CATEGORIZING AGE GROUP BY SKETCHING ABILITY

A Thesis

by

XIEN EHIMIYEN THOMAS

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

MASTER OF SCIENCE

Chair of Committee, Tracy Hammond
Co-Chair of Committee, Richard Furuta
Committee Member, Jeffery Liew
Head of Department, Scott Schaefer

August 2020

Major Subject: Computer Science

TABLE OF CONTENTS

Page

LIST OF FIGURES

ABSTRACT

From an early age, children begin developing critical motor control skills that will be used for the rest of their lives. While everyday activities like standing or walking require gross motor control, fine motor control over the hands is vital to healthy development. Fine motor skills contribute significantly to reading, writing, crafting, drawing, and more, all of which are important for communication and school readiness. Pediatricians can evaluate a child's gross and fine motor skills using questionnaires and activities. For example, there may be periods when a pediatrician meets with a child to see if their tracing abilities are getting better or worse through sketching questionnaires. Usually, these questionnaires ask the adults to draw with their child. However, it is particularly difficult to fully evaluate a child's drawings through a handful of sketches created in just these meetings. We propose to create a sketching system that will collect all drawing data from parents and children that can then automatically evaluate and differentiate a child's sketch from an adult's using only their strokes. We believe that each sketching stroke is unique and includes artifacts of the user's age. Working with sketches from children between 2–5 and adults over 18, we build different statistical features to determine age groups and train a system by analyzing different stroke patterns. A system capable of automatically categorizing users into age groups can enable new solutions for the assessment of fine motor skills as well as enable novel applications related to collaborative learning software and age-based authentication.

DEDICATION

I would like to thank my family, my mom Sandra Walker, my father Friday Thomas, my brother Narada Walker, my older brother Ontreal Walker, and my Aunt Barbara. Without my friends and family, I would not be what I am today. I would like to also thank my lord and savior Jesus Christ for his spiritual guidance.

ACKNOWLEDGEMENTS

I would like to thank my advisor Dr.Tracy Hammond and the Sketch Recognition Lab for their guidance and support throughout my thesis. Without my advisor and fellow lab members, especially Dr. Paul Taele, Dr. Anna Stepanova, Dr. Blake Williford, Seth Polsley, Larry Powell, Raniero Lara-Garduño, Josh Cherian, Jung-In Koh, Duc Hoang, and Samantha Ray, I would not have been able to finish my thesis. I would like to thank my committee members: Dr. Richard Furuta and Dr. Jeffrey, for your support. I also would like to thank Dr. Dilma Da Silva; without her continuous support and energy, I would not have made it this far in my graduate career. Final thanks go to my department and Texas A&M University for a great experience.

1. INTRODUCTION

It is essential to begin evaluating healthy childhood development from an early age, as a child's development determines readiness for school and other activities. Critical school-readiness skills include the ability to hold a pencil, use scissors, and open small items like glue [109]. Multiple studies have shown that delayed development can lead to many complications later in life, including difficulty handling small objects, reduced reading and writing achievements, and other complications through preteen years [12, 33, 89, 98]. Because evaluating children's development is so important, there are several widely-used assessments for this purpose, many of which rely on motor control [45, 56, 62, 97, 112].

1.1 Gross and Fine Motor Control

Typically, there are two forms of motor control that assessments consider: gross and fine [46, 97]. Gross motor control refers to vast movements of the body like waving, jumping, or walking and can be tested in different ways [21, 54, 69]. For instance, gait analysis alongside gross motor function is used by Kurtz et al. to understand if a brain injury is present in a child [54]. Understanding gross motor control can also help in the early identification of autism in children, as shown in work by Ozonoff et al. [69].

Gross motor control contributes to significant movements of the child and their development skills, whereas fine motor control ties to the child's school-readiness because it is vital for everyday activities like cutting, writing, painting, coloring, and handling small items [89]. There is a neurological aspect of fine motor skills that indicates later reading and math achievements [12, 33, 61]. If not correctly identified, delayed fine motor control can lead to the decline of a child's capabilities to retain new information or focus on school-based activities [13]. Additionally, their ability to communicate and write legibly or adequately can suffer [12].

1.2 Assessments

Most often, parents consult pediatricians for all aspects of their child's health, including development. Pediatricians rely on both questionnaires and parent involvement to evaluate children's motor delays, if any [82]. The Ages and Stages Questionnaire (ASQ) and Parents' Evaluations of Developmental Status (PEDS) are two ways to evaluate the development of children [9, 32, 97]. Another standard measure is the Bayley Mental Development Index score. Most of these tools assess children's behavioral, mental, and physical health in ages ranging from four months to five years old. Assessments can be generally classified as indirect or direct; indirect methods can be done by parents at home, while direct methods are most completed with an expert such as a pediatrician.

Questionnaires are screening tools to understand if children's motor control is developing well depend upon their age and other factors [9,97]. Usually, screening tools can be answered by the parents as an indirect evaluation. Still, they can only capture some of the activities of a child compared to direct evaluation with experts [48]. Indirect evaluation is more efficient and less costly, but a parent's response might be biased or missing problems that a pediatrician may identify through direct assessment [48]. Although indirect screenings alleviate the burden on pediatricians, 17% of the children that go through screenings would still need professional help for their motor delay [68]. Also, direct development evaluations are more correlated to Bayley index scores than conventional indirect approaches [102]. Therefore, direct evaluation approaches with pediatricians are more reliable in discovering motor delays in a child.

1.3 Sketching and Development

One way of evaluating fine motor control is through drawing—the small movements of our hands while drawing or writing correlates tightly with fine motor control [15,29,93]. For fine motor control, it is important to evaluate children's sketches comparatively to their parents [33]. It is recommended to have a parent draw with their child; this type of involvement can be associated with the child's performance at schools [94]. Although it is suitable for parents to interact with the child in drawing activities, if a parent saves multiple images of their drawings together, it could be hard to evaluate only the child's sketches during direct evaluations, while disregarding the adult's sketches. Not only can it be hard, but it is also time-consuming to review each drawing.

Pediatricians can use these drawings along with standard questionnaires to comprehensively evaluate the development of a child. Given that the actual likelihood of motor delay in children is relatively low [68], direct evaluation can be limiting on a pediatrician's time. The limited amount of time available for pediatricians could affect children, who do have motor delay problems, of having shorter evaluations. The average pediatrician may manage as many as 2,000 children [60]. Therefore, the best practice would be to have an indirect evaluation screening method that helps lower the ratio between children and pediatricians while yielding useful data [48,102].

1.4 Modernizing Fine Motor Assessment through Sketch Recognition

Instead of writing on paper, this work proposes that the child use a pen-enabled device so the device can process their strokes and automatically determine which strokes are from the child and which strokes are from the parent. Prior work in the field of sketch recognition extracts features from the user's sketches for automatic classification and the evaluation of their drawings [26,36,38,49,53,92,95,104,105].

This research could provide insight as to how to classify each stroke belonging to an adult or a child. The pediatrician could potentially use the information given to differentiate the strokes and use the child's strokes to evaluate if the child has any motor delay.

Research has shown that by age two, 90% of children have moderate ability when operating touchscreen devices [20]. A proliferation of studies and articles have reflected on the impacts of these devices on children's learning [16, 44, 66], social development [64], and parental interactions[1,2]. Research also shows that using tablets or other forms of a mobile device can enhance the children's motivation in learning and collaboration [57]. Children often interact with tablet devices because there are many educational applications. Educational apps are the third most popular category on Apple's App Store, behind gaming and business usage[3]. With hundreds of thousands of educational apps, one can imagine ample uses for a system that can automatically identify whether the parent or child is currently using the device. For example, an application that helps a child learn to write could ask a parent to show the child what to do and then wait until the child has written something before evaluating it.

1.5 Case Scenarios

1.5.1 Without system

A parent is asked to fill out a questionnaire about their child's developmental process to the best of their ability. Questions about drawing skills are included in the "Fine Motor Skills" section. These questions ask the parent how well the child draws a particular shape, or if the child has any difficulties holding a pen. The next

[1] https://www.littlethings.com/reasons-not-to-give-children-technology/
[2] http://www.bbc.co.uk/guides/z3tsyrd
[3] https://www.statista.com/statistics/270291/popular-categories-in-the-app-store/

day the parent meets with a pediatrician to evaluate their child's motor skills. The pediatrician reviews what the parent provided and then discusses the score given by the questionnaire. When reviewing the sketches, the pediatrician may need to repeatedly ask if the stroke belongs to the parent or the child. The parent may not remember or remembers incorrectly, which hinders the time and performance of the pediatrician's evaluation. The pediatrician is unable to determine accurately how well the child is developing, and because the meeting took longer than expected, other meetings will be impacted. To prepare for their next session, the parent must record their child's activity correctly.

1.5.2 With system

With the proposed system, the parent could still fill out a questionnaire, but the labeling system is automated. The system would extract x and y coordinates and timestamps from the sketches to process the data into meaningful features [58,74,78]. The system will then use the extracted features to classify each stroke. Afterward, it would save the information of the strokes and mark the sketch as belonging to a child or an adult. Afterward, it would save the information of the strokes and mark the sketch belonging to a child or an adult, which enables the ability to examine the children separately from adults for proper evaluation of the child's ability. Because the system saves a record of each sketch, the pediatrician can reference all shapes, in addition to asking the parent if the child can draw a specific shape. In this scenario, the strokes that the child made were no longer evaluated only by the parent but also the pediatrician, which reduces a potential bias introduced by the parent. The resulting score is less ambiguous than the outcome of a standard ASQ. The meeting is beneficial for everyone and finishes in a timely manner.

1.6 Design goals

It is vital to evaluate the fine motor control of children. Sketching is an essential method to measure fine motor control, but this requires a more direct approach for the pediatrician to monitor a child's sketch. Direct approaches can be problematic because pediatricians may not have enough time to evaluate each child directly, and for those who need care the most, indirect methods can be inadequate. These problems need solutions which support children's use of touchscreen devices for sketching in the comfort of their own homes, while also providing a better evaluation of fine motor control to enhance their opportunities to learn and develop. This work proposes to create a sketching system that can differentiate strokes from a child from those of an adult using characteristics of their strokes. The goal is to support many further applications beyond assessment, such as educational tools or child authentication.

1.7 Research Questions

R1. How accurately can a smart sketch interface differentiate adults and children using only single strokes from free-form sketches?

Building on Kim's work [49], this study seeks to verify the findings that sketch features embed characteristics of a user's age. While Kim demonstrated classification among specific age groups using strokes and sketches, this work is focused on verifying that work by differentiating adults from children using single strokes in free-form sketches on a smart drawing interface for generic users.

R2. What are the top features that distinguish adult from child sketches?

In accordance with the goal to enable many new applications such as fine motor assessment or enhanced interactive learning, it is important to identify the key age-based characteristics of sketches. These top features are critical to creating machine learning algorithms that distinguish users based on age.

2. RELATED WORK

Quantifying motor skills is crucial in determining the healthy development of a child [98]. Pediatricians use various automated assessments to work with children and their parents to improve children's motor abilities. One example during children's evaluation may involve parents and their children collaborating to perform sketching activities with traditional means [57]. For this type of drawing interaction, existing systems and tools can assist by better informing parents of their children's' motor abilities—by classifying their strokes as either child or adult skill levels—through the following areas.

- **Screening tools** can provide an assessment by evaluating children's fine motor skills.
- **Children's interaction with tools and devices** discusses how children learn motor skills and how educators use technology in the classroom.
- **Sketch recognition** reviews the use of intelligent processing methods to analyze stroke-based input.
- **Differentiation** refers to similar techniques to distinguish between users on a single device.

2.1 Screening Tools

One reasonably known screening tool that parents use is the Ages and Stages Questionnaire (ASQ) [88]. Pediatricians use questionnaires to measure and keep a record of children's motor development. Motor development includes children's communication skills, problem-solving skills, and other skills that require the muscles in the body. Many cultures tested the validity of this tool under different languages [25, 45, 80]. In turn, the ASQ became reliable and widespread for parents and pediatricians to use. The ASQ consists of 21 intervals in five different areas for children of ages 2 to 48 months: personal, gross motor control, fine motor control, problem-solving, and communication. The answered questions provide an approximate estimate of how much motor delay a child may have compared to other children that took the tool previously. Squires et al. [88] discovered that approximately 10–20% of children have a delay, and this screening tool lessened the number of children visiting their pediatrician [87]. However, pediatricians primarily use the Ages and Stages Questionnaire for indirect evaluation, which does not include experts to be part of the screening process. Instead, children's screening process is done by non-experts that need to complete several activities with the child that they might not record. A record of the activities done by children during the screening process can shorten the time of visits with a pediatrician.

The Parents' Evaluations of Developmental Status (PEDS) is another screening tool used for early development [32]. The survey is similar to ASQ, as the questionnaire asks the parent questions that relate to the child's fine motor skills, gross motor skills, problem-solving, and communication skills. PEDS has an additional section for interpreting the results of the screening tools. For example, if there is a single concern, the interpretation form asks for another screening, and if the child failed the assessment again, then there is additional testing. In this way, parents have a continuous record of the child's development.

2.2 Children Interaction with Tools and Devices

Gual et al. studied the interaction of children with standard classroom objects [30]. The problem is in understanding the motor development of many children from different grade levels. In total, they used 253 children for their tasks. Each task contributes to a multitude of motor composite areas: fine motor precision and fine motor integration. Activities similar to cutting, writing, and folding require the precision of the hands and fingers because most fine motor functionalities heavily involve the muscles in our fingers. Taking similar motor functions, Gual discovered that motor skills develop with age, and with this information, children cannot perform some activities. Therefore, there is a distinct difference once children mature.

A series of papers expected that tablets are used more often in the classrooms [14, 18, 55, 66, 77]. Researchers used tablets to observe students' behavior while learning a foreign language [81]. One challenge is how much it affects the teaching styles of teachers and professors. Teachers and professors could eventually change their teaching habits that are more accustomed to class learning, which is similar to work by Montrieux et al. [63]. Montrieux created a Google Form to evaluate a set of teachers for their teaching styles and received 123 responses. Their findings discuss that teachers switch into two categories based on how active the students interact with their tablets. That is, teachers can either be innovative teachers or instrumental teachers, where both styles of teaching are highly interactive with the students and their tablet devices. The study supports that children can be highly interactive when completing activities.

2.3 Sketch Recognition

Academics and researchers have applied sketch recognition to perform many different tasks [24, 34, 37, 40, 65, 90, 91, 110, 111, 113] to interpret a user's intentions. In a high-level view, a sketch is a collection of strokes where strokes are a collection of points. Recognizing drawings require multiple inferences about how a user draws a shape [72–74]. Some systems attempt to understand the user's intentions when they create a diagram [1, 2, 10, 17, 41, 59], by using these tools, there are fewer errors, and ideas are allowed to flow more freely without the concern of beauty in the diagram. Other systems that manage diagrams must be different in other domains where symbols in one domain do not translate in other domains [28, 31, 70]. To accommodate this, some researchers create systems to ease the development of new sketch recognition systems [39]. Usually, these systems can be incorporated into smart user interfaces for easier interaction [7,11]. There are also other uses to sketch recognition for quantifying different algorithms [8,76]. The understanding of human intention is achieved by using gesture-based, vision-based, or geometric-based sketch recognition.

Valentine et al. created a sketching interface called Mechanix, a tutoring system for mechanical engineers [5,6,95,96]. Their system lowers the workload of professors and teaching assistants (TA) when they grade exams, homework, and quizzes. In introductory engineering courses, the disproportionate ratio between students and instructors results in grading, taking a significant portion of the instructors' time. Worse still, the students receive their marked exams and homework after a long delay. Not having immediate feedback could cause students to make the same mistakes during exams or evaluations. Students need to properly learn different concepts in introductory courses to continue their career and understand advanced concepts. Valentine approached this problem with a sketching interface and discussed the various concepts shown in introductory engineering courses. The user would draw as they would on pen and paper, and the system gives immediate feedback if the user drew a concept incorrectly. At the time, 111 students were used in their study, and the students found the system to be very helpful, especially the immediate feedback they received. The contribution was that an intelligent user interface was made using sketch recognition algorithms and that students learned the basic concepts of mechanical engineering while benefiting from automatic feedback. Valentine et al. system demonstrate the value of sketch recognition algorithms in determining a user's intention during a sketch.

There are several systems already that attempt to measure drawing ability and try to improve the way people sketch [19, 22, 35, 42, 103]. In work by Williford et al., they wanted to improve design and communication skills for engineering students when sharing and visualizing their ideas [105]. Often when engineers sit together to share ideas, they are either writing on a whiteboard or chalkboard to visualize or communicate. Without proper design principles, engineers will have a difficult time persuading or convincing others of their idea. Learning to perform basic lines, circles, and ellipses can be treated as lessons in design. The study collected 80 students over two semesters between the ages of 18-19 years. The study consisted of a group using Persketchtivity and another group using pen and paper. The traditional pen and paper group was given homework problems. The group using Persketchtivity was given the same problems but through the application. Students using the application had increased sketching and visualization skills. They were able to appreciate the real-time feedback from Persketchtivity. An intelligent system such as Persketchtivity aids student's design and used sketch recognition to recognize their strokes. Various recognizers can form shapes using strokes to make several observations in a sketch. Similar works like iCanDraw [23] and EyeCanDraw [19] also provide aid or sketching instructions when drawing a human figure. We build upon these works, looking both at the measurement methods and the features they use for recognition.

There are numerous examples in the field of sketch recognition interfaces that target children learning how to sketch for developmental learning [50–53, 71, 101], and these examples could also relate to techniques for improving the classification of children using touchscreen gestures [3, 79, 83–86, 100, 107, 108]. Two systems, TAY-ouKi [101] and KimCHI [50, 51, 53], provide more examples that introduce the basic drawing concepts to children. Kim et al. created a fine motor skill framework based on children overall drawing skills. There were two types of classifiers: KimCHI, which classifies the overall drawing skill of children and KimCHI2, which classify drawing skills based on curvature. The motivation for their work is to assess child school-readiness and social behavior. The problem with current solutions is that they use human experts to evaluate full sketches of child drawings. Looking through the child drawings can be cumbersome and prone to human error. The sketch recognition algorithm that they used to show certain features that paper assessment applications cannot show as the time of sketching and the curvature of their sketch. In their first application, KimCHI, had children draw from a template beside them and generate features that help determine the amount of motor delay. In their second study, they generated different features that relate to the curvature to understand children's fine motor skills. Using 10-fold classification on over 150 children, the KimCHI classifier scored an f-measure of 0.904 from distinguishing between adults and children. Also, using the same data, KimCHI2 uses corner-finding and curvature to achieve an f-measure of 0.744 [49, 53]. For both KimCHI and KimCHI2, both approaches use sketch-based features that are used for the entire sketch instead of a single stroke.

2.4 Differentiation

One of Hang et al.'s approaches to extract features uses two of the user's fingers. Using two fingers would restrict normal movement that one finger can do alone [43]. For example, to draw a square using two fingers, both fingers would have to draw while spreading apart their fingers simultaneously. The technique was first motivated by another paper that used stroke order to distinguishing adults. However, the stroke order could always change, or the user could change the stroke order and create the same result visually. The study showed that they had a high number of false negatives and false positives for authentication. This mainly due to users creating their template from a simple example template. Therefore, templates were too simple to use for distinguishing adults from children. Even if the application had the desired security capabilities, the interaction with the application was too strenuous for the user's fingers. Having a user trace a template using two fingers at the same time would be uncomfortable. The act of using their fingers at the same time would be frustrating and causes some users to draw with only one finger instead of two. Even though their idea of using the distance between two fingers is interesting, it would not be very distinct in the same age group. Using the distance between two fingers in different age groups could prove more useful for differentiating between adults and children. However, the usability of the application could be frustrating for children to use, so the children may not comply with using two fingers.

In the paper "Scribble-a-secret," Oka et al. made a sketch-based password authentication system [67]. Sketching interfaces have increased in popularity since the rise of hand-held computing devices. In their approach, they wanted a way to authenticate an individual using free-form sketches. Then from these sketches, the system would use the edge orientation from the set of strokes to authenticate. Each of the 87 participants drew a free-form sketch of a unique subject. They had to recreate their free form sketch several times for the application to extract features from that same sketch. At least ten sketches from a single user need to be collected to use edge orientation. The author took the average of each edge orientation and used this as the user template. Validating the accuracy of the algorithm, the author compared sketches with 86 participants and showed a low percentage of false negatives and false positives. Edge orientation is an interesting approach and has the advantage of better performance if the password is complicated. As the password becomes complex, the accuracy increases, but users would have a difficult time remembering their password.

3. METHODOLOGY

In this work, the proposed system is a sketch recognition system able to differentiate between adults and children. The system has been used to collect sketching information from users and pre-processed for the shape recognition algorithm. Using the messy sketches from the users, this system attempts to extract traditional sketch recognition features and other features using vision-based techniques. This chapter explains the system, the user studies, the activities, the features, and the classifiers in detail. Chapter 5 discusses the best features and validation methods.

3.1 Data Collection

This section details the collection of data from various users informing the development of the "Sketch Pals" system. Initial studies informed the design of the system to make necessary improvements before the full user study. These user studies had adults and children perform three different activities to collect strokes from tasks with prompts varying in specificity with respect to what the users needed to draw. The sketches from these activities are saved into the device for further pre-processing.

3.1.1 Design

Sketches were collected using a custom sketching application called "Sketch Pals" that ran on a Nexus 7 pen-enabled tablet device. Nexus 7 was the platform of choice because it is a cheap tablet that can be found in a common household. Initial studies took place to evaluate the feasibility and design direction of the study. These users were tasked with sketching different figures without a template to express their creativity fully. Thus the interface design went through several iterations after receiving feedback. Sketch Pals, shown in Figure 3.1, is a simple interface with functions like erasing, changing brush size, clearing a canvas, and saving sketches. The top bar shows different functionalities of the application, represented as icons. From left to right, the "New" button (a plus sign) that will save the current drawing and clear the canvas. Next, the "Clear" button (trash can) clears the canvas. The next two buttons switch the mode of the stylus: the "Draw" button (pen) switches to a pen, and the "Erase" button (rotating rectangle) switches to an eraser. Lastly, the "Save" button (floppy disk) prompts the saving process. This process asks the user if they are 18 or older. If yes, it will ask the user to name their drawing. If no, it will ask the user to call an adult over to save for them. Figures 3.2 and 3.3 are examples of sketches produced by adult and child participants, respectively. Each stroke in each sketch is in a different color.

3.1.2 User Studies

The classifiers were trained using data from a total of 39 participants, as shown in Table 3.1. Adults were over the age of 18, and the children were between the ages of 3 and 5. Each participant sketched on a pen-enabled device and had to complete three different activities. In the first activity, the participants draw for 5 minutes. They may express anything through a sketch. In the second activity, the user picks a shape to replicate for 2 minutes. They can select from a square, triangle, trapezoid, or a star. The last activity is the first activity in addition to testing each button of the system. To preserve the participants' creativity, adults were asked to work alone, and children drew without any help from their teacher or parent. The overall session includes three activities; the last activity, however, saves the sketch differently depending on whether the participant is a child or an adult. For child participants, the researcher took the device and labeled the session a child session. For adult participants, the system would label their session as an adult, and they

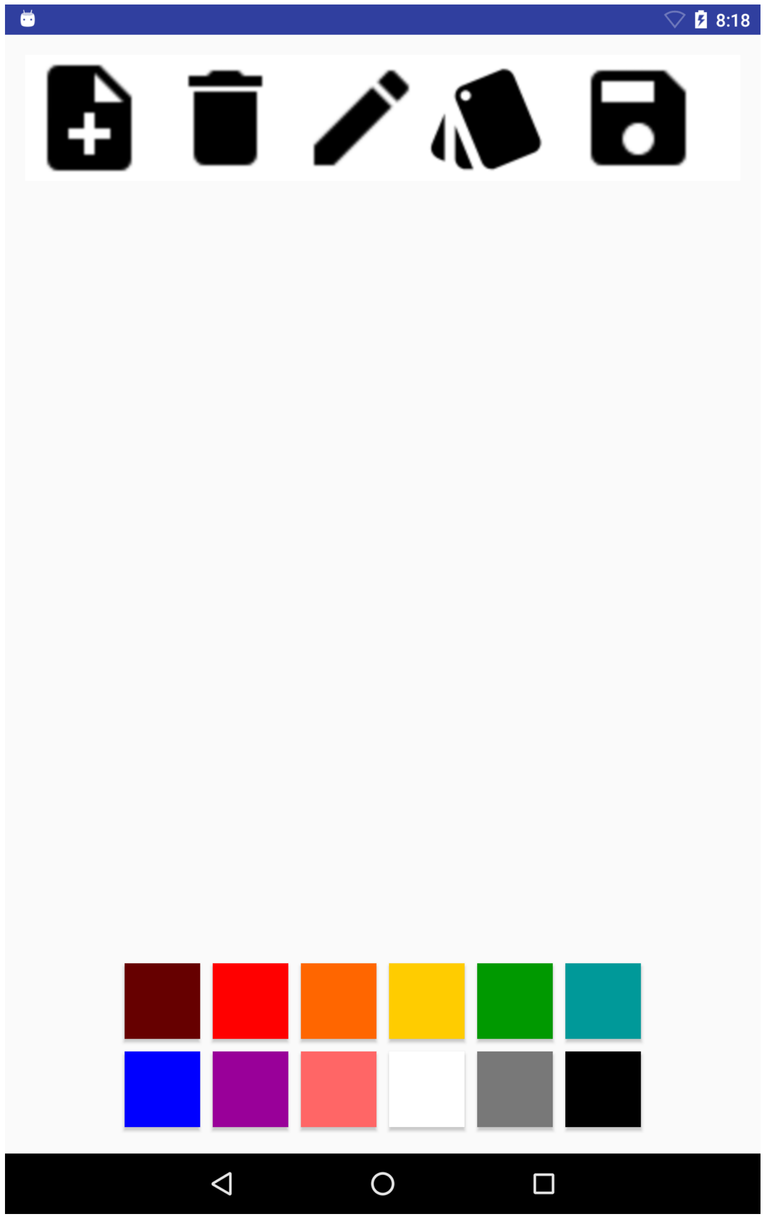

Figure 3.1: SketchPals Application

Figure 3.2: Example of an adult sketch

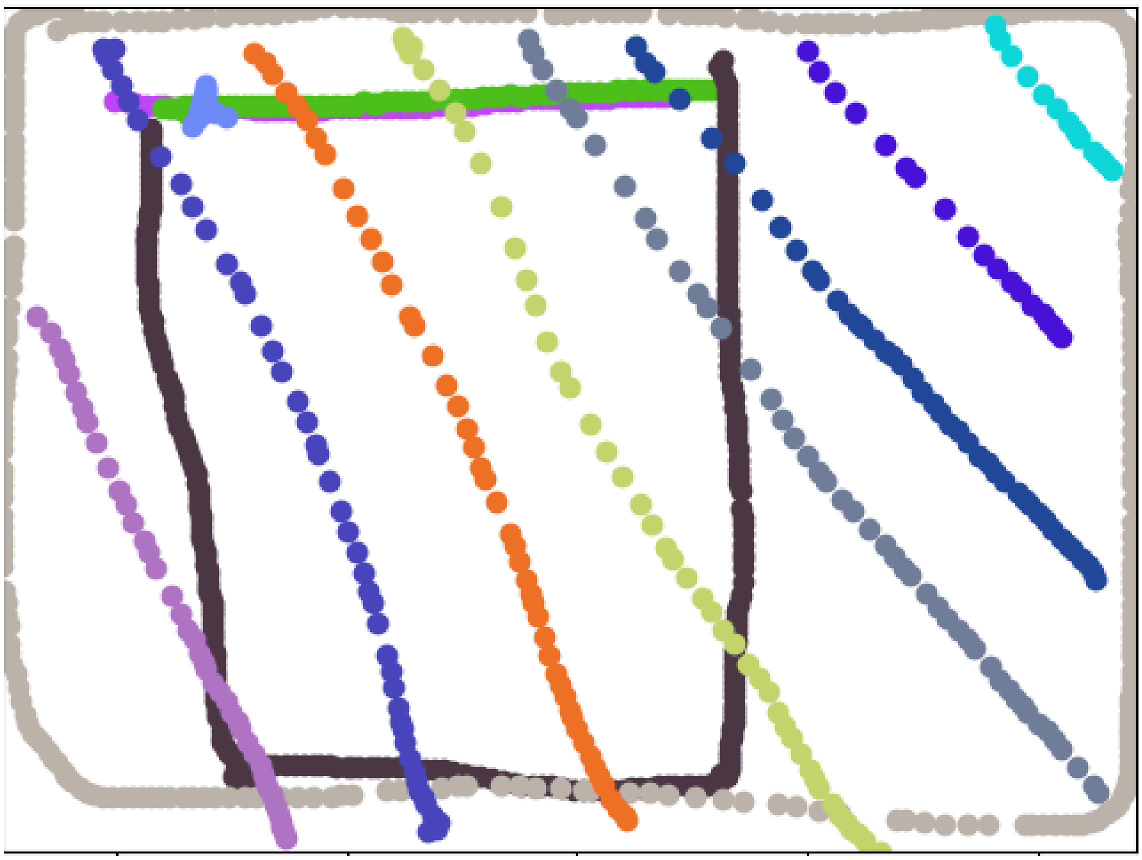

Figure 3.3: Example of a child sketch

Demographics of all participants			
Group	Adults 18+	Children 3-5	Overall
# of Volunteers	25	14	39
# of Sketches	50	28	78
# of Strokes	1185	1888	3073

Table 3.1: The basic Demographics of all users in the study

would get to choose the name of the file.

3.1.3 Data Structure

Sketch data was collected from over several sessions. Each session consists of three sketches from each activity. When a user begins to sketch, every pen-down to pen-up motion defines as a stroke. In every stroke, there is a collection of points. While the user is sketching, the system assigns each point a stroke number. Once a pen-up to pen-down motion starts, the system increments the stroke number. After the user completes the sketch, the system saves the sketch into a CSV file for later processing. The sketches that the user cleared and deleted are also collected. These unwanted sketches help determine the intentions in the sketch during analysis.

Data collection involves sampling the state of the canvas throughout the sketch, collecting x and y coordinate data and timestamps the refresh rate of the screen. Initially, the sampling rate was once per second, but some features were not calculating features accurately because the time between points was not short enough: strokes would look different but have the same calculated features. Increasing the sampling rate to the refresh rate of the screen addressed this issue. The data for each collected stroke has the timestamps in milliseconds, the x and y coordinates of the points, the color of the points, the x and y velocity of the stroke, and the stroke number.

The color shows us the difference between regular strokes and erased strokes. The velocity of the stroke is used for some sketching features. The stroke number label each point in which stroke it belongs. These are the pieces of information that are collected from the tablet device and used in the feature generation process.

3.2 Pre-processing

Resampling was not used during the feature generation process. Resampling, shown in Figure 3.4, is a typical pre-processing step in sketch recognition to standardize the number of points in the sketch by making the distance between the points equidistant. Resampling can be used to upsample, downsample, or remove the effects of an inconsistent sampling rate. Some sketch recognition algorithms rely on resampling to standardize the input. For example, template matching works best when the input sketches have the same number of points as the template as missing data points, or abnormal clusters of data points would cause these template matching to fail. In a similar vein, corner finding algorithms that rely on the geometry of the distance between points being smaller over corners than over a straight segment in the sketch could fail without equidistant points. However, this work benefits from using the raw, messy sketches over the resampled ones. In this work, understanding the differences in the production of strokes between the different categories of users can give some insight. Differences in the strokes could be removed during the resampling process. By not resampling, the system keeps each stroke unique to the user.

Through many tribulations, some files contain the entire session, meaning that a single file contains multiple sketches. The elapsed time was calculated between each stroke. Roughly, the time between strokes is 1–2 seconds, and it takes 3–5 minutes between sketches. If the time difference between strokes is below 1 minute, there is an assumption that the next stroke is in the same sketch. Otherwise, the stroke

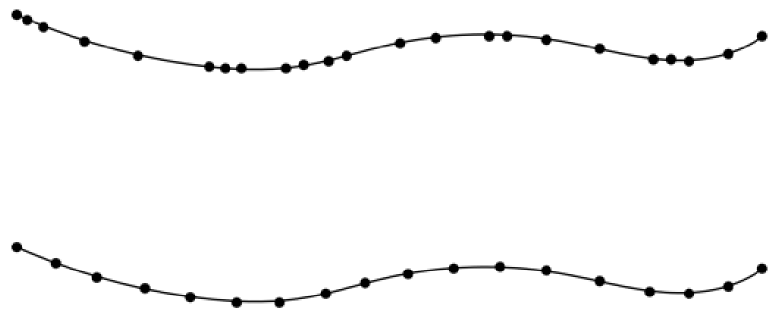

Figure 3.4: Example of a resampled sketch and an original sketch

belongs to another sketch.

The system needed another method for distinguishing strokes by utilizing the velocity, direction, and the elapsed time of the stroke. If the stroke is too fast, the stroke could break into two separate strokes. First, the system extracts velocity from the CSV file. If the total velocity increases between two strokes, then this is where the error occurs. The system subtracts the velocity from the end of one stroke and the beginning of the next stroke. If the difference is below 15, then the two strokes should be in the same stroke. This study also considered the change of direction (angle) between a set of points. The angle gave a considerable amount of leeway to the threshold, depending on if the change is greater or lower than the previous changes. From these heuristics, the system estimated when stokes was a single stroke.

3.3 Shape Recognition

The primary goal for the second activity versus the other activities is to have a set of sketches that follow a template. The second activity in the user studies consists of users tracing different shapes. The system process the traced shape and classifies the shape by using a template-matching algorithm called $P (pronounced p-dollar) [4, 99, 106]. Each template consists of a point cloud, i.e., a set of points consisting of x and y coordinate values. The user's sketch is compared to each template using Hausdorff distance to find the best template. Hausdorff distance matches each point in the input sketch to its corresponding point in the template sketch by finding the pairs of points with the smallest Euclidean distance. The template with the smallest total distance is the predicted shape. For this algorithm to work, the input must go through a pre-processing phase where the sketch is scaled, translated to the origin, and resampled in order to standardize the size, location, and point regularity.

3.4 Optimization and Feature Generation

After the sketches were collected and pre-processed, features were extracted and used later for classification. Generating features involves small calculations that may be shared across features. Instead of needlessly repeating calculations and slowing down the overall algorithm, these calculations can be computed beforehand and then passed to another function to finishing generating the features, as shown in Algorithm 1. In the "send to features" function, list, change in the distance (distDeltas), change in time (timeDeltas), and change in angle (angleDeltas) have been initialized. After initializing the values, the system computes the bounding box of the stroke shown in Algorithm 2. "DistDeltas" gets the distance between each point in a stroke. "TimeDeltas" calculates the elapsed time between each point in a stroke. "AngleDeltas" calculates the angles between points in a stroke. Then some features needed the diagonal Length (diagLen) and the bounding box angle (boundboxAng). To calculate the bounding box, first, the system finds the lowest and highest x, y coordinates. Then the system subtracts the minimum from the maximum x and y values to get the width and height, respectively.

After calculating these repetitive values for optimization purposes, the system passed these values through another function called "features." Given all the deltas

Algorithm 1 Feature Extraction

```
function sendToFeatures(stroke, length)
    list = []
    distDeltas = []
    timeDeltas = []
    angleDeltas = []
    boundingBox = getBoundingBox(stroke)
    for i = 0; i < length; ++i do
        distDeltas[i] = distChangeAtPoint(i, stroke)
        timeDeltas[i] = timeChangeAtPoint(i, stroke)
    end for
    for i = 0; i < length; ++i do
        angleDeltas[i] = angleChangeAtPoint(i, distDeltas, length)
    end for
    diagLen = sqrt(boundingBox.height^2 + boundingBox.width^2)
    boundboxAng = tan^-1(boundingBox.width / boundingBox.height)
    list = features(diagLen, boundboxAng, distDeltas, timeDeltas, angleDeltas)
    return list
end function
```

between points, "features" generates the full feature set in the order shown in Algorithm 3. The corresponding equations for each feature listed in Equations 3.1–3.17.

The main program receives the full computed feature set, where the program starts to train the classifiers. The program computes the last set of features Zernike moments 1–8 in a different data structure that only includes x, y, and time. More information about calculating each feature used in the "features" function can be found below.

- **Distance between first and last point** is the distance between the first and last point of a stroke.

$$\sqrt{(x_{p-1} - x_0)^2 + (y_{p-1} - y_0)^2} \tag{3.1}$$

Algorithm 2 Bounding Box

```
function getBoundingBox(stroke)
    list = []
    width = 0
    height = 0
    xlow = −∞
    xhigh = ∞
    ylow = −∞
    yhigh = ∞
    while nextpoint.next! = null do
        xAxis = nextpointx
        yAxis = nextpointy
        if xAxis < xlow then
            xlow = xAxis
        end if
        if xAxis > xhigh then
            xhigh = xAxis
        end if
        if yAxis < ylow then
            ylow = yAxis
        end if
        if yAxis > yhigh then
            yhigh = yAxis
        end if
        nextpoint = nextpoint.next
    end while
    width = abs(xhigh − xlow)
    height = abs(yhigh − ylow)
    list.add(width)
    list.add(height)
    return list
end function
```

Algorithm 3 Feature Generation

function FEATURES($diagLen, boundboxAng, disDeltas, timeDeltas, angleDeltas, stroke, length, boundingbox$)

$percent10Dist = totalLength * 0.10$

$ptDistFL = firstLastPtDist(stroke)$

$totalLength = totalLen(stroke)$

$totalAngle = totalAng(length, angleDeltas)$

$totalAbsAngle = totalAbsAngle(length, angleDeltas)$

$smooth = smoothness(length, angleDeltas)$

$maxspeed = maxSpeed(length, distDeltas, timeDeltas)$

$totalTime = totalTime(stroke, length)$

$aspect = abs(\frac{3.14}{4.0} - boundingboxAng)$

$curve = curviness(angleDeltas)$

$totAngDivLen = totalAngleDivTotalLength(totalAngle, totalLength)$

$dens1 = density1(totalLength, ptDistFL)$

$dens2 = density2(totalLength, diagonalLength)$

$openess = openess(ptDistFL, diagonalLength)$

$boundingArea = boundingBoxArea(boundingBox)$

$logArea = log(boundingArea)$

$totalAngleDivTotalAbsAngle =$
$totalAngleDivTotalAbsAngle(totalLength, totalAbsAngle)$

$logLength = log(totalLength)$

$logAspect = log(aspect)$

$percent5 = strokePercentEnd(stroke, 0.05, length)$

$percent3 = strokePercentEnd(stroke, 0.03, length)$

$percent2 = strokePercentEnd(stroke, 0.02, length)$

$NDDE = newNDDE(length, stroke)$

$DCR = newDCR(length, stroke)$

$distf10 = grabF10Dist(percent10Dist, length, newDistDeltas, stroke)$

$distb10 = grabB10Dist(percent10Dist, length, newDistDeltas, stroke)$

return $listoffeatures$

end function

- **Total stroke length** is the total distance in a stroke. Let n be the point in a stroke and let δ be the change in x and in y, for example $\delta x_p = x_(p-1) - x_(p)$.

$$\sum_{n=1}^{n-2} \sqrt{\delta x_n^2 + \delta y_n^2} \tag{3.2}$$

- **Total angle traversed** is the summation of angles in a stroke. Let θ be the angle.

$$\sum_{n=1}^{n-2} \theta_n \tag{3.3}$$

- **Total absolute angle traversed** is the summation of angles in a stroke.

$$\sum_{n=1}^{n-2} abs(\theta_n) \tag{3.4}$$

- **Smoothness** is the sum of the squared values of the angles in a stroke.

$$\sum_{n=1}^{n-2} \theta_n^2 \tag{3.5}$$

- **Max speed** is the sum of the squared values of the angles in a stroke.

$$\max \frac{\delta x_n^2 + \delta y_n^2}{\delta t_n^2} \tag{3.6}$$

- **Total time** is the total amount of time of the stroke.

$$t_{n-1} - t_0 \tag{3.7}$$

- **Aspect ratio** is the absolute value of 45 degrees minus the angle of the bounding box.

$$abs(45^o - BoundingBoxAngle) \tag{3.8}$$

- **Curviness** is the summation of the absolute value of all angles below 19 degrees in a stroke.

$$\sum_{n=1}^{n-2} abs(\theta'_n) \text{ where } \theta' = \{\theta \mid \theta < 19^o\} \tag{3.9}$$

- **Total angle divided by length** is the total angle divided by the length of the stroke.

$$\frac{TotalAngle}{StrokeLength} \tag{3.10}$$

- **The distance between the first and last sampling points divided by stroke length**.

$$\frac{FirstandLastPointDistance}{StrokeLength} \tag{3.11}$$

- **Diagonal of the stroke's bounding box divided by the stroke length**.

$$\frac{DiagonalofBoundingBox}{StrokeLength} \tag{3.12}$$

- **Total angle divided by total absolute angle**.

$$\frac{TotalAngle}{TotalAbsoluteAngle} \tag{3.13}$$

- **Log of the area of the bounding box**.

$$Log(AreaofBoundingBox) \tag{3.14}$$

- **Total angle of the last 5%, 3%, and 2% of a stroke**. Let v be the new set of points of the last percentage of a stroke.

$$\sum_{v=1}^{v-2} \theta_v \tag{3.15}$$

- **First 10% of the total distance**. Let v be the new set of points of the first 10% of the stroke.

$$\sum_{v=1}^{v-2} \sqrt{\delta x_v^2 + \delta y_v^2} \tag{3.16}$$

- **Last 10% of the total distance**. Let v be the new set of points of the last 10% of the stroke.

$$\sum_{v=1}^{v-2} \sqrt{\delta x_v^2 + \delta y_v^2} \tag{3.17}$$

3.5 Classifiers

The goal of this study is to determine critical features and classify each stroke by its category: adult or child. Using the data presented in Table 3.1, the system was trained and validated using five different classifiers to differentiate between children's strokes and adult strokes: Random Forest, Decision Tree, Zero-Rule, Naive Bayes, and Support Vector Machines (SVMs). For each algorithm, leave-one-out validation was done so that the classifiers do not overfit for the test data. The baseline algorithm is the Zero-Rule, which picks the majority class without regard for any features. Random Forest algorithm constructs a model based on many randomized Decision Trees. Decision Tree algorithm constructs a model by a series of cascading questions in form of a tree structure using features and values. Naive Bayes is based on Bayes' algorithm; it calculates the probability of a sample to be of a specific class based on the sample's features. SVMs generates a line to separate the data into two classes.

For classifying the data set, the system labeled each participant based on their previous response if they are 18 or older. If the user responds “yes,” then every stroke is labeled as “A” for adult and “C” for a child if the user answer “no.” The responses was then used to compare against labels that the machine learning algorithm picked. In this study, the classifiers were trained using the Sklearn library in Python [75]. To prevent the Random Forest and Decision Tree classifiers from overtraining in one class, there needed to be a way of handling an imbalanced data set. Typically each sample has an equal weight of 1.0, but in this study, there is an imbalanced amount of users from each class. In essence, there are $1,185$ samples in class “A” and $1,888$ in class “C.” When the classifier is training, it can weight each class differently based on the number of samples in each category using Equation 3.18; where w_i is the weight of class i, n is the total amount of samples in the data set, k is the number of classes, and n_i is the number of instances in class i. Once the classifier finishes training, it is validated by having each user test against a set of classified users.

$$w_i = \frac{n}{kn_i} \tag{3.18}$$

4. RESULTS

The results show the performance of the system in determining strokes between child and adult users. F1-score is calculated across the entire data set as a means of evaluating system integrity. F1-score is the harmonic mean between precision and recall: precision measures the reliability of the model's results and recall measures how well the model can detect each class. Another indicator of performance, accuracy, provides less information and can give undesirable results. Consider the scenario where the data has a 90/10 split between the two classes. If the classifier only predicts the majority class, it will get 90% accuracy. However, analyzing the precision and recall of both classes shows that the minority class had a recall of zero. This phenomenon is called the accuracy paradox: the model gets good accuracy but had bad results.

The features derived from this data set were used to validate whether the system could distinguish children from adults. 10-fold cross-validation and leave-one-out cross-validation are commonly used techniques in classification problems to test the model on every part of the data set. 10-fold cross-validation randomly splits the data into ten parts, training ten models with a different part as the test set each time. Leave-one-out cross-validation splits the data such that each part contains data from a single user and uses the same train/test scheme as its 10-fold counterpart. It is important that any user who appears in the training set does not also appear in the test set to make the training and test sets completely independent. Therefore, leave-one-out cross-validation better suits this problem and is the used validation method.

In the study, the best two classifiers were Random Forest and Decision Tree.

Table 4 shows the scores for each classifier used in this work. The Random Forest algorithm is able to determine children and adults with a precision of 0.85, a recall of 0.85, and F1-score of 0.84. Using the Decision Tree algorithm, it had a precision of 0.85, a recall of 0.85, and F1-score of 0.86. Arguably, the Random Forest algorithm is many decision trees, but it is worth mentioning the top features from the top two classifiers. Figures 4.1 and 4.3 shows the confusion matrix for the Random Forest and the Decision Tree classifiers. The confusion matrix shows how many strokes were classified correctly and incorrectly in the data set. From this information, both classifiers had a better time classifying children strokes more than adult strokes. Figures 4.2 and 4.4 are a list of features that were important for both classifiers to differentiate between children and adults. Figures 4.1 and 4.2 are the top ten critical features for the Random Forest and Decision Tree classifiers. Top features were distinguished by how much the classifier can separate the data by category using that particular feature. The best three features for the Random Forest classifier were Long's density metrics: the diagonal length divided by the total length, total angle divided by the total length, and the log of the area of the stroke's bounding box. The best features for the Decision Tree classifier were the diagonal length divided by the total length, total angle divided by the total length, and Zernike's 8th moment.

Table 3. Machine Learning Algorithms			
Classifiers	F1-score	Precision	Recall
Zero-Rule	0.47	0.38	0.61
Naive Bayes	0.57	0.59	0.54
SVM	0.74	0.74	0.74
Random Forest	0.84	0.85	0.85
Decision Tree	0.86	0.85	0.85

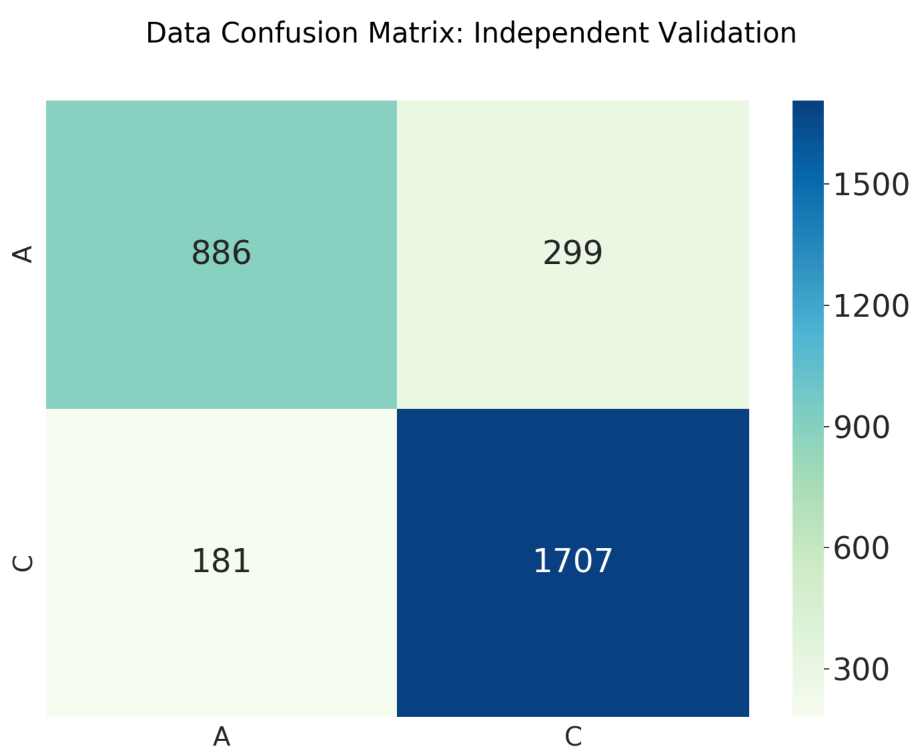

Figure 4.1: Confusion Matrix for the Random Forest Algorithm

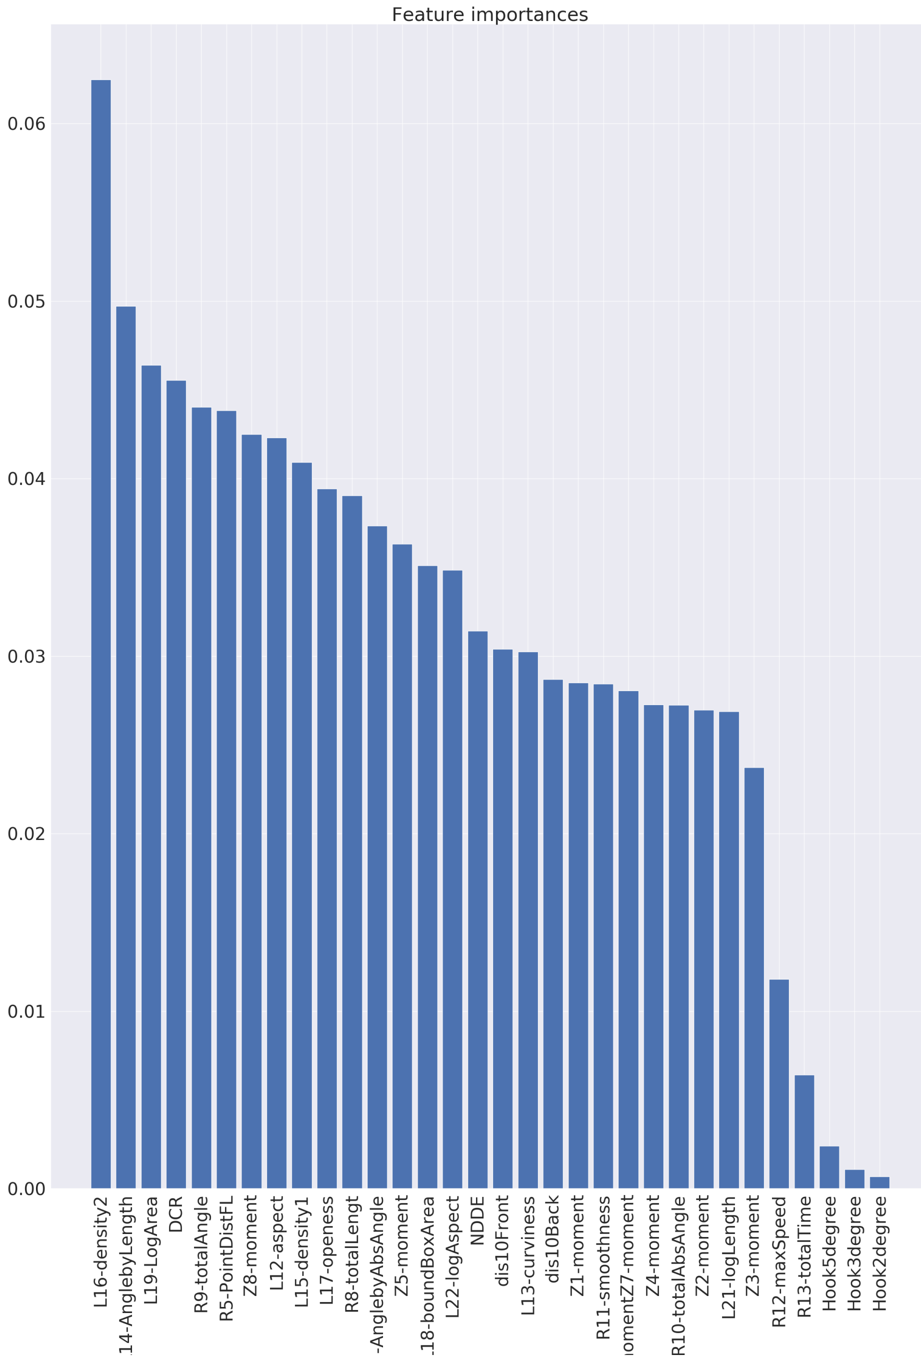

Figure 4.2: Important features selected of the Random Forest Algorithm

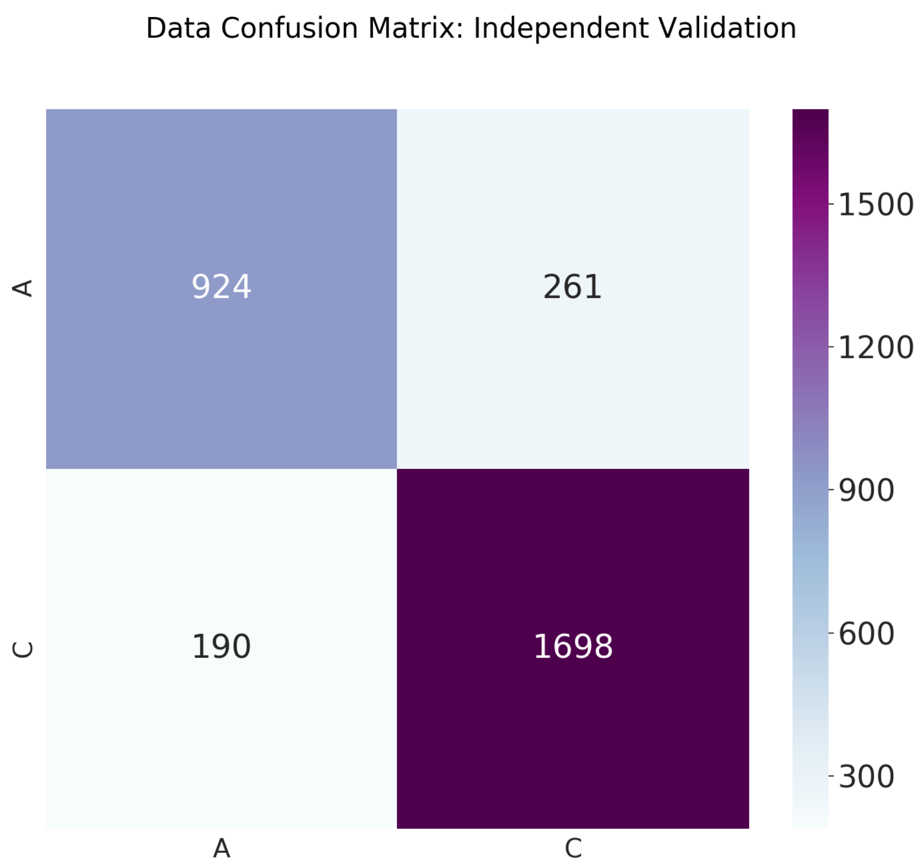

Figure 4.3: Confusion Matrix for the Decision Tree Algorithm

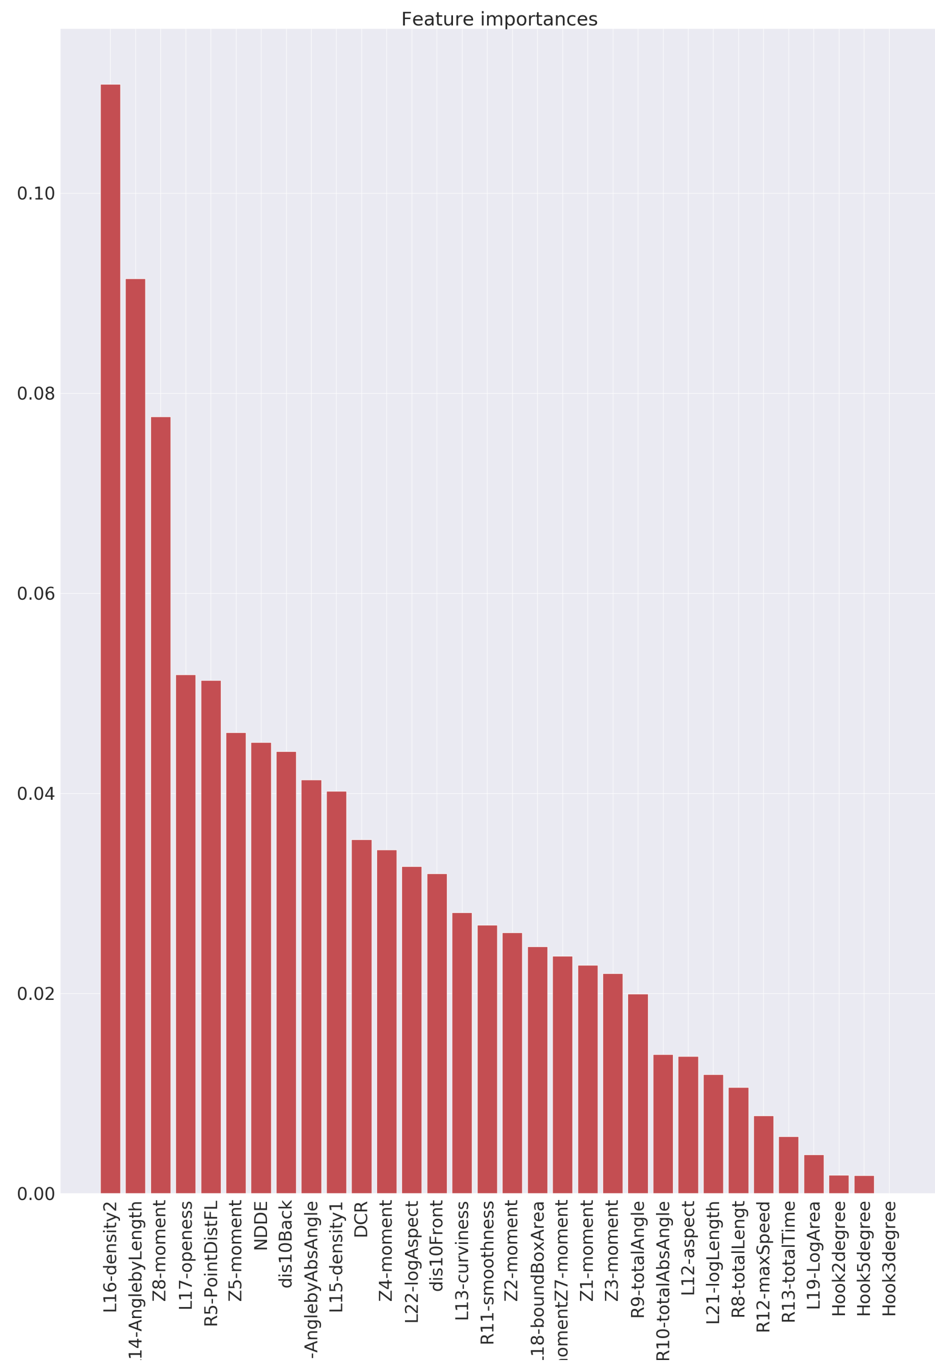

Figure 4.4: Important features selected for the Decision Tree Algorithm

1	The stroke length divided by the diagonal length of the bounding box
2	Total angle divided by length
3	Log of area
4	DCR
5	Total Angle
6	Distance between first and last point
7	Zernike moment 8
8	Aspect ratio
9	The distance between the first and last sampling points divided by stroke length
10	Openness

Table 4.1: Random Forest top 10 features

1	The stroke length divided by the diagonal length of the bounding box
2	Total angle divided by length
3	Zernike moment 8
4	Openness
5	Distance between first and last point
6	Zernike moment 5
7	NDDE
8	10% distance at the end of a stroke
9	Total angle divided by total absolute angle
10	The distance between the first and last sampling points divided by stroke length

Table 4.2: Decision Tree top 10 features

5. DISCUSSION

This section reviews the essential features that contributed to our system's decision-making models. In this work, the system employed techniques from two different modalities to extract possible features: sketch recognition and computer vision. The system determined that one of Long's density metrics is the most important feature for classification. Equation 5.1 shows Long's density metric as the stroke length divided by the diagonal length of the bounding box. In other words, the metric shows how far the user traveled in a single stroke relative to the overall stroke's size. We hypothesize that a child's sketch strokes are longer; when a child draws a house, the child maximizes the amount of space used. Adults' sketches are comparatively minimalistic, e.g., adults' strokes are precise and take less space. A pictorial demonstration of the density metric is in Figure 5.1.

$$\frac{\sum_{p=1}^{p-2} \sqrt{\delta x_p^2 + \delta y_p^2}}{\arctan(\frac{y_max - y_min}{x_max - x_min})} \tag{5.1}$$

The system found that the total angle divided by the total stroke length, seen in Equation 5.2, another one of Long's density metrics, is also a significant indicator. Figure 5.2 is a pictorial demonstration of the density metric, and it shows how much change in direction a stroke has. The values are more significant if the user sketches a spiral versus a straight line. We hypothesize that a child's sketches are more "spiral" than adults. That is, children draw curves and circles frequently, making spiraling motions as they sketch. While not widely explored, this behavior has been acknowledged as perhaps helping to indicate motor development differences in autistic children [27]. In conjunction with longer strokes, one can infer that children fill in their sketches using spirals to use as much space as possible.

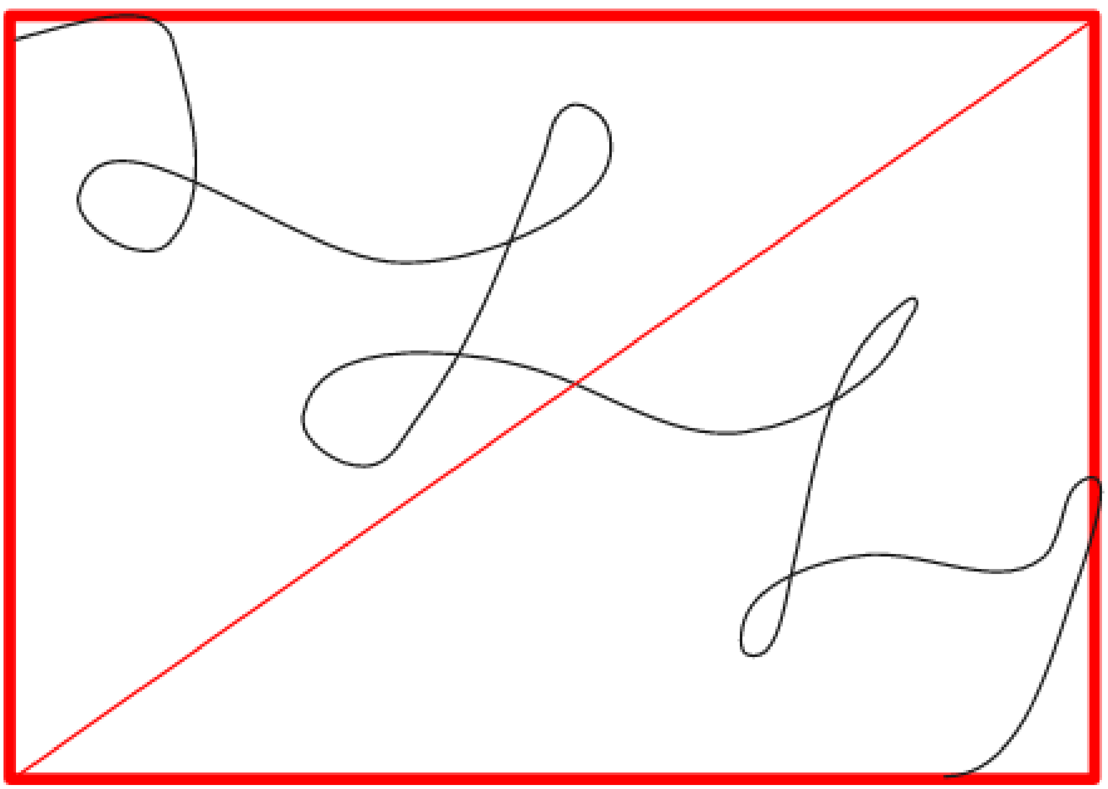

Figure 5.1: A demonstration of a bounding box and stroke length ratio

$$\sum_{p=1}^{p-2} \frac{\theta_p}{\sqrt{\delta x_p^2 + \delta y_p^2}} \tag{5.2}$$

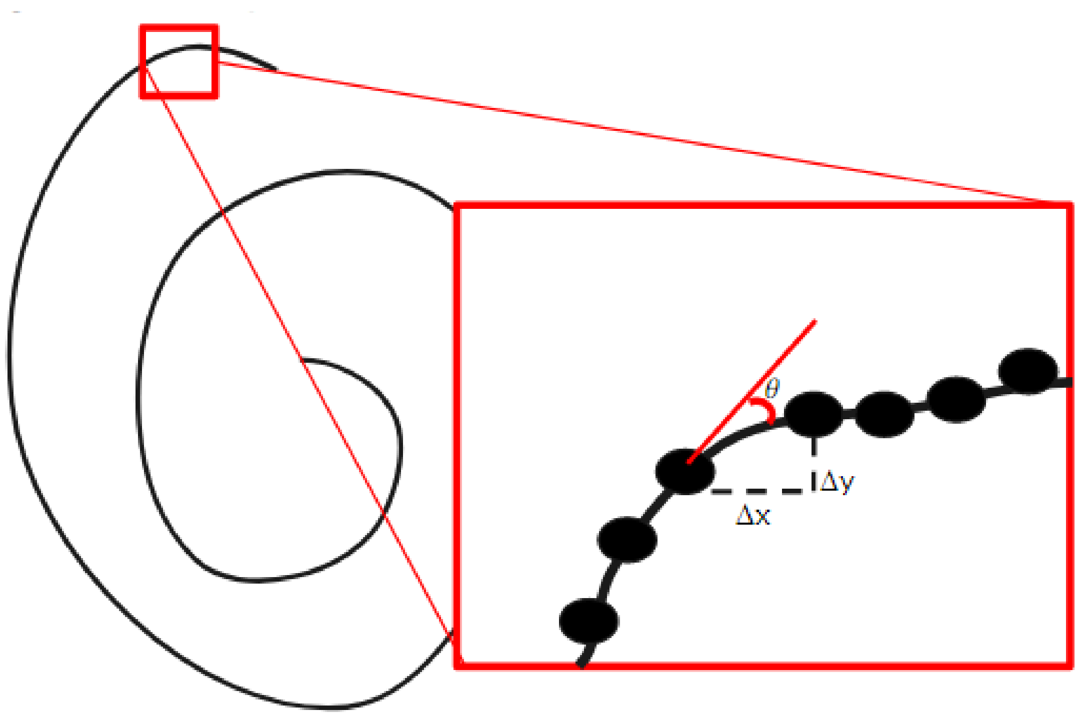

Figure 5.2: A demonstration of a spiral

The third most important feature was Zernike moment 8. Zernike moments are optical features that highlight different characteristics of an image based on how they would appear through a figurative "lens." In practice, a moment is calculated from Zernike polynomials, mathematical filters representing how light passes through varying lenses [47]. For example, the first moment is constant and does not produce any meaningful values. The second and third moments are different modifications of the first moment but are viewed as the vertical tilt (sine) and the horizontal tilt (cosine) shown respectively in Equation 5.3 and Equation 5.4, where p is the radial distance. Equation 5.5 is the equation for Zernike moment 8, which is a measure of the light bending along the horizontal axis.

Usually, Zernike moments are applied to images in computer vision techniques, but in this case, visual representations of individual strokes rather than full sketches were used. From single strokes, the results showed that adults' sketches align more along the horizontal axis than children's. One explanation may be that adults are more predisposed to axial alignment than children. These findings are also consistent with the previous discussion of children's tendency to sketch with spiraling behavior or in a way that fills more space. In both cases, Zernike moment 8 further defends that adults are relatively more strict and precise than children.

- **Zernike moment 2**: tilt (vertical tilt)

$$2psin\theta \tag{5.3}$$

- **Zernike moment 3**: tip (horizontal tilt)

$$2pcos\theta \tag{5.4}$$

- **Zernike moment 8**: horizontal coma; light bending in the horizontal axis

$$\sqrt{8}(3p^3 - 2p)cos\theta \tag{5.5}$$

Another important aspect of the drawings that was observed relates to the presence of "hooks" in the sketches. In the field of sketch recognition, hooks and tails refer to small, curved marks at the very beginning or ending of strokes, such as those left by putting the pen down a little before starting the next stroke while still moving from the previous one. From a recognition standpoint, it is beneficial to remove these artifacts [74]. However, as some preliminary work by Kim demonstrated [49], there may be some value in keeping them for other types of analysis. For example, Figure 5.3 shows that the participant created a hook at the end of the outline of an "A," and projecting along the direction of the hook shows that it clearly points to the beginning of the cross through the "A," as seen in Figure 5.4. Note that to clearly show order, line segments in Figure 5.4 have been colored green if they are in the first 10% of the beginning of a stroke and red if they are in the last 10% end of a stroke. The projected line is generated from the line of best fit formed between the final point in the stroke and the point of max curvature along the hook. We hypothesize this phenomenon is an artifact correlating with cognitive development of planning, or more precisely, as an indication of pre-planning. The ability to plan and understand the concept of future and consequences develops as children grow, so there is an assumption that pre-planning from participants could be helpful in future works in differentiating between adults and children.

From a sketch recognition standpoint, children often draw in erratic behavior when given a choice to draw without a template. Their sketches are inexperienced in that they may have low precision in shapes or lines. Adults are more precise and clean when they sketch. Adults would often want to create beautiful pictures with perfect lines and curves when drawing , which influences other works that beautify strokes [95,105]. Not only are these differences interesting from a childhood development perspective, e.g., measuring creativity when not using a template or building

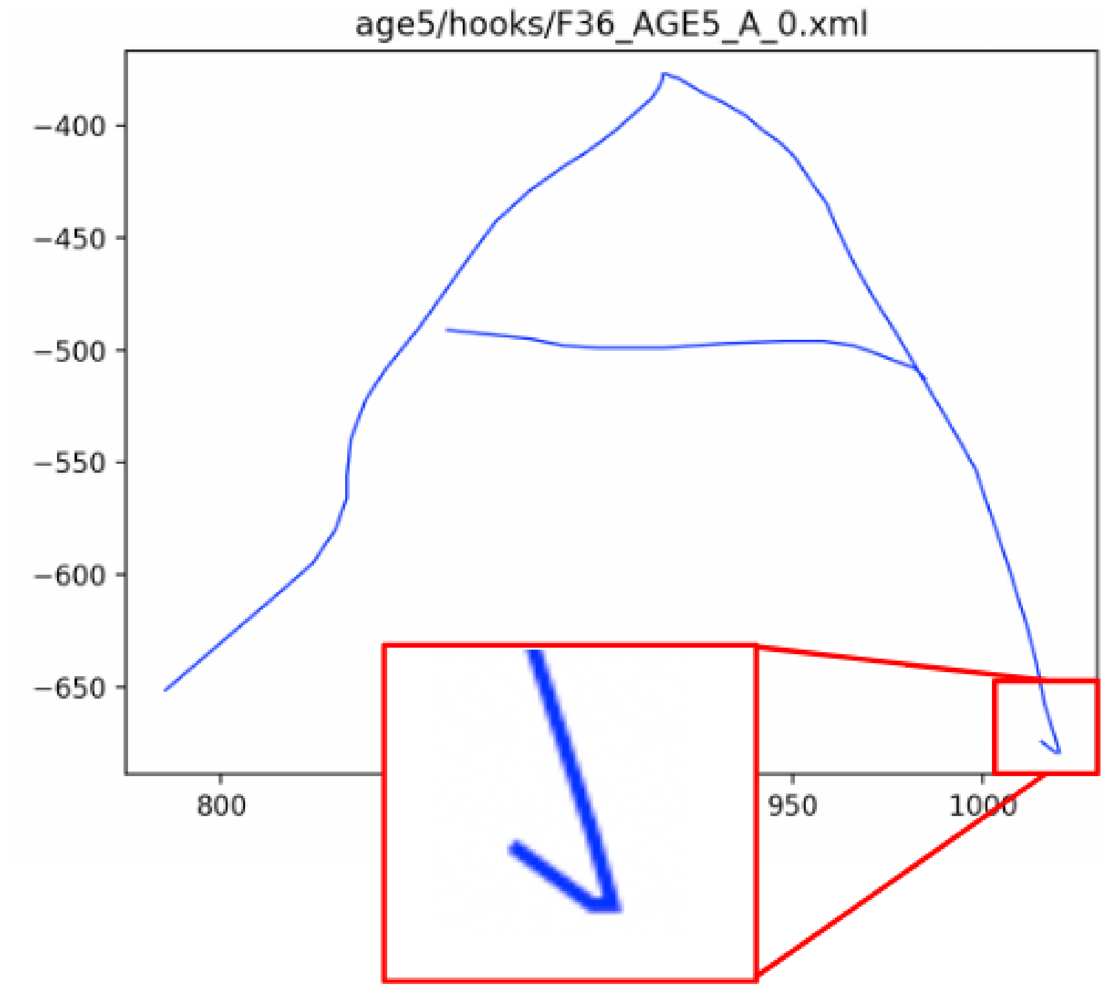

Figure 5.3: A closer look at a hook

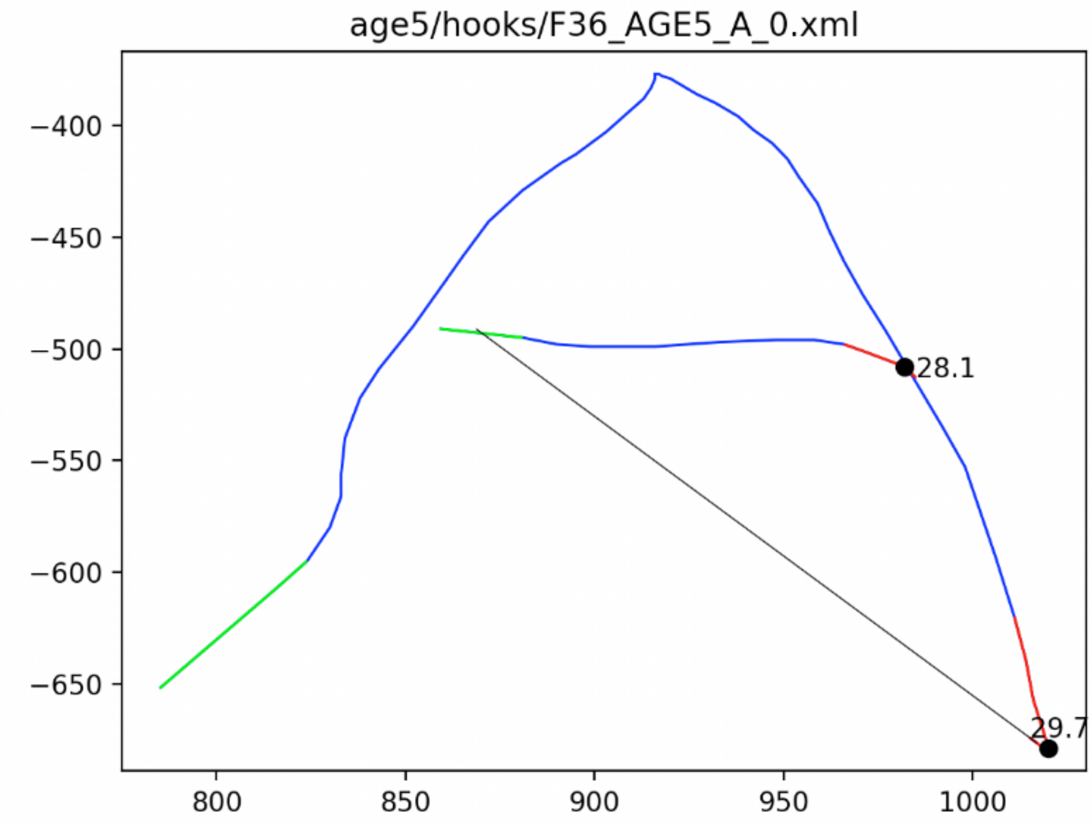

Figure 5.4: Hook leading to the next stroke

assessment tools based on cognitive behaviors, but they also support the primary goal of differentiating adults and children using sketch recognition principles.

6. FUTURE WORK

A goal of this project was to develop a system that can authenticate a child user on a device. The current system has an F1-score of 0.86 in differentiating between adults and children. These results should be sufficient to support a system that could serve as the basis of a more personalized and effective security framework. Most modern security systems require users to answer security questions, provide usernames and passwords, and/or complete a biometric scan to confirm that the machine is providing the correct information to the right person. This only works when people lock their computer, but many times they will leave their computer unlocked in their home, allowing children to have access to their parent's information. Our system could easily be extended to look at several strokes grouped together, creating an accurate and effective means of preventing such scenarios.

It would also be interesting to see if sketch recognition features can be used to differentiate between different age groups at the stroke level. Some preliminary work towards this goal was done by Kim, who was able to distinguish between children's fine motor skill developmental stages; however, work still needs to be done to do this at a stroke level and identify specific ages [49]. Systems built on this recognition could facilitate further personalization, allowing device interaction to be customized to specific age groups.

7. CONCLUSION

Children learn key concepts and abilities as they grow, which may impact their problem-solving ability, their academic performance, and, ultimately, their careers. It is imperative to evaluate children's development as they grow, and measuring healthy progression of motor control is one important metric. These skills are used in school-based activities, and children can fall behind without proper evaluation. Fortunately, there are ways to quantify the children writing academic achievements with the help of pediatricians. Using pediatricians can help evaluate a child's gross and fine motor skills that may affect how the child grows and learns. Since fine motor skills can contribute to reading, writing, crafting, and drawing, it is crucial to quantify the child's motor delay. However, the ratio between pediatricians and children is too high for pediatricians. This high ratio can prevent the assessment of children that need guidance. Screenings such as the "Ages and Stages Questionnaire" and the PEDS score form can help lower the ratio, but it does not shorten the duration or complexity of direct evaluations.

In this work, a system was trained to collect sketch data from children and adults to classify the sketches between the two users automatically. The results show that sketches can be successfully differentiated and that the sketch data can provide valuable information about the intentions of the user. There was a collection of $3,073$ strokes from 14 child participants and 25 adult participants on the Nexus 7 tablet. The study consisted of sketching anything they can imagine, tracing shapes, and lastly sketching from their imagination again while using each button. Sketch recognition features and vision techniques were extracted from the user's strokes. Once the set of features were calculated from each stroke, different models were constructed based on varying machine learning algorithms: Zero-Rule, Naive Bayes, Support Vector Machines, Random Forest, and Decision Tree. Out of all the classification machine learning algorithms, the Decision Tree performed the best. Using leave-one-out validation, the performance of the Decision Tree algorithm after validation was a total F1-score of 0.86 with high precision and recall.

Certain features were discussed that were important to the system derived from traditional sketch recognition features. The most important feature was Long's density metric, stroke length divided by bounding diagonal box length, which describes how space was used in a stroke. The second most important feature was Long's density metric, total angle divided by the stroke length, which shows the change in direction in the substrokes of a stroke. Lastly, Zernike moment 8 was an important feature. The intuition of Zernike moment 8 is that strokes drawn in the horizontal axis typically differ between children and adults.Hook features were also discussed that were previously used in other works. More work on these hook features can help in differentiating adults from children.

To conclude, there were two questions this study sought to address.

- *How accurately can a smart sketch interface differentiate adults and children using only single strokes from free-form sketches?*
- *What are the top three important features in determining adults and children?*

Regarding the first question, results showed that the system was able to differentiate adults and children with an F1-score of 0.86. By testing with leave-one-out cross-validation, the results correlate to the real world where the system would not be re-trained for new users. Not only does this verify findings from Kim's work [49] by categorizing age groups based on sketching features, but it also expands the results to free-form sketches and supports a wider number of applications.

In answer to the second question, the essential features were two of Long's density metrics—length divided by bounding box diagonal, total angle traversed divided by total stroke length—and Zernike Moment 8.

REFERENCES

[1] Christine Alvarado and Randall Davis. Resolving ambiguities to create a natural computer-based sketching environment. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence—Volume 2*, IJCAI'01, pages 1365–1371, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

[2] Christine Alvarado and Randall Davis. Dynamically constructed bayes nets for multi-domain sketch understanding. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, IJCAI'05, pages 1407–1412, San Francisco, CA, USA, 2005. Morgan Kaufmann Publishers Inc.

[3] Lisa Anthony. Physical dimensions of children's touchscreen interactions: Lessons from five years of study on the mtagic project. *International Journal of Human-Computer Studies*, 128:1–16, 2019.

[4] Lisa Anthony and Jacob O. Wobbrock. A lightweight multistroke recognizer for user interface prototypes. *Proceedings of Graphics Interface 2010. Canadian Information Processing Society*, 2010.

[5] Olufunmilola Atilola, Martin Field, Erin McTigue, Tracy Hammond, and Julie Linsey. Evaluation of a natural sketch interface for truss fbds and analysis. In *Frontiers in Education Conference (FIE)*, pages S2E–1 – S2E–6, Rapid City, SD, USA, October 2011. IEEE. ISBN: 978-1-61284-468-8.

[6] Olufunmilola Atilola, Stephanie Valentine, Hong-Hoe (Ayden) Kim, David Turner, Erin McTigue, Tracy Hammond, and Julie Linsey. Mechanix: A natural sketch interface tool for teaching truss analysis and free-body dia-

grams. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing (AIEDAM)*, 28-2:169–192, May 2014. ISSN: 1469-1760.

[7] Danilo Avola, Andrea Buono, Giorgio Gianforme, and Stefano Paolozzi. A novel recognition approach for sketch-based interfaces. In *Proceedings of the Fifteenth International Conference on Image Analysis and Processing*, ICIAP '09, pages 1015–1024, Berlin, Heidelberg, 2009. Springer-Verlag.

[8] Danilo Avola, Andrea Del Buono, Giorgio Gianforme, Stefano Paolozzi, and Rui Wang. Sketchml a representation language for novel sketch recognition approach. In *Proceedings of the Second International Conference on PErvasive Technologies Related to Assistive Environments*, PETRA '09, pages 31:1–31:8, New York, NY, USA, 2009. ACM.

[9] Diane Bricker, Jane Squires, Linda Mounts, L Potter, Robert Nickel, Elizabeth Twombly, and Jane Farrell. Ages and stages questionnaire. *Paul H. Brookes: Baltimore*, 1999.

[10] Florian Brieler and Mark Minas. A model-based recognition engine for sketched diagrams. *Journal on Visual Languages and Computing*, 21(2):81–97, April 2010.

[11] Klaus Broelemann, Xiaoyi Jiang, and Angela Schwering. Automatic understanding of sketch maps using context-aware classification. *Expert Systems and Applications*, 45(C):195–207, March 2016.

[12] Claire E Cameron, Laura L Brock, William M Murrah, Lindsay H Bell, Samantha L Worzalla, David Grissmer, and Frederick J Morrison. Fine motor skills and executive function both contribute to kindergarten achievement. *Child development*, 83(4):1229–1244, 2012.

[13] Jane Case-Smith. Fine motor outcomes in preschool children who receive occupational therapy services. *American Journal of Occupational Therapy*, 50(1):52–61, 1996.

[14] Sabine S Chebli and Marc J Lanovaz. Using computer tablets to assess preference for videos in children with autism. *Behavior analysis in practice*, 9(1):50–53, 2016.

[15] Isabelle D Cherney, Clair S Seiwert, Tara M Dickey, and Judith D Flichtbeil. Children's drawings: A mirror to their minds. *Educational psychology*, 26(1):127–142, 2006.

[16] Barbie Clarke and Siv Svanaes. An updated literature review on the use of tablets in education. *Tablets for Schools. UK: Family Kids & Youth*, 2014.

[17] Gennaro Costagliola, Mattia De Rosa, and Vittorio Fuccella. Local context-based recognition of sketched diagrams. *Journal of Visual Languages and Computing*, 25(6):955–962, December 2014.

[18] Leslie J Couse and Dora W Chen. A tablet computer for young children? exploring its viability for early childhood education. *Journal of research on technology in education*, 43(1):75–96, 2010.

[19] Danielle Cummmings, Francisco Vides, and Tracy 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*, pages 97–106. Eurographics Association, 2012.

[20] Christakis DA. Interactive media use at younger than the age of 2 years: Time to rethink the american academy of pediatrics guideline? *JAMA Pediatrics*, 168(5):399–400, 2014.

[21] Malcolm B Dick, Rodman W Shankle, Richard E Beth, Cordula Dick-Muehlke, Carl W Cotman, and Mary-Louise Kean. Acquisition and long-term retention of a gross motor skill in alzheimer's disease patients under constant and varied practice conditions. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 51(2):P103–P111, 1996.

[22] Daniel Dixon, Manoj Prasad, and Tracy 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)*, pages 897–906, Atlanta, GA, USA, April 10-15, 2010. ACM. ISBN: 978-1-60558-929-9.

[23] Daniel Dixon, Manoj Prasad, and Tracy 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*, pages 897–906, New York, New York, USA, April 2010. ACM Press.

[24] Maha El Meseery, Mahmoud Fakhr El Din, Samia Mashali, Magda Fayek, and Nevin Darwish. Sketch recognition using particle swarm algorithms. In *Proceedings of the Sixteenth IEEE International Conference on Image Processing*, ICIP'09, pages 1997–2000, Piscataway, NJ, USA, 2009. IEEE Press.

[25] Jorina Elbers and Andrew Macnab. The ages and stages questionnaires: feasibility of use as a screening tool for children in Canada. *Canadian Journal of Rural Medicine*, 13(1):9, 2008.

[26] Brian David Eoff and Tracy Hammond. Who dotted that ''i"?: context free user differentiation through pressure and tilt pen data. In *Proceedings of Graphics Interface 2009*, pages 149–156. Canadian Information Processing Society, 2009.

[27] Amanda Fleury. *Fractal dynamics of circle drawing in children with ASD*. PhD thesis, Citeseer, 2011.

[28] Luoting Fu and Levent Burak Kara. From engineering diagrams to engineering models: Visual recognition and applications. *Journal on Computer Aided Design*, 43(3):278–292, Mar. 2011.

[29] Christina T Fuentes, Stewart H Mostofsky, and Amy J Bastian. Children with autism show specific handwriting impairments. *Neurology*, 73(19):1532–1537, 2009.

[30] David Gaul and Johann Issartel. Fine motor skill proficiency in typically developing children: On or off the maturation track? *Human movement science*, 46:78–85, 2016.

[31] Leslie Gennari, Levent Burak Kara, Thomas F. Stahovich, and Kenji Shimada. Combining geometry and domain knowledge to interpret hand-drawn diagrams. *Journal of Computer Graphics*, 29(4):547–562, August 2005.

[32] Frances P Glascoe. Early detection of developmental and behavioral problems. *Pediatrics in Review*, 21(8):272–280, 2000.

[33] David Grissmer, Kevin J Grimm, Sophie M Aiyer, William M Murrah, and Joel S Steele. Fine motor skills and early comprehension of the world: two new school readiness indicators. *Developmental psychology*, 46(5):1008, 2010.

[34] A. Hall, C. Pomm, and P. Widmayer. A combinatorial approach to multi-domain sketch recognition. In *Proceedings of the Fourth Eurographics Workshop on Sketch-based Interfaces and Modeling*, SBIM '07, pages 7–14, New York, NY, USA, 2007. ACM.

[35] Tracy Hammond, Shalini Ashok Kumar, Matthew Runyon, Josh Cherian, Blake Williford, Swarna Keshavabhotla, Stephanie Valentine, Wayne Li, and Julie Linsey. It's not just about accuracy: Metrics that matter when modeling expert sketching ability. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 8 (19):1–47, July 2018.

[36] Tracy Hammond and Randall 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*, page 8 pages, Menlo Park, California, USA, July 28-August 1, 2002. AAAI.

[37] Tracy Hammond and Randall Davis. Automatically transforming symbolic shape descriptions for use in sketch recognition. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, AAAI'04, pages 450–456. AAAI Press, 2004.

[38] Tracy Hammond and Randall Davis. Ladder, a sketching language for user interface developers. *Computers & Graphics*, 29-4:518–532, August 2005.

[39] Tracy Hammond and Randall Davis. Creating the perception-based ladder sketch recognition language. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems (DIS)*, pages 141–150, Aarhus, Denmark, Denmarkl, August 16-20, 2010. ACM. ISBN: 978-1-4503-0103-9.

[40] Tracy Hammond, Krzysztof Gajos, Randall Davis, and Howard Shrobe. An agent-based system for capturing and indexing software design meetings. In *Proceedings of International Workshop on Agents In Design, WAID*, volume 2, page 18 pages, Cambridge, MA, USA, September 2002. MIT.

[41] Tracy Hammond and Brandon Paulson. Recognizing sketched multistroke primitives. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 1-4:1–34, October 2011. ISSN: 2160-6455.

[42] Tracy Hammond, Manoj Prasad, and Daniel Dixon. Art 101: Learning to draw through sketch recognition. In *Proceedings of 10th International Symposium on Smart Graphics, Lecture Notes In Computer Science 6133*, volume 633, pages 277–280, Banff, Canada, Canada, June 24-26, 2010. Springer Berlin Heidelberg. ISBN: 978-3-642-13543-9.

[43] Alina Hang. *Exploiting autobiographical memory for fallback authentication on smartphones*. PhD thesis, lmu, 2016.

[44] Christothea Herodotou. Young children and tablets: A systematic review of effects on learning and development. *Journal of Computer Assisted Learning*, 34(1):1–9, 2018.

[45] Monica Juneja, Mugdha Mohanty, Rahul Jain, and Siddarth Ramji. Ages and stages questionnaire as a screening tool for developmental delay in indian children. *Indian pediatrics*, 49(6):457–461, 2012.

[46] Jorien M Kerstjens, Arend F Bos, Elisabeth MJ ten Vergert, Gea de Meer, Phillipa R Butcher, Sijmen A Reijneveld, et al. Support for the global feasibility of the ages and stages questionnaire as developmental screener. *Early human development*, 85(7):443–447, 2009.

[47] Alireza Khotanzad and Yaw Hua Hong. Invariant image recognition by zernike moments. *IEEE Transactions on pattern analysis and machine intelligence*, 12(5):489–497, 1990.

[48] Carolyn R Kilday, Mable B Kinzie, Andrew J Mashburn, and Jessica V Whittaker. Accuracy of teacher judgments of preschoolers' math skills. *Journal of Psychoeducational Assessment*, 30(2):148–159, 2012.

[49] Hong-Hoe Kim. *A Fine Motor Skill Classifying Framework To Support Children's Self-Regulation Skills And School Readiness*. PhD thesis, Texas A&M, 2016.

[50] Hong-hoe Kim, Paul Taele, Jinsil Hwaryoung Seo, Jeffrey Liew, and Tracy Hammond. A novel sketch-based interface for improving children's fine motor skills and school readiness. In *Expressive '16: Proceedings of the Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*, Expressive '16, pages 69–78. Eurographics Association, May 2016.

[51] Hong-hoe Kim, Paul Taele, Stephanie Valentine, and Tracy Hammond. Developing intelligent sketch-based applications for children's fine motor sketching skill development. In *Proceedings of the 2014 International Conference on Intelligent User Interfaces Workshop on Sketch: Pen and Touch Recognition*, IUI SKETCH '14, February 2014.

[52] Hong-Hoe (Ayden) Kim, Stephanie Valentine, Paul Taele, and Tracy Hammond. Easysketch: A sketch-based fine motor skill recognizing educational interface for children emerging technology research strand. In A. Adler T. Hammond, S. Valentine and M. Payton, editors, *The Impact of Pen and Touch Technology on Education*, Human-Computer Interaction Series, chapter 4, pages 35–46. Springer, Switzerland, 2015.

[53] Honghoe Kim, Paul Taele, Stephanie Valentine, Erin McTigue, and Tracy Hammond. KimCHI: a sketch-based developmental skill classifier to enhance

pen-driven educational interfaces for children. In *Proceedings of the International Symposium on Sketch-Based Interfaces and Modeling*, pages 33–42. ACM, 2013.

[54] Johann P Kuhtz-Buschbeck, Birgit Hoppe, Mukaddes Gölge, Mona Dreesmann, Ute Damm-Stünitz, and Annegret Ritz. Sensorimotor recovery in children after traumatic brain injury: analyses of gait, gross motor, and fine motor skills. *Developmental medicine and child neurology*, 45(12):821–828, 2003.

[55] Alexis R Lauricella, Ellen Wartella, and Victoria J Rideout. Young children's screen time: The complex role of parent and child factors. *Journal of Applied Developmental Psychology*, 36:11–17, 2015.

[56] PLD Lee and Y Becher. Orientation on the east asia-pacific early child development scales (eap-ecds). In *Asia-Pacific Regional Early Childhood Development (ECD) Conference*. ARNEC (Asia-Pacific Regional Network for Early Childhood) Secretariat., 2017.

[57] Sandy C Li, Jacky WC Pow, Emily ML Wong, and Alex CW Fung. Empowering student learning through tablet pcs: A case study. *Education and Information Technologies*, 15(3):171–180, 2010.

[58] A Chris Long Jr, James A Landay, Lawrence A Rowe, and Joseph Michiels. Visual similarity of pen gestures. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 360–367. ACM, 2000.

[59] Sébastien Macé and Eric Anquetil. Eager interpretation of on-line hand-drawn structured documents: The dali methodology. *Journal on Pattern Recognition*, 42(12):3202–3214, December 2009.

[60] W Marder and G Gaumer. Reexamination of the adequacy of physician supply made in 1980 by the graduate medical education national advisory committee (gmenac) for selected specialties. *Abt Associates, Inc*, 1991.

[61] Clarence C McCormick, Janice Nelson Schnobrich, S Willard Footlik, and Betty Poetker. Improvement in reading achievement through perceptual-motor training. *Research Quarterly. American Association for Health, Physical Education and Recreation*, 39(3):627–633, 1968.

[62] Dana Charles McCoy, Evan D Peet, Majid Ezzati, Goodarz Danaei, Maureen M Black, Christopher R Sudfeld, Wafaie Fawzi, and Günther Fink. Early childhood developmental status in low-and middle-income countries: national, regional, and global prevalence estimates using predictive modeling. *PLoS Medicine*, 13(6):e1002034, 2016.

[63] Hannelore Montrieux, Ruben Vanderlinde, Tammy Schellens, and Lieven De Marez. Teaching and learning with mobile technology: A qualitative explorative study about the introduction of tablet devices in secondary education. *PloS one*, 10(12):e0144008, 2015.

[64] Holly Carrell Moore and Jennifer Keys Adair. "i'm just playing ipad": Comparing prekindergarteners' and preservice teachers' social interactions while using tablets for learning. *Journal of Early Childhood Teacher Education*, 36(4):362–378, 2015.

[65] Shahzad Nabeel, Brandon Paulson, and Tracy Hammond. Urdu qaeda: Recognition system for isolated urdu characters. In *Proceedings of the Workshop on Sketch Recognition at the 14th International Conference of Intelligent User Interfaces (IUI)*, page 4 pages, Sanibel, FL, USA, February 8-11, 2009. ACM.

[66] Michelle M Neumann. An examination of touch screen tablets and emergent literacy in australian pre-school children. *Australian Journal of Education*, 58(2):109–122, 2014.

[67] Mizuki Oka, Kazuhiko Kato, Yingqing Xu, Lin Liang, and Fang Wen. Scribble-a-secret: Similarity-based password authentication using sketches. In *2008 19th International Conference on Pattern Recognition*, pages 1–4. IEEE, 2008.

[68] Committee on Children with Disabilities et al. Developmental surveillance and screening of infants and young children. *Pediatrics*, 108(1):192–195, 2001.

[69] Sally Ozonoff, Gregory S Young, Stacy Goldring, Laura Greiss-Hess, Adriana M Herrera, Joel Steele, Suzanne Macari, Susan Hepburn, and Sally J Rogers. Gross motor development, movement abnormalities, and early identification of autism. *Journal of autism and developmental disorders*, 38(4):644–656, 2008.

[70] Rachel Patel, Beryl Plimmer, John Grundy, and Ross Ihaka. Ink features for diagram recognition. In *Proceedings of the Fourth Eurographics Workshop on Sketch-based Interfaces and Modeling*, SBIM '07, pages 131–138, New York, NY, USA, 2007. ACM.

[71] Brandon Paulson, Brian Eoff, Aaron Wolin, Joshua Johnston, and Tracy Hammond. Sketch-based educational games: "drawing" kids away from traditional interfaces. In *Proceedings of the Seventh International Conference on Interaction Design and Children*, IDC '08, pages 133–136, New York, NY, USA, 2008. ACM.

[72] Brandon Paulson and Tracy Hammond. Recognizing and Beautifying Low-level Sketch Shapes with Two New Features and Ranking Algorithm. In *Twentieth*

Annual ACM Symposium on User Interface Software and Technology Posters, number September in IUI '08, pages 1–10. ACM, 2007.

[73] Brandon Paulson and Tracy 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)*, page 2 pages, Newport Rhode Island, USA, October 7-10 2007. ACM.

[74] Brandon Paulson and Tracy Hammond. Paleosketch: Accurate primitive sketch recognition and beautification. In *Proceedings of the 13th International Conference on Intelligent User Interfaces (IUI)*, page 10 pages, Canary Islands, Spain, Spain, January 13-16, 2008. ACM. ISBN: 978-1-59593-987-6.

[75] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python . *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[76] Florian Perteneder, Martin Bresler, Eva-Maria Grossauer, Joanne Leong, and Michael Haller. cluster: Smart clustering of free-hand sketches on large interactive surfaces. In *Proceedings of the Tewnty-Eighth Annual ACM Symposium on User Interface Software & Technology*, UIST '15, pages 37–46, New York, NY, USA, 2015. ACM.

[77] Janet C Read and P Markopoulos. Child-computer interaction. *International Journal of Child-Computer Interaction*, 1(1):2–6, 2013.

[78] Dean Rubine. Specifying gestures by example. *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*, 1991.

[79] Karen Rust, Meethu Malu, Lisa Anthony, and Leah Findlater. Understanding child-defined gestures and children's mental models for touchscreen tabletop interaction. In *Proceedings of the 2014 Conference on Interaction Design and Children*, IDC '14, pages 201–204, New York, NY, USA, 2014. ACM.

[80] David Rydz, Myriam Srour, Maryam Oskoui, Nancy Marget, Mitchell Shiller, Rena Birnbaum, Annette Majnemer, and Michael I Shevell. Screening for developmental delay in the setting of a community pediatric clinic: a prospective assessment of parent-report questionnaires. *Pediatrics*, 118(4):e1178–e1186, 2006.

[81] Perihan Savas. Tablet pcs as instructional tools in english as a foreign language education. *Turkish Online Journal of Educational Technology-TOJET*, 13(1):217–222, 2014.

[82] Luisa Schonhaut, Iván Armijo, Marianne Schönstedt, Jorge Alvarez, and Miguel Cordero. Validity of the ages and stages questionnaires in term and preterm infants. *Pediatrics*, 131(5):e1468–e1474, 2013.

[83] Alex Shaw. Human-centered recognition of children's touchscreen gestures. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, ICMI '17, pages 638–642, New York, NY, USA, 2017. ACM.

[84] Alex Shaw and Lisa Anthony. Analyzing the articulation features of children's touchscreen gestures. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, ICMI '16, pages 333–340, New York, NY, USA, 2016. ACM.

[85] Alex Shaw and Lisa Anthony. Toward a systematic understanding of children's touchscreen gestures. In *Proceedings of the 2016 CHI Conference Extended*

Abstracts on Human Factors in Computing Systems, CHI EA '16, pages 1752–1759, New York, NY, USA, 2016. ACM.

[86] Alex Shaw, Jaime Ruiz, and Lisa Anthony. Comparing human and machine recognition of children's touchscreen stroke gestures. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, ICMI '17, pages 32–40, New York, NY, USA, 2017. ACM.

[87] Ajay Singh, Chia Jung Yeh, and Sheresa Boone Blanchard. Ages and stages questionnaire: a global screening scale. *Boletín Médico Del Hospital Infantil de México (English Edition)*, 74(1):5–12, 2017.

[88] Jane Squires, Diane Bricker, Elizabeth Twombly, Robert Nickel, Jantina Clifford, Kimberly Murphy, Robert Hoselton, LaWanda Potter, Linda Mounts, and Jane Farrell. *Ages & Stages Questionnaires: A Parent-Completed Child Monitoring System*. Brookes Publishing, Baltimore, MD, USA, third edition, jun 2009.

[89] Roger A Stewart, Audrey C Rule, and Debra A Giordano. The effect of fine motor skill activities on kindergarten student attention. *Early Childhood Education Journal*, 35(2):103–109, 2007.

[90] Paul Taele, Laura Barreto, and Tracy Hammond. Maestoso: An intelligent educational sketching tool for learning music theory. In *The Twenty-Seventh Annual Conference on Innovative Applications of Artificial Intelligence at AAAI (IAAI 2015)*, page 7 pages, Austin, Texas, USA, January 27-29, 2015. AAAI.

[91] Paul Taele and Tracy Hammond. Lamps: A sketch recognition-based teaching tool for mandarin phonetic symbols i. In *Journal of Visual Languages & Computing (JLVC)*, volume 21-2, pages 109–120. Elsevier, April 2010.

[92] Eugene M Taranta and Joseph J LaViola Jr. Math boxes: A pen-based user interface for writing difficult mathematical expressions. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*, pages 87–96. ACM, 2015.

[93] Anna H Tielsch and Patricia Jackson Allen. Listen to them draw: screening children in primary care through the use of human figure drawings. *Pediatric nursing*, 31(4), 2005.

[94] David R Topor, Susan P Keane, Terri L Shelton, and Susan D Calkins. Parent involvement and student academic performance: A multiple mediational analysis. *Journal of prevention & intervention in the community*, 38(3):183–197, 2010.

[95] Stephanie Valentine, Francisco Vides, George Lucchese, David Turner, Honghoe Kim, Wenzhe Li, Julie Linsey, and Tracy Hammond. Mechanix: A sketch-based tutoring system for statics courses. In *Twenty-Fourth IAAI Conference*, 2012.

[96] Stephanie Valentine, Francisco Vides, George Lucchese, David Turner, Hong-Hoe (Ayden) Kim, Wenzhe Li, Julie Linsey, and Tracy Hammond. Mechanix: A sketch-based tutoring and grading system for free-body diagrams. *AI Magazine*, 34-1:55–66, January 2013.

[97] Roshanak Vameghi, Firoozeh Sajedi, Adis Kraskian Mojembari, Abbas Habiollahi, Hamid Reza Lornezhad, and Bahram Delavar. Cross-cultural adaptation, validation and standardization of ages and stages questionnaire (asq) in iranian children. *Iranian journal of public health*, 42(5):522, 2013.

[98] Deborah Lowe Vandell, Jay Belsky, Margaret Burchinal, Laurence Steinberg, Nathan Vandergrift, and NICHD Early Child Care Research Network. Do

effects of early child care extend to age 15 years? results from the nichd study of early child care and youth development. *Child development*, 81(3):737–756, 2010.

[99] Lisa Anthony Vatavu, Radu-Daniel and Jacob O. Wobbrock. Gestures as point clouds: a $ p recognizer for user interface prototypes. *Proceedings of the 14th ACM international conference on Multimodal interaction*, 2012.

[100] Radu-Daniel Vatavu, Lisa Anthony, and Quincy Brown. Child or adult? inferring smartphone users' age group from touch measurements alone. In Abascal J., Barbosa S., Fetter M., Gross T., Palanque P., and Winckler M., editors, *Lecture Notes in Computer Science*, volume 9299 of *1*, pages 1–9. Springer, Cham, 2015.

[101] Francisco Vides, Paul Taele, Hong-hoe Kim, Jimmy Ho, and Tracy Hammond. Intelligent feedback for kids using sketch recognition. In *Proceedings of the ACM SIGCHI 2012 Conference on Human Factors in Computing Systems Workshop on Educational Interfaces, Software, and Technology*, CHI EIST '12, May 2012.

[102] Robert G Voigt, Antolin M Llorente, Craig L Jensen, J Kennard Fraley, William J Barbaresi, and William C Heird. Comparison of the validity of direct pediatric developmental evaluation versus developmental screening by parent report. *Clinical pediatrics*, 46(6):523–529, 2007.

[103] Blake Williford, Matthew Runyon, Adil Malla, Wayne Li, Julie Linsey, and Tracy Hammond. Zensketch: A sketch-based game for improving line work. In *The ACM SIGCHI Annual Symposium on Computer-Human Interaction in Play (CHI PLAY), Student Game Design Competition*, pages 591–598, Amster-

dam, The Netherlands, October 15-18, 2017. ACM. ISBN: 978-1-4503-5111-9, DOI: 10.1145/3130859.3130861.

[104] Blake Williford, Matthew Runyon, Adil Hamid Malla, Wayne Li, Julie Linsey, and Tracy Hammond. Zensketch: A sketch-based game for improving line work. In *Extended Abstracts Publication of the Annual Symposium on Computer-Human Interaction in Play*, pages 591–598. ACM, 2017.

[105] Blake Williford, Paul Taele, Trevor Nelligan, Wayne Li, Julie Linsey, and Tracy Hammond. Persketchtivity: an intelligent pen-based educational application for design sketching instruction. In *Revolutionizing Education with Digital Ink*, pages 115–127. Springer, 2016.

[106] Jacob O Wobbrock, Andrew D Wilson, and Yang Li. Gestures without libraries, toolkits or training: a $1 recognizer for user interface prototypes. In *Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 159–168. ACM, 2007.

[107] Julia Woodward, Alex Shaw, Aishat Aloba, Ayushi Jain, Jaime Ruiz, and Lisa Anthony. Tablets, tabletops, and smartphones: Cross-platform comparisons of children's touchscreen interactions. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, ICMI '17, pages 5–14, New York, NY, USA, 2017. ACM.

[108] Julia Woodward, Alex Shaw, Annie Luc, Brittany Craig, Juthika Das, Phillip Hall, Jr., Akshay Holla, Germaine Irwin, Danielle Sikich, Quincy Brown, and Lisa Anthony. Characterizing how interface complexity affects children's touchscreen interactions. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 1921–1933, New York, NY, USA, 2016. ACM.

[109] Cheryl Wright, Marissa Diener, and Susan C Kay. School readiness of low-income children at risk for school failure. *Journal of Children and Poverty*, 6(2):99–117, 2000.

[110] Erelcan Yanik and Tevfik Metin Sezgin. Active learning for sketch recognition. *Journal of Computers & Graphics*, 52(C):93–105, November 2015.

[111] Bo Yu. Recognition of freehand sketches using mean shift. In *Proceedings of the Eighth International Conference on Intelligent User Interfaces*, IUI '03, pages 204–210, New York, NY, USA, 2003. ACM.

[112] Jane C Zahn. Differences between adults and youth affecting learning. *Adult Education*, 17(2):67–77, 1967.

[113] Lisha Zhang and Zhengxing Sun. An experimental comparison of machine learning for adaptive sketch recognition. *Applied Mathematics and Computation*, 185(2):1138–1148, February 2007.