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.

- An analysis of how motivations of individuals with different sketching skill levels change and evolve. This analysis may be useful for other researchers, educators, and technologists involved in design education and/or art education.

- A framework for how various approaches to sketch-based gameplay can target these motivations and encourage sketching practice for individuals with different skill levels.

- The results of a semester-long case study of two sketch-based gameplay implementations in a university sketching course and a high school art course. We gathered interesting insights about how students engaged with our approaches and how they were motivated.

- A general-purpose rectilinear perspective accuracy sketch recognition algorithm with strong performance and novel sketch recognition features.

- The promising results of a user study and exploration of providing intelligent real-time feedback to novices using the algorithm. On average, participants could improve their perspective accuracy by a statistically significant amount by using the intelligent feedback, and this improvement carried over when the feedback was removed.

- Insights about the optimal way to utilize the different forms of feedback for faded scaffolding to best help novices learn perspective sketching skills.

- A discovery that using ambiguous stimuli in the beginning of the ideation process resulted in less fixation on humanoid forms and provoked more unusual forms. Other effects included cultivating holistic thinking, stimulating exploration, enjoyment, resistance to a loss of control and freedom, and a consensus that the stimuli are most useful in the earliest stages of the creative process.

- Implications for further research in utilizing ambiguous stimuli to augment human creativity and provoke more exploration, along with design implications for related co-creative systems and creativity support tools.

- A stronger understanding of the potential of an ITS for sketching fundamentals to improve the educational experience in the studio for both teachers and students.

12. FUTURE WORK

The methods described in this dissertation, while promising, have room for improvement. This section describes various improvements that can be made in the future including improving feedback, algorithms, building more lessons for other fundamentals, making the creative challenges more engaging, adding more games, and exploring more visualizations of sketch data and sketching performance.

12.1 More Nuanced Feedback

In Chapter 10 we discussed how the students found the feedback useful, but were aware that it did not tell the full story. In order to provide feedback that is increasingly human-like, we must move beyond simple metrics computed from gesture recognition, and use data-driven machine learning approaches.

The perspective accuracy recognition algorithm described in chapter 8 and the feedback it provides could be made more robust and accurate by collecting more training data and fine-tuning the machine learning associated with it. As was mentioned, it could also be improved by modeling the context of the strokes (e.g. grouping strokes of a window together or grouping sidewalk lines together). This could better prepare the algorithm for integration into a commercial product.

12.2 More Interactive Lessons

This dissertation only briefly touched on shading, but once the fundamentals of line work and perspective are mastered, value and shading becomes an important aspect of learning to draw. New interactive lessons that provide feedback on the proper amount of pressure to apply when shading could accomplish this. This was partially explored in the *DrawMyPhoto* work discussed in Chapter 6.

12.3 Challenges

In Chapter 10 we discussed how the students did not find the creative challenges particularly engaging. However, in order to truly improve drawing ability, an individual must put in the work and start to draw actual objects. There are a number of ways in which they could be made more engaging including:

- A better digital sketching experience
- Utilization of ambiguous stimuli to help individuals avoid fixation and think of tangential ideas
- Better scaffolding of challenges such that individuals have a starting point
- More sophisticated and human-like feedback

12.4 Sketch-based Gameplay

The motivation study described in Chapter 7 unveiled nuances of what motivates individuals with different skill levels and how sketch-based gameplay can help. However, we are only scratching the surface with games like *ZenSketch*. Many more drawing fundamentals can be gamified and this will be exciting to explore in the future.

12.5 Data Visualization

In this dissertation we only briefly touched on the possibilities of visualizing sketch performance data. Future opportunities include better visualizations of sketching performance over time so that individuals know how they are progressing, or perhaps more importantly, if they are regressing.

13. CONCLUSION

This dissertation explored a variety of ways in which artificial intelligence and interactive digital tools can help people improve their drawing ability *holistically*. In order to do so, the focus was to cultivate not just the fine motor skills related to drawing but also the common psychological needs of individuals learning drawing. Those needs include self-efficacy [107] and motivation to practice sketching. This was primarily accomplished through sketch-based gameplay and real-time feedback driven by sketch recognition.

This dissertation also explored a method to cultivate creativity and avoid fixation by utilizing ambiguity. As an individual progresses in their drawing ability, they naturally can become more creative, particularly if they have been trained in a proper way and have strong fundamentals.

Lastly, a deployment of the software in three Texas high schools and at the Georgia Institute of Technology allowed us to more closely understand how an intelligent tutoring system for drawing fundamentals can be integrated into existing curricula. We believe this may be the best possible way to learn drawing as human instructors could continue to do what they do best with the added benefit of a tool that greatly improves the confidence and motivation of students.

While the methods herein are by no means perfect and have room for improvement, they serve as a strong foundation for digital tool(s) for helping people of all ages improve their drawing ability. The mission associated with this dissertation is “Anyone can learn to draw” and that mission is just beginning.

REFERENCES

[1] K. G. Short, G. Kauffman, and L. H. Kahn, "'I just need to draw:' responding to literature across multiple sign systems," *The Reading Teacher*, vol. 54, pp. 160–171, Oct 2000. ISSN-0034-0561.

[2] D. A. Schön, "Designing as reflective conversation with the materials of a design situation," *Knowledge-Based Systems*, vol. 5, no. 1, pp. 3–14, 1992.

[3] L. Moss, "How drawing and driving are alike," *American Artist*, vol. 64, pp. 42–47, Jan 2000.

[4] K. E. daSilva, "Drawing on experience: Connecting art and language," *Primary Voices K-6*, vol. 10, no. 2, pp. 2–9, 2001. ERIC=EJ634470.

[5] I. Schiferl, "Both sides now: Visualizing and drawing with the right and left hemispheres of the brain," *Studies in Art Education*, vol. 50, no. 1, pp. 67–82, 2008.

[6] R. M. Yasin, L. Halim, and A. Ishar, "Effects of problem-solving strategies in the teaching and learning of engineering drawing subject," *Asian Social Science*, vol. 8, no. 16, pp. 65–79, 2012.

[7] E. Jones, "The correlation of visual memory and perception of perspective with drawing ability," *School and Society*, pp. 174–176, 1922.

[8] S. Sorby, "Educational research in developing 3-d spatial skills for engineering students," *International Journal of Science Education*, vol. 31, pp. 459–480, Feb 2009.

[9] M.-A. Mace and T. Ward, "Modeling the creative process: A grounded theory analysis of creativity in the domain of art making," *Creativity Research Journal*, vol. 14, no. 2, pp. 179–192, 2002.

[10] M. Suwa and B. Tversky, "What do architects and students perceive in their design sketches? a protocol analysis," *Design Studies*, vol. 18, no. 4, pp. 385–403, 1997.

[11] B. Tversky, "What do sketches say about thinking," in *2002 AAAI Spring Symposium, Sketch Understanding Workshop, Stanford University, AAAI Technical Report SS-02-08*, pp. 148–151, 2002.

[12] B. Tversky, M. Suwa, M. Agrawala, J. Heiser, C. Stolte, P. Hanrahan, D. Phan, J. Klingner, M.-P. Daniel, P. Lee, *et al.*, "Sketches for design and design of sketches," in *Human Behaviour in Design*, pp. 79–86, Springer, 2003.

[13] B. Buxton, *Sketching User Experiences: Getting the Design Right and the Right Design*. Burlington, Massachusetts, USA: Morgan Kaufmann, 1 ed., 2007.

[14] B. Buxton, *Sketching User Experiences: Getting the Design Right and the Right Design*. Burlington, Massachusetts: Morgan Kaufmann, 2010. ISBN=9780123740373.

[15] M. Sturdee, S. Mann, and S. Carpendale, "Sketching sustainability in computing," in *Proceedings of the 2019 on Creativity and Cognition*, (San Diego, California, USA), pp. 29–40, ACM, 2019.

[16] M. Lewis, M. Sturdee, J. Alexander, J. Van Dijk, M. K. Rasmussen, and T. Hoang, "Sketchingdis: Hand-drawn sketching in hci," in *Proceedings of the 2017 ACM Conference Companion Publication on Designing Interactive Systems*, (Edinburgh), pp. 356–359, ACM, 2017.

[17] R. G. Ehrenberg, D. J. Brewer, A. Gamoran, and J. D. Willms, "Class size and student achievement," *Psychological Science in the Public Interest*, vol. 2, pp. 1–30, May 2001.

[18] A. Bandura, *Self-Efficacy: The Exercise of Control*. New York, New York, USA: Worth Publishers, 1^{st} ed., 1997. ISBN=9780716728504.

[19] L. A. Linnenbrink and P. R. Pintrich, "Motivation as an enabler for academic success," *School Psychology Review*, vol. 31, no. 3, pp. 313–327, 2002.

[20] R. Hickman and S. Lord, "An examination of adolescents' self-efficacy, engagement and achievement in representational drawing," *Australian Art Education*, vol. 32, no. 2, pp. 73–85, 2009.

[21] K. A. Flannery and M. W. Watson, “Perceived competence in drawing during the middle childhood years,” *Visual Arts Research*, pp. 66–71, 1991.

[22] R. Richards, R. Richards, *et al.*, “My drawing sucks!’: Children’s belief in themselves as artists,” in *NZARE/AARE Joint Conference 2003*, NZARE/AARE, 2003.

[23] C. Brühwiler and P. Blatchford, “Effects of class size and adaptive teaching competency on classroom processes and academic outcome,” *Learning and Instruction*, vol. 21, no. 1, pp. 95–108, 2011.

[24] T. R. Kelley, “Design sketching: A lost skill: It is not enough to have students know how to create design sketches but to also know the purpose of sketching in design,” *Technology and Engineering Teacher*, vol. 76, no. 8, p. 8, 2017.

[25] K. D. Regmi, “Global construction of literacy policies for ‘least developed countries:’ focus on ethiopia, nepal, and sierra leone,” *Adult Education Quarterly*, p. 0741713619837350, 2019.

[26] M. Pleck, M. McGrath, G. Bertoline, D. Bowers, and M. Sadowski, “Factors affecting the engineering design graphics curriculum: Past, present, future,” in *Proceedings of the NSF Symposium on Modernization of the Engineering Design Graphics Curriculum*, pp. 43–52, 1990.

[27] I. de Vere, G. Melles, A. Kapoor, *et al.*, “Sketchfest: Emphasising sketching skills in engineering learning,” in *DS 74: Proceedings of the* 14^{th} *International Conference on Engineering & Product Design Education (E&PDE12) Design Education for Future Wellbeing, Antwerp, Belguim, 06-07.9. 2012*, pp. 443–448, 2012.

[28] M. Graves, “Architecture and the lost art of drawing,” *The New York Times*, vol. 1, 2012.

[29] D. Sleeman and J. S. Brown, *Intelligent tutoring systems*. London: Academic Press, 1982.

[30] A. T. Corbett, K. R. Koedinger, and J. R. Anderson, “Intelligent tutoring systems,” in *Handbook of human-computer interaction*, pp. 849–874, Elsevier, 1997.

[31] K. R. Koedinger, J. R. Anderson, W. H. Hadley, and M. A. Mark, "Intelligent tutoring goes to school in the big city," *International Journal of Artificial Intelligence in Education (IJAIED)*, vol. 8, pp. 30–43, 1997.

[32] G. S. Plano, *The Effects of the Cognitive Tutor Algebra on Student Attitudes and Achievement in a 9th Grade Algebra Course*. PhD thesis, Seton Hall University, South Orange, NJ, 2004.

[33] J. A. Kulik and J. Fletcher, "Effectiveness of intelligent tutoring systems: a meta-analytic review," *Review of Educational Research*, vol. 86, no. 1, pp. 42–78, 2016.

[34] R. Vesselinov and J. Grego, "Duolingo effectiveness study," *City University of New York, USA*, vol. 28, 2012.

[35] R. A. Sottilare, A. Graesser, X. Hu, and H. Holden, *Design recommendations for intelligent tutoring systems: Volume 1-learner modeling*, vol. 1. US Army Research Laboratory, 2013.

[36] R. Davis, "Position statement and overview: Sketch recognition at mit," in *AAAI Sketch Understanding Symposium*, 2002.

[37] T. Hammond and R. Davis, "Tahuti: A geometrical sketch recognition system for uml class diagrams," in *Technical Report SS-02-08: Papers from the 2002 Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium on Sketch Understanding*, (Menlo Park, California, USA), p. 8 pages, AAAI, July 28-August 1, 2002.

[38] T. Hammond, K. Oshiro, and R. Davis, "Natural editing and recognition of uml class diagrams," in *MIT Lab Abstract. Artificial Intelligence Laboratory*, (Cambridge, MA, USA), p. 2 pages, MIT, September 2002.

[39] J.-I. Herrera-Camara and T. Hammond, "Flow2code: from hand-drawn flowcharts to code execution," in *Proceedings of the Symposium on Sketch-Based Interfaces and Modeling*, p. 3, ACM, 2017.

[40] S. Ray, J. I. Herrera-Cámara, M. Runyon, and T. Hammond, "Flow2code: Transforming hand-drawn flowcharts into executable code to enhance learning," in *Inspiring Students with Digital Ink*, pp. 79–103, Springer, 2019.

[41] T. Hammond and R. Davis, "Ladder, a sketching language for user interface developers," *Computers & Graphics*, vol. 29-4, pp. 518–532, August 2005. `https://doi.org/10.1016/j.cag.2005.05.005`.

[42] T. Hammond and R. Davis, "Ladder: A sketch recognition language," in *MIT Lab Abstract. Computer Science and Artificial Intelligence Laboratory (CSAIL)*, (Cambridge, MA, USA), p. 2 pages, MIT, September 2004.

[43] B. Yu and S. Cai, "A domain-independent system for sketch recognition," in *Proceedings of the* 1^{st} *International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*, pp. 141–146, ACM, 2003.

[44] C. Alvarado, M. Oltmans, and R. Davis, "A framework for multi-domain sketch recognition," in *Proceedings of AAAI Spring Symposium on Sketch Understanding*, pp. 1–8, 2002.

[45] B. Paulson and T. Hammond, "Paleosketch: Accurate primitive sketch recognition and beautification," in *Proceedings of the* 13^{th} *International Conference on Intelligent User Interfaces (IUI)*, (Canary Islands, Spain), pp. 1–10, ACM, January 13–16, 2008.

[46] B. Paulson, P. Rajan, P. Davalos, R. Gutierrez-Osuna, and T. Hammond, "What!?! no rubine features?: using geometric-based features to produce normalized confidence values for sketch recognition," in *HCC Workshop: Sketch Tools for Diagramming*, pp. 57–63, 2008.

[47] B. Paulson and T. Hammond, "A system for recognizing and beautifying low-level sketch shapes using ndde and dcr," in *ACM Symposium on User Interface Software and Technology (UIST'07)*, 2007.

[48] T. Hammond, B. Eoff, B. Paulson, A. Wolin, K. Dahmen, J. Johnston, and P. Rajan, “Free-sketch recognition: putting the CHI in sketching,” in *CHI'08 Extended Abstracts on Human Factors in Computing Systems*, pp. 3027–3032, ACM, 2008.

[49] P. Taele, L. Barreto, and T. Hammond, “Maestoso: An intelligent educational sketching tool for learning music theory,” in *Proceedings of the Twenty-Seventh Innovative Applications of Artificial Intelligence Conference*, IAAI '15, (Palo Alto, California, USA), pp. 3999–4005, AAAI, 2015.

[50] P. Taele, L. Barreto, and T. Hammond, “Hashigo: A next-generation sketch interactive system for japanese kanji,” in *Proceedings of the Twenty-First Innovative Applications of Artificial Intelligence Conference*, IAAI '15, (Palo Alto, California, USA), pp. 153–158, AAAI, 2009.

[51] P. Taele and T. Hammond, “Boponoto: An intelligent sketch education application for learning zhuyin phonetic script,” in *Proceedings of the* 21^{st} *International Conference on Distributed Multimedia Systems*, DMS '15, (Skokie, Illinois, USA), pp. 101–107, Knowledge Systems Institute, 2015.

[52] S. Valentine, F. Vides, G. Lucchese, D. Turner, H.-H. A. Kim, W. Li, J. Linsey, and T. Hammond, “Mechanix: A sketch-based tutoring system for statics courses,” in *Proceedings of the Twenty-Fourth Innovative Applications of Artificial Intelligence Conference (IAAI)*, (Toronto, Canada, Canada), pp. 2253–2260, AAAI, July 2012.

[53] S. Valentine, F. Vides, G. Lucchese, D. Turner, H.-h. Kim, W. Li, J. Linsey, and T. Hammond, “Mechanix: a sketch-based tutoring and grading system for free-body diagrams,” *AI Magazine*, vol. 34, no. 1, pp. 55–55, 2013.

[54] S. Valentine, R. Lara-Garduno, J. Linsey, and T. Hammond, “Mechanix: A sketch-based tutoring system that automatically corrects hand-sketched statics homework,” in *The Impact of Pen and Touch Technology on Education*, pp. 91–103, Springer, 2015.

[55] O. Atilola, M. Field, E. McTigue, T. Hammond, and J. Linsey, “Mechanix: a sketch recognition truss tutoring system,” in *ASME 2011 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pp. 645–654, American Society of Mechanical Engineers Digital Collection, 2012.

[56] T. Nelligan, S. Polsley, J. Ray, M. Helms, J. Linsey, and T. Hammond, “Mechanix: a sketch-based educational interface,” in *Proceedings of the 20th International Conference on Intelligent User Interfaces Companion*, pp. 53–56, ACM, 2015.

[57] A. N. Bhat, *Sketchography—Automatic Grading of Map Sketches for Geography Education*. Master's thesis, Texas A&M University, College Station, TX, 2017.

[58] J. J. LaViola, Jr. and R. C. Zeleznik, “Mathpad2: A system for the creation and exploration of mathematical sketches,” in *ACM SIGGRAPH 2004 Papers*, SIGGRAPH 2004, (New York, New York, USA), pp. 432–440, ACM, 2004.

[59] D. Dixon, M. Prasad, and T. Hammond, “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, (New York, New York, USA), pp. 897–906, ACM, 2010.

[60] D. Cummings, F. Vides, and T. Hammond, “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, (Aire-la-Ville, Switzerland, Switzerland), pp. 97–106, Eurographics Association, 2012.

[61] K. Forbus, J. Usher, A. Lovett, K. Lockwood, and J. Wetzel, “Cogsketch: Sketch understanding for cognitive science research and for education,” *Topics in Cognitive Science*, vol. 3, no. 4, pp. 648–666, 2011.

[62] K. D. Forbus, J. M. Usher, A. M. Lovett, K. Lockwood, and J. Wetzel, “Cogsketch: Open-domain sketch understanding for cognitive science research and for education.,” *SBM*, vol. 8, pp. 159–166, 2008.

[63] J. W. Wetzel and K. Forbus, "Automated critique of sketched mechanisms," in *Twenty-First IAAI Conference*, 2009.

[64] A. Lovett, E. Tomai, K. Forbus, and J. Usher, "Solving geometric analogy problems through two-stage analogical mapping," *Cognitive Science*, vol. 33, no. 7, pp. 1192–1231, 2009.

[65] A. Lovett, K. Forbus, and J. Usher, "A structure-mapping model of raven's progressive matrices," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 32, 2010.

[66] A. Lovett and K. Forbus, "Cultural commonalities and differences in spatial problem-solving: A computational analysis," *Cognition*, vol. 121, no. 2, pp. 281–287, 2011.

[67] M. de los Angeles Chang and K. D. Forbus, "Using quantitative information to improve analogical matching between sketches," in *Twenty-Fourth IAAI Conference*, 2012.

[68] E. Sagi, D. Gentner, and A. Lovett, "What difference reveals about similarity," *Cognitive Science*, vol. 36, no. 6, pp. 1019–1050, 2012.

[69] K. Forbus, D. Gentner, B. Jee, B. Sageman, and D. Uttal, "Drawing on experience: Use of sketching to evaluate knowledge of spatial scientific concepts," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 31, 2009.

[70] P. Yin, K. D. Forbus, J. Usher, B. Sageman, and B. D. Jee, "Sketch worksheets: A sketch-based educational software system," in *Twenty-Second IAAI Conference*, 2010.

[71] C. Lassheikki, I. Nokelainen, Y. Shan, M. Kaos, M. Raevaara, H. Valja, and P. Hämäläinen, "Fair and encouraging—developing an automatic assessment system in aalto drawing," in *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*, (Melbourne, VIC, Australia), pp. 337–344, ACM, 2018.

[72] S. Keshavabhotla, B. Williford, E. Hilton, P. Taele, W. Li, J. Linsey, and T. Hammond, "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, (Los Angeles, CA, USA), pp. 1–11, ACM, July 29-30 2017.

[73] Y. Kuribayashi, E. Yamaga, and T. Sasaki, “Handwritten cube recognition and evaluation based on perspectives for sketch training system,” in *2019 IEEE/SICE International Symposium on System Integration (SII)*, pp. 34–39, IEEE, 2019.

[74] J. Fernquist, T. Grossman, and G. Fitzmaurice, “Sketch-sketch revolution: An engaging tutorial system for guided sketching and application learning,” in *Proceedings of the* 24^{th} *annual ACM Symposium on User Interface Software and Technology*, (Santa Barbara, California, USA), pp. 373–382, ACM, 2011.

[75] J. W. Hennessey, H. Liu, H. Winnemöller, M. Dontcheva, and N. J. Mitra, “How2sketch: generating easy-to-follow tutorials for sketching 3d objects,” *Symposium on Interactive 3D Graphics and Games*, 2017.

[76] R. Schmidt, A. Khan, K. Singh, and G. Kurtenbach, “Analytic drawing of 3D scaffolds,” in *ACM SIGGRAPH Asia 2009 Papers*, SIGGRAPH Asia ’09, (New York, NY, USA), pp. 149:1–149:10, ACM, 2009.

[77] R. Schmidt, A. Khan, G. Kurtenbach, and K. Singh, “On expert performance in 3d curve-drawing tasks,” in *Proceedings of the* 6^th *Eurographics Symposium on Sketch-Based Interfaces and Modeling*, SBIM ’09, (New York, NY, USA), pp. 133–140, ACM, 2009.

[78] S.-H. Bae, R. Balakrishnan, and K. Singh, “Everybodylovessketch: 3d sketching for a broader audience,” in *Proceedings of the* 22^{nd} *Annual ACM Symposium on User Interface Software and Technology*, UIST ’09, (New York, NY, USA), pp. 59–68, ACM, 2009.

[79] Y. J. Lee, C. L. Zitnick, and M. F. Cohen, “Shadowdraw: Real-time user guidance for freehand drawing,” in *ACM Transactions on Graphics (TOG)*, vol. 30, (New York, New York, USA), p. 27, ACM, 2011.

[80] E. Iarussi, A. Bousseau, and T. Tsandilas, “The drawing assistant: Automated drawing guidance and feedback from photographs,” in *Proceedings of the* 26^{th} *Annual ACM Symposium on User Interface Software and Technology*, UIST ’13, (New York, New York, USA), pp. 183–192, ACM, 2013.

[81] J. Xie, A. Hertzmann, W. Li, and H. Winnemöller, “Portraitsketch: Face sketching assistance for novices,” in *Proceedings of the* 27^{th} *Annual ACM Symposium on User Interface Software and Technology*, UIST ’14, (New York, New York, USA), pp. 407–417, ACM, 2014.

[82] L. Benedetti, H. Winnemöller, M. Corsini, and R. Scopigno, “Painting with bob: Assisted creativity for novices,” in *Proceedings of the* 27^{th} *Annual ACM Symposium on User Interface Software and Technology*, UIST ’14, (New York, New York, USA), pp. 419–428, ACM, 2014.

[83] N. Davis, C.-P. Hsiao, Y. Popova, and B. Magerko, “An enactive model of creativity for computational collaboration and co-creation,” in *Creativity in the Digital Age*, pp. 109–133, Springer, 2015.

[84] G. Fischer, E. Giaccardi, H. Eden, M. Sugimoto, and Y. Ye, “Beyond binary choices: Integrating individual and social creativity,” *International Journal of Human-Computer Studies*, vol. 63, no. 4-5, pp. 482–512, 2005.

[85] A. Kantosalo, J. M. Toivanen, P. Xiao, and H. Toivonen, “From isolation to involvement: Adapting machine creativity software to support human-computer co-creation.,” in *ICCC*, pp. 1–7, 2014.

[86] G. N. Yannakakis, A. Liapis, and C. Alexopoulos, “Mixed-initiative co-creativity.,” in *FDG*, 2014.

[87] D. Long, M. Jacob, and B. Magerko, “Designing co-creative ai for public spaces,” in *Proceedings of the 2019 on Creativity and Cognition*, pp. 271–284, ACM, 2019.

[88] P. T. Ravikumar, “Notational communication with co-creative systems: Studying improvements to musical coordination,” in *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*, pp. 518–523, ACM, 2017.

[89] A. Das and B. Gambäck, “Poetic machine: Computational creativity for automatic poetry generation in bengali.,” in *ICCC*, pp. 230–238, 2014.

[90] D. Long, M. Jacob, N. Davis, and B. Magerko, "Designing for socially interactive systems," in *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*, pp. 39–50, ACM, 2017.

[91] N. Davis, C.-P. Hsiao, K. Y. Singh, L. Li, S. Moningi, and B. Magerko, "Drawing apprentice: An enactive co-creative agent for artistic collaboration," in *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, pp. 185–186, ACM, 2015.

[92] N. Davis, C.-P. Hsiao, K. Yashraj Singh, L. Li, and B. Magerko, "Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent," in *Proceedings of the* 21^{st} *International Conference on Intelligent User Interfaces*, pp. 196–207, ACM, 2016.

[93] Robert Angel, "*Pictionary*." Game [Board Game], 1985.

[94] OMGPop and Zynga, "*DrawSomething*." Game [Mobile], February 2012.

[95] B. Paulson, B. Eoff, A. Wolin, J. Johnston, and T. Hammond, "Sketch-based educational games: Drawing kids away from traditional interfaces," in *Proceedings of the* 7^{th} *international conference on Interaction Design and Children (IDC)*, (Chicago, IL, USA), pp. 133–136, ACM, July 2008. ISBN=978-1-59593-994-4.

[96] H.-H. Kim, S. Valentine, P. Taele, and T. Hammond, "Easysketch: A sketch-based fine motor skill recognizing educational interface for children emerging technology research strand," in *The Impact of Pen and Touch Technology on Education* (T. Hammond, S. Valentine, A. Adler, and M. Payton, eds.), Human-Computer Interaction Series, ch. 4, pp. 35–46, Switzerland: Springer Publishing Company, Incorporated, 2015.

[97] F. Vides, P. Taele, H.-H. A. Kim, and T. Hammond, "Intelligent feedback for kids using sketch recognition," in *ACM SIGCHI 2012 Conference on Human Factors in Computing Systems Workshop on Educational Interfaces, Software, and Technology*, (Austin, Texas, USA), ACM, 2012.

[98] F. Vides, "Tayouki: A sketch-based tutoring system for young kids," MS Master's Thesis, Texas A&M University (TAMU), College Station, TX, USA, August 2012. Advisor: Tracy Hammond, `http://hdl.handle.net/1969.1/ETD-TAMU-2012-08-11497`.

[99] C. T. Truong, D.-H. Nguyen-Huynh, and M.-T. Tran, "How do we teach young children new concepts via sketching?," in *Distributed, Ambient, and Pervasive Interactions* (N. Streitz and P. Markopoulos, eds.), (Limenas Chersonisou, Greece), pp. 292–303, Springer International Publishing, 2014.

[100] R. Wasson, D. Mould, R. Biddle, and C. S. Martinez, "A sketching game for art history instruction," in *Proceedings of the International Symposium on Sketch-Based Interfaces and Modeling*, SBIM '13, (New York, NY, USA), pp. 23–31, ACM, 2013.

[101] E. Law and L. Von Ahn, "Input-agreement: A new mechanism for collecting data using human computation games," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (Boston, Massachusetts, USA), pp. 1197–1206, ACM, 2009.

[102] G. Johnson and E. Y.-L. Do, "Games for sketch data collection," in *Proceedings of the* 6^{th} *Eurographics Symposium on Sketch-Based Interfaces and Modeling*, (New Orleans, LA), pp. 117–123, ACM, 2009.

[103] I. Iacovides, J. Aczel, E. Scanlon, J. Taylor, and W. Woods, "Motivation, engagement and learning through digital games," *International Journal of Virtual and Personal Learning Environments (IJVPLE)*, vol. 2(2), pp. 1–16, 2011.

[104] N. J. Whitton, *An Investigation into the Potential of Collaborative Computer Game-Based Learning in Higher Education*. PhD thesis, Edinburgh Napier University, 2007.

[105] P. Wouters, C. Van Nimwegen, H. Van Oostendorp, and E. D. Van Der Spek, "A meta-analysis of the cognitive and motivational effects of serious games.," *Journal of Educational Psychology*, vol. 105, no. 2, p. 249, 2013.

[106] K. Vanlehn, "The behavior of tutoring systems," *International Journal of Artificial Intelligence in Education*, vol. 16, no. 3, pp. 227–265, 2006.

[107] A. Bandura, "Guide for constructing self-efficacy scales," *Self-Efficacy Beliefs of Adolescents*, vol. 5, no. 1, pp. 307–337, 2006.

[108] B. J. Zimmerman, S. Bonner, and R. Kovach, *Developing self-regulated learners: Beyond achievement to self-efficacy.* American Psychological Association, 1996.

[109] B. J. Zimmerman, "Self-efficacy: An essential motive to learn," *Contemporary Educational Psychology*, vol. 25, no. 1, pp. 82–91, 2000.

[110] P. R. Pintrich and E. V. De Groot, "Motivational and self-regulated learning components of classroom academic performance.," *Journal of Educational Psychology*, vol. 82, no. 1, p. 33, 1990.

[111] J. A. Livingston, *Metacognition: An Overview.* PhD thesis, University at Buffalo, Buffalo, NY, 2003.

[112] N. Dabbagh and A. Kitsantas, "Personal learning environments, social media, and self-regulated learning: A natural formula for connecting formal and informal learning," *The Internet and Higher Education*, vol. 15, no. 1, pp. 3–8, 2012.

[113] Y.-J. Joo, M. Bong, and H.-J. Choi, "Self-efficacy for self-regulated learning, academic self-efficacy, and internet self-efficacy in web-based instruction," *Educational Technology Research and Development*, vol. 48, no. 2, pp. 5–17, 2000.

[114] C.-H. Wang, D. M. Shannon, and M. E. Ross, "Students' characteristics, self-regulated learning, technology self-efficacy, and course outcomes in online learning," *Distance Education*, vol. 34, no. 3, pp. 302–323, 2013.

[115] M. Barak, "Motivating self-regulated learning in technology education," *International Journal of Technology and Design Education*, vol. 20, no. 4, pp. 381–401, 2010.

[116] J. León, J. L. Núñez, and J. Liew, "Self-determination and stem education: Effects of autonomy, motivation, and self-regulated learning on high school math achievement," *Learning and Individual Differences*, vol. 43, pp. 156–163, 2015.

[117] R. M. Ryan and E. L. Deci, “Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being.,” *American Psychologist*, vol. 55, no. 1, p. 68, 2000.

[118] B. S. Bloom, “Learning for mastery. instruction and curriculum. regional education laboratory for the carolinas and virginia, topical papers and reprints, number 1.,” *Evaluation Comment*, vol. 1, no. 2, p. n2, 1968.

[119] L. Vygotsky, “Interaction between learning and development,” *Readings on the Development of Children*, vol. 23, no. 3, pp. 34–41, 1978.

[120] R. Bartle, “Hearts, clubs, diamonds, spades: Players who suit muds,” *Journal of MUD Research*, vol. 1, no. 1, p. 19, 1996.

[121] C. Phillips, D. Johnson, M. Klarkowski, M. J. White, and L. Hides, “The impact of rewards and trait reward responsiveness on player motivation,” in *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play*, (Melbourne, VIC, Australia), pp. 393–404, ACM, 2018.

[122] T. W. Malone, “Toward a theory of intrinsically motivating instruction,” *Cognitive Science*, vol. 5, no. 4, pp. 333–369, 1981.

[123] V. J. Shute, “Stealth assessment in computer-based games to support learning,” *Computer Games and Instruction*, vol. 55, no. 2, pp. 503–524, 2011. ISBN=978-1617354083.

[124] L. A. Sharp, “Stealth learning: Unexpected learning opportunities through games.,” *Journal of Instructional Research*, vol. 1, pp. 42–48, 2012.

[125] S. Robertson and T. Bertling, *How to Draw: Drawing and Sketching Objects and Environments from your Imagination*. Culver City, California, USA: Design Studio Press, 1^{st} ed., 2013. ISBN=978-1933492735.

[126] M. Tovey, S. Porter, and R. Newman, “Sketching, concept development and automotive design,” *Design Studies*, vol. 24, no. 2, pp. 135–153, 2003.

[127] J. A. Koncelik and K. Reeder, *Conceptual Drawing*. Boston, Massachusetts, USA: Course Technology, 1^{st} ed., 2008. ISBN=978-1418080976.

[128] B. Edwards, *The New Drawing on the Right Side of the Brain*. New York, New York, USA: Tarcher, 2^{nd} ed., 1999. ISBN=978-1585429219.

[129] B. Edwards, *Drawing on the Right Side of the Brain*. ACM, 1997.

[130] B. Dodson, *Keys to drawing*. Cincinnati, OH, USA: North Light Ohio, 1985.

[131] M. Willenbrink and M. Willenbrink, *Drawing for the absolute beginner: a clear & easy guide to successful drawing*. F+W Media, Inc., 2006.

[132] C. W. Garcia, *Drawing for the Absolute and Utter Beginner, Revised*. Watson-Guptill, 2018.

[133] E. Norling, *Perspective made easy*. Courier Corporation, 1999.

[134] P. Van Sommers, *Drawing and Cognition: Descriptive and Experimental Studies of Graphic Production Processes*. New York, New York, USA: Cambridge University Press, 1984.

[135] S. Tobias, J. D. Fletcher, and A. P. Wind, "Game-based learning," in *Handbook of Research on Educational Communications and Technology*, pp. 485–503, Springer, 2014.

[136] K. Sjolen and A. MacDonald, *Learning Curves: An Inspiring Guide to Improve Your Design Sketch Skills*. Sundsvall, Sweden: Keeos Design Books, 1 ed., 2011.

[137] N. L. L. H. Williford, Taele, "Persketchtivity: an intelligent pen-based educational application for design sketching instruction," *Revolutionizing Education with Digital Ink*, pp. 115–127, 2016.

[138] B. Williford, "Sketchtivity: Improving creativity by learning sketching with an intelligent tutoring system," in *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*, (Singapore), pp. 477–483, ACM, 2017.

[139] D. Weibel and B. Wissmath, "Immersion in computer games: The role of spatial presence and flow," *International Journal of Computer Games Technology*, vol. 2011, p. 6, 2011.

[140] J. Chen, "Flow in games (and everything else)," *Communications of the ACM*, vol. 50, no. 4, pp. 31–34, 2007.

[141] K. Salen and E. Zimmerman, *Rules of Play: Game Design Fundamentals*. MIT press, 2004.

[142] B. Williford, M. Runyon, A. H. Malla, W. Li, J. Linsey, and T. Hammond, "Zensketch: A sketch-based game for improving line work," in *Extended Abstracts Publication of the Annual Symposium on Computer-Human Interaction in Play*, pp. 591–598, ACM, 2017.

[143] M. Zhao and S.-C. Zhu, "Artistic rendering of portraits," in *Image and Video-Based Artistic Stylisation*, pp. 237–253, New York, NY: Springer, 2013.

[144] C. Lu, L. Xu, and J. Jia, "Combining sketch and tone for pencil drawing production," in *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*, (Annecy, France), pp. 65–73, Eurographics Association, 2012.

[145] M. I. Norton, D. Mochon, and D. Ariely, "The ikea effect: When labor leads to love," *Journal of Consumer Psychology*, vol. 22, no. 3, pp. 453–460, 2012.

[146] R. W. Belk, "Extended self in a digital world," *Journal of Consumer Research*, vol. 40, no. 3, pp. 477–500, 2013.

[147] R. De Reyna, *How to draw what you see*. New York City, NY: Watson-Guptill, 1996.

[148] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and <achine Intelligence*, vol. 6, pp. 679–698, 1986.

[149] T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu, "An efficient k-means clustering algorithm: Analysis and implementation," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 7, pp. 881–892, 2002.

[150] B. Hangün and Ö. Eyecioğlu, "Performance comparison between opencv built in cpu and gpu functions on image processing operations," *International Journal of Engineering Science and Application*, vol. 1, no. 2, pp. 34–41, 2017.

[151] K. Pfeuffer, K. Hinckley, M. Pahud, and B. Buxton, "Thumb+ pen interaction on tablets.," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, (Denver, CO), pp. 3254–3266, ACM, 2017.

[152] Y. Xin, X. Bi, and X. Ren, "Acquiring and pointing: an empirical study of pen-tilt-based interaction," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (Vancouver, BC), pp. 849–858, ACM, 2011.

[153] Y. Xin, X. Bi, and X. Ren, "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*, (Austin, TX), pp. 801–804, ACM, 2012.

[154] T. M. Amabile, "Children's artistic creativity: Detrimental effects of competition in a field setting," *Personality and Social Psychology Bulletin*, vol. 8, no. 3, pp. 573–578, 1982.

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

[156] S. J. Edwards, D. J. Buckland, and J. D. McCoy-Powlen, *Developmental & Functional Hand Grasps*. Thorofare, New Jersey, USA: SLACK Incorporated, 1 ed., 2002.

[157] B. Shneiderman, "Creativity support tools," *Communications of the ACM*, vol. 45, no. 10, pp. 116–120, 2002.

[158] S. L. Jackson, J. Krajcik, and E. Soloway, "The design of guided learner-adaptable scaffolding in interactive learning environments," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, (New York City, NY), pp. 187–194, ACM Press/Addison-Wesley Publishing Co., 1998.

[159] M. Son, H. Kang, Y. Lee, and S. Lee, "Abstract line drawings from 2d images," in 15^{th} *Pacific Conference on Computer Graphics and Applications (PG'07)*, (Maui, HI), pp. 333–342, IEEE, 2007.

[160] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected

crfs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2018.

[161] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, (Boston, MA), pp. 3431–3440, IEEE, 2015.

[162] B. Williford, A. Doke, M. Pahud, K. Hinckley, and T. Hammond, "Drawmyphoto: Assisting novices in drawing from photographs," in *Proceedings of the 2019 on Creativity and Cognition*, (San Diego, California, USA), pp. 198–209, ACM, 2019.

[163] P. van Sommers, "A system for drawing and drawing-related neuropsychology," *Cognitive Neuropsychology*, vol. 6, pp. 117–164, Mar 1989.

[164] S. Nicholson, "Two paths to motivation through game design elements: Reward-based gamification and meaningful gamification," in *iConference 2013*, (Fort Worth, Texas, USA), pp. 671–672, iSchools, Inc., 2013.

[165] Y.-k. Chou, "Actionable gamification—beyond points, badges, and leaderboards," in *Gamification—Engaging your workforce*, London, United Kingdom: ArkGroup, 2015. ISBN=1511744049.

[166] K. Charmaz, *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*. London, United Kingdom: Sage, 2014. ISBN=978-0857029140.

[167] B. G. Glaser and A. Strauss, "Strauss (1967): The discovery of grounded theory: Strategies for qualitative research," *London: Wiedenfeld and Nicholson*, vol. 81, p. 86, 1978. ISBN=978-0202302607.

[168] B. G. Glaser, "The constant comparative method of qualitative analysis," *Social Problems*, vol. 12, no. 4, pp. 436–445, 1965.

[169] P. R. Pintrich, "A manual for the use of the motivated strategies for learning questionnaire (mslq).," *National Center for Research to Improve Postsecondary Teaching and Learning*

Project on Instructional Processes and Educational Outcomes, vol. OERI-86-0010, 1991. ERIC Number: ED338122.

[170] J. D. Helsel, *Reading Engineering Drawings Through Conceptual Sketching*. Columbus, Ohio, USA: Glencoe/Mcgraw-Hill, 1^{st} ed., 1979. ISBN=978-0070280311.

[171] D. L. Butler and P. H. Winne, "Feedback and self-regulated learning: A theoretical synthesis," *Review of Educational Research*, vol. 65, no. 3, pp. 245–281, 1995.

[172] I. Dunwell, S. de Freitas, and S. Jarvis, "Four-dimensional consideration of feedback in serious games.," in *Digital games and learning*, Curtin Research Publications, pp. 42–62, Perth, Australia: Continuum Publishing, 2011.

[173] Z. O. Toups, A. Kerne, and W. Hamilton, "Motivating play through score," in *In ACM Computer Human Interaction Workshop on Engagement by Design*, (Boston, Massachusetts, USA), ACM, 2009.

[174] S. Zhai, J. Kong, and X. Ren, "Speed-accuracy tradeoff in fitts' law tasks—on the equivalency of actual and nominal pointing precision," *International Journal of Human-Computer Studies*, vol. 61, no. 6, pp. 823–856, 2004.

[175] R. A. Schmidt, D. E. Sherwood, H. N. Zelaznik, and B. J. Leikind, "Speed-accuracy trade-offs in motor behavior: Theories of impulse variability," in *Motor Behavior: Programming, Control, and Acquisition* (H. Heuer, U. Kleinbeck, and K.-H. Schmidt, eds.), pp. 79–123, Berlin, Heidelberg: Springer Berlin Heidelberg, 1985.

[176] T. Hammond, S. P. A. Kumar, M. Runyon, J. Cherian, B. Williford, S. Keshavabhotla, S. Valentine, W. Li, and J. Linsey, "Its not just about accuracy: Metrics that matter when modeling expert sketching ability," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 8, no. 3, p. 19, 2018.

[177] J. C. Rosser, P. J. Lynch, L. Cuddihy, D. A. Gentile, J. Klonsky, and R. Merrell, "The impact of video games on training surgeons in the 21^{st} century," *Archives of Surgery*, vol. 142, no. 2, pp. 181–186, 2007.

[178] E. Hilton, W. Li, S. H. Newton, M. Alemdar, R. Pucha, and J. Linsey, “The development and effects of teaching perspective free-hand sketching in engineering design,” in *ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, (Charlotte, North Carolina), pp. V003T04A013–V003T04A013, American Society of Mechanical Engineers (ASME), 2016.

[179] E. Hilton, M. Paige, B. Williford, W. Li, T. Hammond, and J. Linsey, “Improving the sketching ability of engineering design students,” in *International Conference on Engineering Design (ICED)*, vol. 9: Design Education, (Vancouver, BC, Canada), pp. 217–224, The Design Society, August 21–25 2017. ISBN: 978-1-904670-97-1.

[180] D. Avola, A. Del Buono, G. Gianforme, S. Paolozzi, and R. Wang, “Sketchml a representation language for novel sketch recognition approach,” in *Proceedings of the* 2^{nd} *International Conference on PErvasive Technologies Related to Assistive Environments (PETRA)*, (Corfu, Greece), p. 31, ACM, June 9–13 2009.

[181] A. Wolin, B. Eoff, and T. Hammond, “Shortstraw: A simple and effective corner finder for polylines.,” in *SBM*, pp. 33–40, 2008.

[182] D. Rubine, “Specifying gestures by example,” *SIGGRAPH Computer Graphics*, vol. 25, pp. 329–337, July 1991.

[183] A. C. Long Jr, J. A. Landay, L. A. Rowe, and J. Michiels, “Visual similarity of pen gestures,” in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pp. 360–367, ACM, 2000.

[184] A. Liaw and M. Wiener, “Classification and regression by random forest,” *R News*, vol. 2, no. 3, pp. 18–22, 2002.

[185] H. J. Perkinson, *Learning From Our Mistakes*, vol. 36. International Society for General Semantics, 1979.

[186] M. E. Fonteyn, B. Kuipers, and S. J. Grobe, “A description of think aloud method and protocol analysis,” *Qualitative Health Research*, vol. 3, no. 4, pp. 430–441, 1993.

[187] L. Vygotsky, "Zone of proximal development," *Mind in Society: The Development of Higher Psychological Processes*, vol. 5291, p. 157 pages, 1987.

[188] J. Xing, H.-T. Chen, and L.-Y. Wei, "Autocomplete painting repetitions," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 6, p. 172, 2014.

[189] J. Xing, L.-Y. Wei, T. Shiratori, and K. Yatani, "Autocomplete hand-drawn animations," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, p. 169, 2015.

[190] P. R. Pintrich, C. A. Wolters, and G. P. Baxter, "Assessing metacognition and self–regulated learning," in *Metacognitive Assessment* (G. Schraw, ed.), Lincoln, Nebraska, USA: University of Nebraska Press, 2000.

[191] J. J. Shah, N. Vargas-Hernandez, J. D. Summers, and S. Kulkarni, "Collaborative sketching (c-sketch)—an idea generation technique for engineering design," *The Journal of Creative Behavior*, vol. 35, no. 3, pp. 168–198, 2001.

[192] W. W. Gaver, J. Beaver, and S. Benford, "Ambiguity as a resource for design," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 233–240, ACM, 2003.

[193] T. Hammond, "Dialectical creativity: Sketch-negate-create," in *Studying Visual and Spatial Reasoning for Design Creativity*, pp. 91–108, Springer, 2015.

[194] A. E. Clarke, C. Friese, and R. Washburn, *Situational analysis in practice: Mapping research with grounded theory*. Routledge, 2016.

[195] K. T. Konecki, "Visual grounded theory: A methodological outline and examples from empirical work," *Revija Za Sociologiju*, vol. 41, no. 2, pp. 131–160, 2011.

[196] T. B. Ward, "Structured imagination: The role of category structure in exemplar generation," *Cognitive Psychology*, vol. 27, no. 1, pp. 1–40, 1994.

[197] R. A. Finke, T. B. Ward, and S. M. Smith, *Creative cognition: Theory, research, and applications*. MIT press Cambridge, MA, 1992.

[198] C. Joyce, “The blank page: effects of constraint on creativity,” *London School of Economics*, 2009.

[199] T. B. Ward, M. J. Patterson, and C. M. Sifonis, “The role of specificity and abstraction in creative idea generation,” *Creativity Research Journal*, vol. 16, no. 1, pp. 1–9, 2004.

[200] P. Langley, “User modeling in adaptive interface,” in *UM99 User Modeling*, pp. 357–370, Springer, 1999.

[201] P. Tierney and S. M. Farmer, “Creative self-efficacy: Its potential antecedents and relationship to creative performance,” *Academy of Management Journal*, vol. 45, no. 6, pp. 1137–1148, 2002.

[202] K. Eissen and R. Steur, *Sketching: Drawing Techniques for Product Designers*. Amsterdam, The Netherlands: BIS Publishers, 12 ed., 2009.

[203] B. J. Zimmerman and A. Kitsantas, “The hidden dimension of personal competence: Self-regulated learning and practice,” in *Handbook of Competence and Motivation* (A. Elliot, C. Dweck, and S. Carol, eds.), pp. 509–526, New York, New York, USA: Guilford Publications, 2005.

[204] A. Kukla, *Social constructivism and the philosophy of science*. Routledge, 2013.

[205] M. Csikszentmihalyi, “Flow and the psychology of discovery and invention,” *New Yprk: Harper Collins*, 1996.

[206] H.-h. Kim, P. Taele, J. Seo, J. Liew, and T. Hammond, “Easysketch 2: a novel sketch-based interface for improving children’s fine motor skills and school readiness,” in *Proceedings of the Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*, pp. 69–78, Eurographics Association, 2016.

APPENDIX A

QUESTIONING ROUTES

A.1 Teacher Interviews

Why do you teach sketching in your course?

How do you like to teach sketching in your course?

How did you choose to implement SketchTivity in your course?

What were come challenges you encountered (if any)?

What devices were your students using?

Was their any issues with access to devices?

What did you have your students do?

How engaged were the students?

Were the students competitive with each other?

Did you notice any effects on their confidence?

Did you notice any effects on their motivation to practice?

Did any students continue using it?

Did the instructor profile help you at all?

What would you improve about the tool?

Would you integrate such a tool in to your classroom in the future? Why or why not?

A.2 Student Interviews

How was your confidence in your sketching ability going in to this semester?

How often do you practice sketching?

What motivates you to practice sketching?

What do you sketch generally?

Do you sketch digitally? How often vs paper?

How often did you use SketchTivity?

Walk me through your experiences using SketchTivity

Which features did you use? (for each feature:)

How did [this feature] affect your confidence?

How did [this feature] motivate you?

How did [this feature] influence your creativity?

How do you think the real-time feedback can be improved?

How did viewing your "stats" affect your confidence or motivation?

Did you feel like you improved from using SketchTivity?

If so, do you think those improvements translate to paper and traditional media?

What do you think are the advantages of SketchTivity in learning sketching fundamentals?

What do you think are the disadvantages of SketchTivity in learning sketching fundamentals?

What would you like to see changed or added to SketchTivity?