A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input

***Abstract*—Gesture recognition algorithms help designers create intelligent user interfaces for a number of application areas. However, these recognition algorithms are usually designed to recognize the gestures of adults, not children, and as such they generally do not perform as well for children as adults. Recognition of younger children's gestures is particularly poor when compared to recognition of older children's and adults' gestures. Researchers have begun to examine the aspects of children's gesture articulation patterns that make recognition difficult. This paper extends the initial work examining child-specific recognition approaches by considering general purpose approaches and how they might apply to the problem of recognizing children's touchscreen gestures. This paper presents a survey of existing recognition and analysis techniques for gestures of both adults and children from a human-centered perspective, highlighting ways in which improved recognition can lead to a better experience for children using touchscreen gestures in a variety of contexts.**

***Index Terms*—children, drawing, gesture, learning, recognition**

I. Introduction

CHILDREN are increasingly using touchscreen devices in a variety of contexts (Common Sense Media, 2013). Prior work investigating children's touchscreen interactions has shown clear benefits of designing touchscreen applications based on children's specific interaction patterns, which are largely distinct from those of adults (Anthony et al., 2012a; Arif and Sylla, 2013; Hiniker et al., 2015; Vatavu et al., 2015b). In particular, children's gesture interactions are quite different from those of adults. For example, children tend to use more ink, produce larger gestures, and take longer to create their gestures than adults (Shaw and Anthony, 2016a). We use the term 'gesture' to refer to a series of one or more single-finger strokes on a touchscreen to produce a shape, letter, number, or other symbol, as demonstrated by Figure 1. This definition is consistent with a large body of previous work (Anthony et al., 2013b; Rubine, 1991; Vatavu et al., 2013a, 2013b, 2012; Woodward et al., 2016). Because the recognition of these single-touch gestures is a fundamentally different problem than recognizing other types of gestures, we focus only on single-touch gestures in this work. However, we believe that many of our conclusions are also applicable to other types of gestures produced by children.

Gesture recognition algorithms have achieved accuracy rates as high as 99% for adults (Alimoglu and Alpaydin, 1997; Cho, 2006; Olsen et al., 2007; Taranta II and LaViola Jr., 2015; Vatavu et al., 2012), but rates are much lower for children. Woodward et al. (Woodward et al., 2016) report rates as low as 64% for 5-year-old children. Children are less consistent than adults in their gesture interactions in a number of ways (Anthony et al., 2013b; Shaw and Anthony, 2016a), but there has been little work examining why recognition rates are so much lower for children's gestures than those of adults, or how

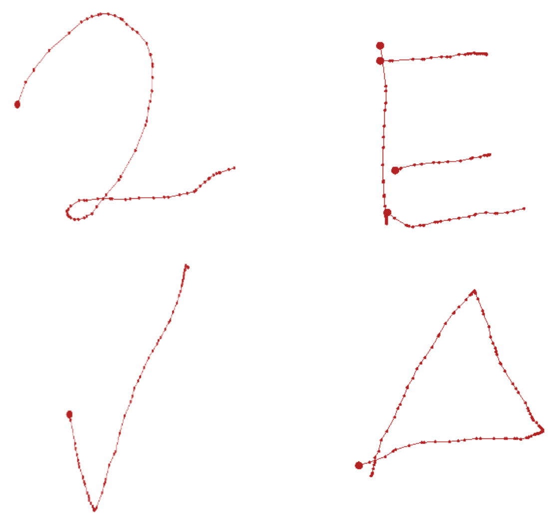

Figure 1. Four examples of touchscreen gestures. These gestures include the number "2", the letter "E", the "checkmark" symbol, and the "triangle" shape, and were all produced by the same 10-year-old participant in Woodward et Al.'s (Woodward et al., 2016) study of interface complexity.

to improve them. The low recognition rates point to the need for additional work examining the behaviors exhibited by children when making these gestures.

The primary focus of this paper is to examine children's touchscreen gestures, which we accomplish by conducting a survey to learn about specific interaction patterns and children's use of touchscreen technology. We also identify open areas for improving gesture recognition. Much of the work done with children's touchscreen gestures to this point has adapted prior work on adults' gestures (Anthony, 2019; Kim et al., 2013; Shaw and Anthony, 2016a; Woodward et al., 2017, 2016). To motivate future work on children's gestures, it is important to understand existing work on adults' gestures and what aspects have and have not been applied to children's gestures. Furthermore, in this survey we discuss what can be learned from examining children's gestures, both in terms of specific interaction patterns and more general observations about children's use of touchscreen technology. Toward this end, in this survey, we present an examination of the field of touchscreen gesture recognition and analysis for both children and adults through the lens of how it is relevant to children's gesture interactions.

Later in this paper, we also present examples of ways in which gesture recognition is being used in educational technology. These examples help motivate continued work on recognizing children's gestures as a means of improving the state of these pedagogical systems. An example of such an educational system is the iOS application *abc PocketPhonics*

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Source	Count
ACM Digital Library	69
Google Scholar	42
IEEE Explore	13
Other (textbook, app, etc.)	10
TOTAL	**134**

Table 1. A breakdown of the sources of the papers included in this survey.

("abc PocketPhonics: letter sounds & writing + first words," 2018), which uses recognition to evaluate children's gesture input. Based on the result, the system provides feedback on how the child can improve their writing. Having the ability to recognize the children's gestures enables the system to give more meaningful guidance to the child.

While there have been many surveys on gesture recognition, few have focused on recognizing children's gestures. In particular, there have been numerous surveys on 3D hand, face, and body gesture recognition (Cheng et al., 2016; Konstantinos G. Derpanis, 2004; Mitra and Acharya, 2007; Murthy and Jadon, 2009; Ravindran, 2010), but very few addressing touchscreen gestures. Zhai et al.'s (Zhai et al., 2012) survey concerns touchscreen gestures, but focuses on issues regarding the design of the gestures, only briefly discussing recognition. Olsen et al. (Olsen et al., 2009) briefly discuss stroke gesture recognition, though their survey focuses on approaches to sketch-based modeling, which involves the creation and recognition of more complex drawings (such as circuit diagrams for engineering courses), which is less necessary for children's applications. Additionally, none of the surveys mentioned above have focused on children's gestures. We fill this gap in the gesture interaction literature by explicitly examining children's gestures and by surveying stroke gesture recognition methods from the perspective of how they can be applied to children's gestures.

We found the papers selected for this survey using the following methodology: we searched the ACM Digital Library, Google Scholar, and IEEE Explore for "touchscreen gesture," "touchscreen gesture recognition," "gesture recognition children," and "children touchscreen gesture." We examined papers that were returned and include all those that concerned touchscreen gesture recognition (not necessarily for children) or children's interactions with touchscreens. We also examined relevant papers cited by the papers we found in our search using the same criteria. Table 1 shows a summary of the number of papers from each source.

The outline for the rest of this paper is as follows: Section II provides an overview of the types of gesture-based applications commonly used by children as well as the potential benefits of using those applications. Section III discusses children's touchscreen gestures, including recognition and classification of those gestures. In Section IV, we discuss types of recognizers and offer a comparison between the different categories. Section V concerns gesture features and analyses researchers use to investigate the factors that influence recognition. Finally, Section VI discusses open areas for future work. Figure 2 shows a glossary of important terms that the reader may find useful.

- *Feature* – Any quantitative measure of some aspect of a gesture.
- *Articulation Feature* – A measure used to quantify some aspect of the way a user creates a gesture.
- *Bounding Box* – The smallest rectangle with vertical and horizontal sides that can completely enclose a gesture.
- *Elicitation* – An experimental technique by which gesture data is collected from users.
- *Gesture* – A series of one or more strokes on a touchscreen to make a letter, number, shape or symbol.
- *Guessability* – The quality of a gesture that allows a user to know what the gesture refers to without explicit familiarity with the gesture.
- *Referent* – The action or symbol that a gesture is intended to represent.
- *Training set* – The gestures used to define the model that is used to recognize test gestures.
- *User-dependent* – A type of gesture recognition experiment in which the recognizer is trained on gestures from the same user it is tested on.
- *User-independent* – A type of gesture recognition experiment in which the recognizer is trained on gestures from different users than it is tested on.

Figure 2. A glossary of important terms from our paper.

II. GESTURE BASED APPLICATIONS

Touchscreen applications have the potential to provide children with experiences that help them gain real world knowledge and problem-solving skills through intuitive interactions (Lovato and Waxman, 2016). In this section, we begin by discussing some of these potential cognitive and developmental benefits of using touchscreen applications, particularly those that employ gesture recognition. We also provide further context by discussing several types of applications commonly used by children that often employ gesture recognition. In particular, we discuss the use of gestures in games and educational applications. It is important to note that while we provide examples of existing gesture-based applications for children, most such systems typically support limited gesture sets and highly constrained domains. Thus, there is still much room for improving these systems. Better recognition will both enhance existing systems and enable the creation of new systems that were not previously possible.

As a motivational example, consider the following scenario. A child is using an application designed to teach young children to draw shapes like circles and squares. The system prompts the child to draw a diamond, but instead she draws a square. The recognition algorithm, which is designed to account for common idiosyncrasies in children's gestures, classifies the drawing as a square and offers feedback, showing a guide on how to draw the diamond. This example shows a clear instance in which improved recognition rates can lead to adaptive

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

scaffolding to support children's learning, particularly in the context of learning to draw shapes, letters, numbers, and symbols.

A. Benefits of Recognition Based Applications

The use of touchscreen devices and computers in general can offer cognitive benefits in many contexts. While parents and educators may worry about the harmful effects of children spending too much time using computers, several prior studies have shown that children can learn valuable real-world knowledge and problem-solving skills when using certain applications (Flewitt et al., 2015; Huber et al., 2016; Lauricella et al., 2010). For example, Lauricella et al. (Lauricella et al., 2010) showed how computers could effectively convey information to children in a study in which children were shown the location of hidden objects with a digital 3D representation of a room, then asked to retrieve the objects in the actual room. Children performed better when shown the digital representation in an interactive mouse-and-keyboard game than when shown a video of the objects being hidden. Flewitt et al. (Flewitt et al., 2015) showed the positive effects of supplementing schooling with iPad applications. The authors note that children using iPad applications as part of their schooling achieved measurable gains in literacy over a two-month period. Huber et al. (Huber et al., 2016) examined 4- to 6-year-old children's ability to learn from touchscreen applications by using an application to teach them to solve a variant of the Towers of Hanoi problem, then having them solve the problem in the real world. The interactive nature of touchscreen devices provides the opportunity for a system to provide feedback based on a child's interactions that would normally require oversight by a skilled adult. The results of the study show that children are able to transfer problem-solving skills from their interaction with both traditional mouse-and-keyboard input and touchscreen devices to real world scenarios. Based on this body of literature, we can see that interactions with computers, including touchscreen devices, have the potential to offer cognitive and developmental benefits to children who use them. In this paper, we discuss ways in which gesture interaction can be used in applications that provide these kinds of benefits. Our goal in this work is to motivate one method by which these recognition-based systems can be improved. We now discuss two of the most common types of recognition-based applications for children: games and educational systems.

B. Games

A number of touchscreen applications for children use gesture recognition. In particular, mobile games are often used by children, and many of these applications make use of touchscreen gestures. A 2017 report by Common Sense Media stated that children ages 5 to 8 spend an average of 24 minutes a day playing games out of an average of one hour and 2 minutes of mobile usage (Common Sense Media, 2017). The authors reported that playing games was the second most common usage of mobile devices by children, behind watching videos at 25 minutes. Thus, playing games accounted for 38.7% of children's mobile usage. Even in the 2- to 4-year-old age group, the survey reports that children spend an average of 16 minutes per day playing mobile games out of a total of 58 minutes of mobile usage (27.6% of usage). Clearly, children spend a non-negligible amount of time playing games on mobile devices. Thus, developers of some games could improve children's experiences with their apps by better understanding children's gestures. Improving gesture recognition could help improve children's experience as well as allowing for more complex interactions in games.

Several mobile games for children make use of touchscreen gestures. Anthony et al. (Anthony et al., 2012a) conducted a survey of 23 games for children on the Android marketplace and found that 6 of them used touchscreen gestures as a form of interaction. This suggests the number of games that use gesture-based applications for children is low, but this could be due to low recognition rates, and they provide useful examples of how touchscreen gestures are used. We walk through some examples of games that could be played by children to provide some concrete examples. Consider the application Magic Touch: Wizard for Hire (Nitrome, 2015), in which the user must create gestures to cast spells. Figure 3 shows a screenshot of the game in which the user is drawing a "V" gesture to cast a spell to destroy the corresponding balloon. Other games use a similar interaction mechanism of creating gestures to cast spells or perform other actions (Toast Games, 2016). Other popular games like Candy Crush Saga (King.com Ltd., 2017) and Temple Run (Imangi, 2011) make use of a simple swipe touch gesture. Beyond the obvious benefit of reducing children's frustration when using these games, improving recognition of children's gestures in games can help enable beneficial new experiences, especially in the context of education. As an example, Williford et al.'s (Williford et al., 2017) ZenSketch application uses recognition as part of a game to help engineering students develop their freehand drawing skills. Games offering a wide variety of benefits, particularly development of motor skills, could be developed with improved recognition.

Figure 3. A screenshot of Magic Touch: Wizard for Hire (Nitrome, 2015), a mobile application for children in which the user draws a gesture to cast a spell that pops a balloon.

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

C. Educational Applications

A number of educational applications make use of touchscreen gestures. The lower number of educational applications may be due to the expectation that the application will provide measurable benefits to the children using it, which presents difficulty since recognition rates for children's gestures are so low (Woodward et al., 2016). Thus, in order to expand the potential benefits of educational applications using gesture interaction, it is important that research be devoted to improving recognition of children's gestures.

Gesture-based educational systems may also have the benefit of improving education in some contexts through their natural support of tactile learning (the acquisition of knowledge through physical activity (Gardner, 1985)). In other words, the physical act of performing a touchscreen gesture may offer benefits that other modalities do not. In the case of children, physically drawing the letters, numbers, shapes, and symbols can offer children the opportunity to develop motor skills that may also apply to pencil and paper interaction. One study supporting tactile learning in adults using touchscreens was that of Appert and Zhai (Appert and Zhai, 2009). In the study, the researchers showed that stroke-based gesture shortcuts could be equally effective as keyboard-based shortcuts, with the added benefits of aiding in recognition and recall due to the gestures' natural resemblance to their intended action. However, continued work is needed to better understand this potential benefit of gesture interaction, particularly as it may relate to children.

A common class of applications helps children develop their skills with creating letters, numbers, and shapes by tracing a guide and providing feedback. Recognition allows the system to provide more detailed feedback than a simple check against the provided tracing template. For example, the system might say, "it looks like you drew the number '8' instead of the letter 'A'" after a failed attempt to recognize an "A." This model of feedback and resubmit has been shown to improve learning outcomes compared to traditional methods (Malmi and Korhonen, 2004). Other educational drawing applications instruct the child to draw on a blank canvas (without tracing), then use a recognition algorithm to evaluate the input. Lanna and Oro (Crescenzi Lanna and Grané Oro, 2019) reported that children employ gestures when using drawing and coloring apps on touchscreen devices.

Schuler's (Shuler, 2009) report on educational technology points out several key benefits of educational applications for children. According to the report, these applications enable children to learn anywhere at any time without the need for an instructor to be present. These applications also enable personalized learning experiences, which can be provided by adaptive algorithms that are not feasible in traditional classroom interactions. However, to take advantage of these potential benefits, it is important that the applications be designed with children's interactions in mind. A system's ability to offer an educational experience depends on its ability to evaluate the users' input, so improved recognition is necessary for gesture-based systems to reach their full potential.

These examples of use of recognition in educational applications helps illustrate how touchscreen devices can serve as an aid to children's learning rather than a hindrance. Improving recognition would enable these systems to better understand the children's intent and to provide more appropriate feedback. Thus, the system can go a step beyond simply recognizing the child's gestures and allow them to correct poor gesture articulations in the same way that a teacher may help them to correct malformed handwriting. Improved recognition can offer richer interaction experiences for children using educational applications.

III. Children's Touchscreen Interactions and Gesture Recognition

In this section, we describe the prior work that has been conducted to analyze children's touchscreen gestures. However, because the body of work in the area specifically focusing on children has been very limited, we also include work on gesture recognition and interaction that did not specifically target children and discuss how it might apply or extend to children. Furthermore, the recognition techniques that have been applied to children's gestures thus far have been directly derived from prior work on adults' gestures. However, we describe the work through the lens of how it relates to children's gestures.

It has been well documented that children's touchscreen input behaviors are not equivalent to those of adults (Anthony et al., 2012a; Arif and Sylla, 2013; Hiniker et al., 2015; Vatavu et al., 2015b). In fact, several prior studies have shown that even a single touch or swipe can be used to identify whether a user is a child or adult with over 85% accuracy (Cheng et al., 2020; Nguyen et al., 2019; Vatavu et al., 2015a). Touchscreen interactions are also quite different from traditional mouse input, and the effects vary among different ages of children (Findlater et al., 2013), indicating the importance of studying touchscreen interactions in children of specific age groups, as compared to general mouse input interactions. Most commercial hardware devices with touchscreens like iPads or Android tablets are generally designed for adults, but specifically investigating children's interaction patterns allows application designers to improve their applications for children. One factor that makes recognizing children's gestures a difficult problem is that the interaction patterns of children of different ages are also quite different from one another. Younger children (e.g., ages 5 to seven years old), for example, tend to be less consistent in creating gestures than older children (e.g., 8 years old and older) (Brown and Anthony, 2012; Woodward et al., 2016). Figure 4 illustrates the wide variety among gestures from children of different ages. Thus, the age group of intended users of touchscreen applications is an important factor for designers to consider. Prior studies have offered a number of guidelines for designing touchscreen applications for various ages of children (Anthony, 2019; Anthony et al., 2012a; Hiniker et al., 2015; McKnight and Fitton, 2010; Nacher et al., 2015; Soni et al., 2019a; Woodward et al., 2016). For example, Anthony et al. (Anthony et al., 2012a) suggest training age-specific recognizers for recognizing children's gestures, and Woodward et al. (Woodward et al., 2016) suggest using more training examples

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Figure 4. "Diamond" gestures produced by children ages 5 to 10 from Woodward et al.'s study (Woodward et al., 2016), used with permission. Each column represents a different user, and gestures are scaled uniformly to show variation in size.

when training recognizers for younger children. We focus our discussion of existing work on touchscreen interactions and gesture recognition into several major categories: (A) developmentally appropriate prompts and feedback, (B) selection of gestures, (C) gesture elicitation, (D) recognition and classification of children's gestures, and we conclude by discussing the challenges associated with studying children's touchscreen gestures (E).

A. Developmentally Appropriate Prompts and Feedback

Understanding the differences in interaction patterns among children of different ages allows designers to create adaptive experiences for the children using their applications. Several studies have investigated ways in which designers can adapt their applications to make them more suitable for specific age groups of children. For example, Hiniker et al. (Hiniker et al., 2015) showed that designers should consider the age of children who are the target audience when prompting them to make gestures, since 2-year-olds respond better to visual cues than audio cues, but the opposite is true for 5-year-olds. McKnight and Fitton's (McKnight and Fitton, 2010) study of 6- to 7-year-olds found that they make different errors when prompted to provide different types of input, such as , such as press and drag, select, and double click. The amount of time taken to respond

Dataset	Gesture Types
Anthony et al. (Anthony et al., 2012a)	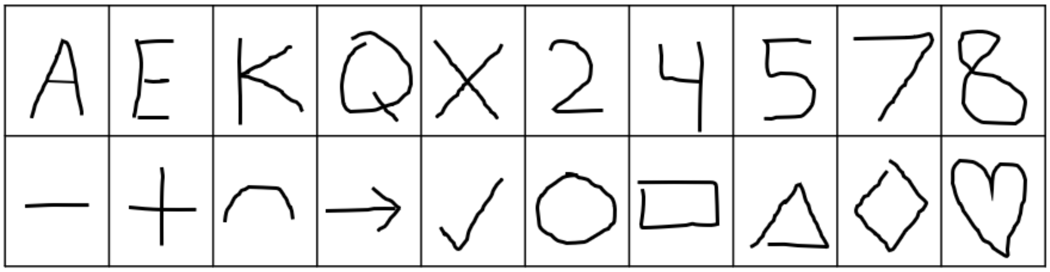
MMG (Anthony and Wobbrock, 2012)	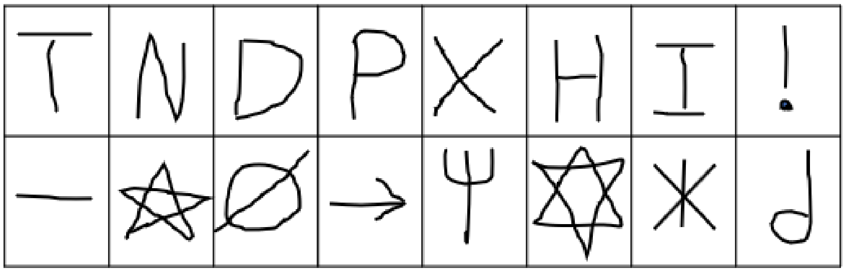
Algebra (Anthony et al., 2012b)	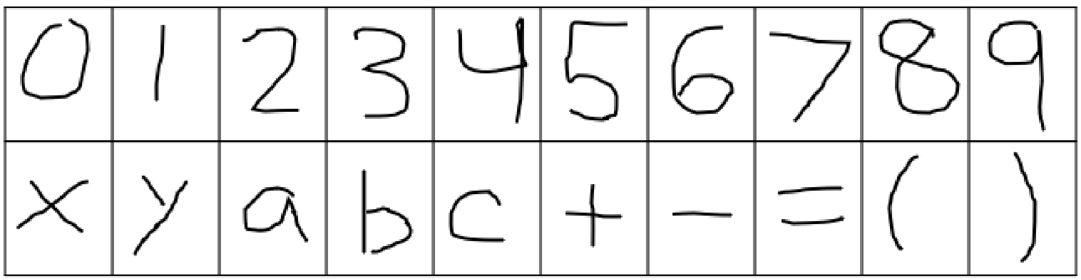
Unistroke (Wobbrock et al., 2007)	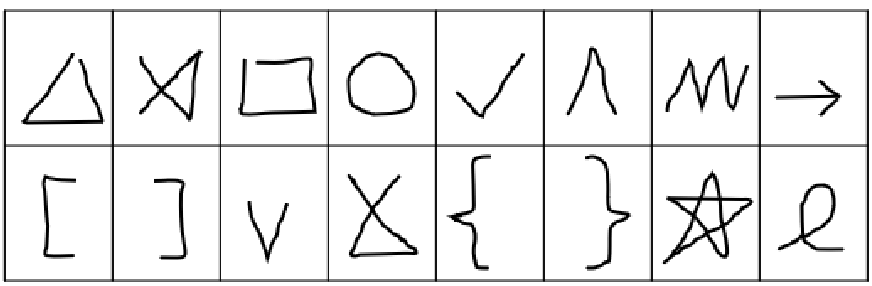
HHReco (Hse and Richard Newton, 2005)	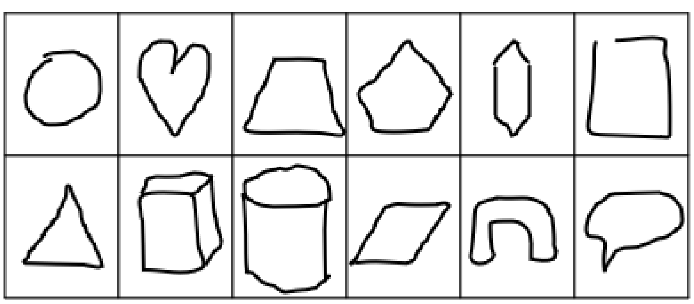
NicIcon (Willems et al., 2009)	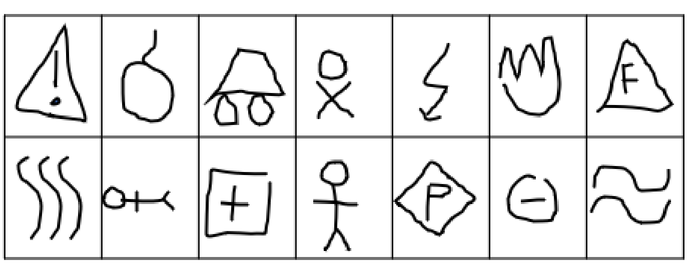

Table 2. Gesture sets used in several previous recognition experiments.

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

to those prompts also varied significantly. A study by Andrén (Andrén, 2011) found that children's use of longer, more sustained strokes in gestures led to improved responses communication with their parents.

Prior work has also found that children benefit from additional feedback when interacting with touchscreens as compared to adults. Anthony et al. (Anthony et al., 2015, 2013a) tested the effect of visual feedback on recognition accuracy of two different template matchers, $N-Protractor (Anthony and Wobbrock, 2012) and $P (Vatavu et al., 2012) (these recognizers are discussed in more detail later in this paper). The study found that children created gestures that were different enough to affect the accuracy of $P (Vatavu et al., 2012) in the presence or absence of visual feedback, but that $N-Protractor (Anthony and Wobbrock, 2012) did not display this sensitivity. Furthermore, a number of different features of the gestures, including the width and height of the gestures, were significantly different in the presence or absence of visual feedback in the younger participants than the older participants. Both younger children and adults reported that they preferred gesture interactions with visual feedback to those without feedback.

B. Selection of Gestures

Another issue faced by application designers when creating gesture-based applications for children is that of selecting an appropriate gesture set that will be easily used by children without being too highly constrained by its simplicity. A study by Nacher et al. (Nacher et al., 2015) on multi-touch gestures in 2- to 3-year-old children suggests that developers underestimate children's ability to perform complex gestures like rotation and scale-up. However, in contrast to Nacher's study, children ages 2 to 4 who participated in a study by Aziz et al. (Abdul Aziz et al., 2013) did have trouble with free rotate, drag and drop, and pinch and spread gestures. Research with adults has also shown that users are better able to remember gesture sets that they themselves define rather than having them predefined (Jego et al., 2013; Nacenta et al., 2013). It is not clear whether this finding would extend to children, given the rapid changes in a child's memory during development, e.g., between the ages of 4 and 8 (Gathercole, 1999), and children's tendency to try novel new gestures when interacting with new devices (Rust et al., 2014; Soni et al., 2019b). Further work is needed to better understand the types of gestures children are best able to remember.

A number of different gesture sets have been used to test the accuracy of gesture recognition algorithms. Most have not been evaluated in recognizing children's gestures, but familiarity with the gesture sets is useful for understanding the overall state of gesture interaction. Table 2 shows some of these gesture sets. The Unistroke set (Wobbrock et al., 2007) was designed specifically for testing general stroke gesture recognizers that were limited to single stroke gestures, and MMG (Anthony and Wobbrock, 2012) was later developed to test general multi-stroke recognizers. HHReco (Hse and Richard Newton, 2005) and NicIcon (Willems et al., 2009) provide domain-specific gesture sets reflecting geometrical and safety symbols, respectively. Anthony et al.'s (Anthony et al., 2012a) gesture set, the only set in this discussion designed specifically for kids, was created based on a survey of psychological and developmental literature as well as existing applications for children, and has been used in several studies on children's gestures (Brown and Anthony, 2012; Shaw and Anthony, 2016a; Vatavu et al., 2015b; Woodward et al., 2016).

C. Gesture Elicitation

Elicitation studies are a method of understanding users' natural gesture execution tendencies. In elicitation studies, participants are asked to perform an action using whatever gesture they feel would be most appropriate to perform that action, and their response is recorded. This process is repeated for a number of participants, and the experimenter then analyzes the responses, looking to see if there is any similarity among the gestures between the different users (Wobbrock et al., 2009). Agreement rates are computed among the responses to quantify the level of agreement among the participants, and a gesture set may be developed based on the most common gestures elicited (Wobbrock et al., 2005). Elicitation studies can give insight into the mental models that users have when interacting with touchscreens, which some studies accomplish using a "think-aloud" protocol, in which participants describe and explain each gesture they make during the study (Rust et al., 2014). A full treatment of gesture elicitation studies is beyond the scope of this work, but we outline some important studies in the domain to illustrate the impact of elicitation studies in Human-Computer Interaction (HCI).

Gesture elicitation were presaged by Wobbrock et al.'s (Wobbrock et al., 2005) study on the guessability of gestures. In Wobbrock et al.'s study, guessability refers to the degree to which a user is able to know what a gesture refers to even without being familiar with the gesture. The authors demonstrated a formalized method of improving guessability of gesture sets for adults by having participants define gestures. In what is generally considered the first elicitation study, Wobbrock et al. (Wobbrock et al., 2009) elicited natural gestures on a multi-touch tabletop device from users by showing them the effect that the gesture has, then asking them to create the gesture that would have that action. The authors report that previous experience with desktop devices strongly influences the way people made gestures on the multi-touch device, and they present a complete gesture set based on their findings.

In another study, Morris et al. (Morris et al., 2010) asked participants to rate different multi-touch tabletop gestures created by researchers as well as multi-touch gestures created by participants from a previous study (Wobbrock et al., 2009). Participants rated the gestures based on their perceived level of match between gesture and referent, and ease in performing the gesture. To prevent bias by the participants, they did not know which gestures were designed by researchers and which were designed by participants in the previous study. These gestures were gathered in an elicitation study (Wobbrock et al., 2009), in contrast to those designed by the researchers, which were reportedly designed based on the researchers' intuition of what

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

would be the most natural interaction for users. In general, gestures that had been created by other participants were blind-rated more highly than those created by the researchers for ease of use. A possible explanation for why the designers' gestures were less preferred to the elicited gestures could be that the designers' gesture set was affected by so-called "expert blind spot" (Nathan et al., 2003). The gesture set designers knew more about how the gesture recognition process worked, which could have led them to create gestures that were more biased in favor of better recognition than natural interaction. Even if the elicitation study was designed without a specific recognizer in mind, the researchers may have been influenced by their knowledge of the recognition process to create gesture sets that could be more easily distinguished between and recognized. Thus, designers should consider the tradeoff of using a gesture set they design versus a gesture set elicited from users, which users may find more natural, but which may pose challenges for recognition (e.g., conflicts, noise).

Elicitation studies help us understand users' natural interaction tendencies, but they suffer from several key limitations. First and foremost, elicitation studies are prone to legacy bias, described by Wobbrock et al. (Wobbrock et al., 2009) (though the term was first used by Morris et al. (Morris et al., 2014)). The concept of legacy bias refers to users' natural predisposition to perform new interactions based on what they have learned from their previous interactions with technology. This effect can result in gesture sets that are informed by these previous experiences rather than more novel (and possibly more natural) interactions. To counter legacy bias, Morris et al. (Morris et al., 2014) introduced three techniques to improve elicitation studies, including *production* (requiring multiple possible gestures for each prompt), *priming* (exposing participants to the technology or examples of gestures produced by experts before the study), and *partners* (having participants take part in elicitation studies in groups).

Another limitation of elicitation studies in gesture interaction is cultural differences among users. Participants in elicitation studies are generally from similar areas of the world, and as such are likely part of the same culture. Thus, it is not clear whether their interactions will generalize to people of other cultures. To tackle this problem, Mauney et al. (Mauney et al., 2010) conducted an elicitation study on a simulated touchscreen device with participants from nine different countries, finding a generally high level of agreement among the users. The study does, however, note that Chinese participants created significantly more symbolic gestures than participants from other countries. Further work is needed with more diverse pools of participants to characterize cultural differences among participants' gesture interactions in elicitation so that designers can better be prepared to select gesture sets for applications across different cultures.

1) Elicitation of Gestures with Children

Several elicitation studies have focused specifically on gestures produced by children. Connell et al. (Connell et al., 2013) conducted an elicitation study examining whole body interaction for menu navigation in children 3 to 8 years old using the Microsoft Kinect. The study reported a low level of consistency among the gestures elicited from the children. Rust et al. (Rust et al., 2014) conducted an elicitation study with both adults and children on a multi-touch tabletop computer, then compared the gestures and offered guidelines for future elicitation studies. Adults and children created similar gestures that were influenced by their previous experience with making touchscreen gestures, but children were more likely to invent new gestures. Soni et al. (Soni et al., 2019b) found similar results in an elicitation study using a spherical touchscreen device. Children were again more likely to invent new gestures than adults. However, there have not been any gesture elicitation studies of children using small screen touchscreen devices like phones or tablets.

D. Recognition and Classification of Children's Gestures

Early work in gesture recognition (Connell and Jain, 2001; Rubine, 1991; Wobbrock et al., 2007) focused on recognizing gestures produced by adults, largely overlooking children. In those studies that do examine recognition of children's gestures, the same recognition algorithms are used for the children as for the adults, even though they generally perform much worse for children's than adults' gestures (Anthony et al., 2015; Woodward et al., 2016). Because studies of adults' gestures have been so important in the development of the algorithms applied to children's gestures, this section discusses studies conducted with both children's and adults' gesture data.

Several studies have examined recognition rates of children's touchscreen gestures. Anthony et al. (Anthony et al., 2015) examined recognition rates for children ages 10 to 17 using two popular template-based recognizers, $N and $P[1] (Anthony and Wobbrock, 2012; Vatavu et al., 2012). Woodward et al.'s (Woodward et al., 2016) study of interface complexity, which also employed $P, examined recognition rates of 5- to 10-year-olds' gestures in both user-dependent and user-independent contexts (see Section 2). User-dependent experiments estimate how well a recognizer can perform given training examples from the person whose gestures it will recognize. User-independent experiments provide a measure of the recognition accuracy that can be expected when the recognizer is used as part as an unconfigured, newly shipped application, without being trained on the specific user of the application. Anthony et al. (Anthony et al., 2015) and Woodward et al.'s (Woodward et al., 2016) studies of 5- to 10-year-olds conclude that the 5-year-olds' gestures are recognized with the least accuracy (approximately 65%), and that accuracy steadily increases for older children, with 10-year-olds having the highest accuracy (approximately 94%). Furthermore, Woodward et al. (Woodward et al., 2016) reported a significant interaction between the user's age and the number of samples used to train recognizers with respect to the accuracy achieved. Based on this finding, the authors recommended using more training samples for younger

[1] Discussed in detail in Section V.

children in order to achieve higher accuracy. However, this recommendation can present a challenge, since it can be time consuming to collect gestures from young children due to their tendency to lose attention during laboratory studies (Brewer et al., 2013). Other methods of collecting gestures, such as longitudinal data collection outside of the lab (Johnson and Turner, 2002), may not present the same difficulty with getting sufficient training data from children.

To establish a baseline for comparing future work, Shaw et al. (Shaw et al., 2017) compared the ability of human viewers to identify gestures produced by children ages 5 to 10, finding a significant difference between human recognition (90.60%) and machine recognition (84.14%). The authors suggest the 90.60% accuracy rate obtained by humans can be used as a benchmark for future work on recognizing children's gestures.

A study worth particular note due to its relevance to this paper is that of Kim et al. (Kim et al., 2013) on KimCHI, a system designed to classify children's developmental skill and gender based on their execution of the digits 0 to 9 and the letters A to F. The system distinguished between preschool versus grade schoolers with 82.7% accuracy and classified the gender of the children in the study based on their sketches with 72.8% accuracy. The authors collected a total of 725 gestures from four adults, twelve 7- to 8-year-olds, and eight 3- to 4-year-olds. The system does not, however, attempt to recognize the gestures. The distance between the age groups included in the study makes it difficult to develop a cohesive understanding to characterize children's touchscreen interactions across all ages. The gap in age groups led Kim et al. to conduct another study which built on their KimCHI framework, introducing EasySketch[2] (Kim et al., 2016), an intelligent interface that combines KimCHI's developmental classifier with a gesture recognizer to help development of children's motor control. The system, which was tested on 70 children from ages 3 to 8, helped children improve their ability to draw by providing a developmentally appropriate interface with adaptive prompts and feedback; it used an algorithm developed by Valentine et al. (Valentine et al., 2012) to perform recognition. The system provides feedback and a 'trace-the-dots' activity to help the child develop their skills. The system quantified children's improvement in drawing ability based on the similarity of each gesture to a predefined template.

E. Challenges in Studying Children's Gestures

Recognition of children's gestures presents several challenges that are less problematic when dealing with adults' gestures. One such challenge is that recognition experiments generally require a large number of samples for training, which can require participants to take part in long studies. Children are prone to lose interest and stop participating in empirical studies if they find them uninteresting (Brewer et al., 2013; Punch, 2002). To deal with this issue, Brewer et al. (Brewer et al., 2013) introduced a method of gamifying gesture collection in which participants are awarded points for completing individual components of the study. After completion of the study, the children have the opportunity to claim a prize, such as a small toy or stickers, based on the number of points they earn. The paper reported an increase from 73% completion without gamification up to 97% completion with gamification, in a study with children ages 5 to 7 years old, indicating a significant benefit in data fidelity can be gained by making an empirical study more engaging for children.

Beyond the challenge of collecting large numbers of samples of gestures from children, there are also challenges in dealing with the data produced by children. For example, several prior studies have reported that younger children sometimes draw the wrong gesture or scribble randomly when prompted to draw some of the gestures (Woodward et al., 2017, 2016). Additionally, Bahamóndez showed that children's writing on a touchscreen device was slower and less legible than traditional on-paper writing. Prior work on children's gestures have considered that giving the participants the ability to erase their gestures is not ideal, since it may encourage them to try to produce 'beautified' gestures rather than more natural ones, leading to less accurate representations of the gestures that would need to be recognized in a real application (Anthony et al., 2013b; Shaw and Anthony, 2016a; Woodward et al., 2016). A method of preventing children from randomly drawing that has been employed in previous studies is to have them produce an example of each of the gestures in the corpus on a sheet of paper, which they can then use as a reference if they feel unsure (Anthony et al., 2015; Brewer et al., 2013; Woodward et al., 2016), preventing the children from getting confused during gesture collection.

Children of different ages can produce very different gestures, due in part to both cognitive and motor development progress. Figure 4 shows an example of how drastic these variations in gestures can be, in data collected during a study by Woodward et al. (Woodward et al., 2016). The diamond gestures are highly variable for the younger children but become more consistent for older children and adults.

The developmental differences among age groups is also reflected by research on fine motor control, which shows that children rapidly develop gross motor skills during their first two years of life (Newell, 1991), and continue to develop fine motor skills for the next several years. From age two to seven, children reach maturity in several motor tasks, like walking and running, and they begin to exhibit more refined motor control (Seifert and Hoffnung, 1987) in their use of their hands, fingers, and feet. Development of these fine motor skills is not only affected by age, but a number of other factors, such as a child's experience and genetic disposition (Kakebeeke et al., 2013). Even among neurotypical children, there is a great deal of individual variation, indicating the importance of studying gesture patterns across ages. Personalized gesture recognition may be necessary to support the range of children's gesture production patterns.

The high level of variability indicates the value of studying children's gestures at the fine-grained level of individual age groups, but this is not possible in all cases due to sample size limitations. Thus, researchers may analyze gesture interactions using groupings in which children of similar developmental levels are analyzed together. Several past studies (Anthony et al., 2015, 2012a; Connell et al., 2013; Kim et al., 2016; Shaw

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Type of Recognizer	Strengths	Weaknesses
Feature Based Statistical Classifier	Easy to implement; high recognition rates for gestures that fit the model defined by the features	Not adaptive; limited by simplicity; relatively low accuracy; may require data in specific format
Template Matcher	Easy to implement and understand even for novice developers; relatively little code needed; high recognition rates with relatively few training examples; fast runtime	Does not scale well for very large gesture sets; must iterate through each template, which can be both time and space consuming
Hidden Markov Model (HMM)	High overhead from segmentation and training; high accuracy	High complexity; difficult to implement for novice developers
Neural Networks	Able to handle very large gesture sets; can adaptively learn and improve from mistakes	Difficult to implement and understand; implementation details can be obscured by complexity
Mixed Methods	Combines positive aspects of other types of recognizers to improve recognition rates	Overhead of combining disparate algorithms can lead to slow runtime and difficulty in programming

Table 3. Several types of recognition algorithms and their strengths and weaknesses.

and Anthony, 2016a) on children's touchscreen interactions have grouped children based on Piaget's (Piaget, 1983) theory of cognitive development, which posits that children undergo four stages of cognitive development. Stage 1 of Piaget's model is the sensorimotor stage, which begins at birth and lasts until age two. Stage 2 is the preoperational stage, which spans from age two until age seven, followed by Stage 3, the concrete operational stage, which begins at age seven and ends at age eleven. Stage 4, the final stage of the model, is the formal operational stage, spanning from age eleven to adulthood. These classifications allow for useful and interesting comparisons, but prior work by Woodward et al. (Woodward et al., 2016) found substantial differences within these groupings, reporting 64% accuracy for 5-year-olds, 79% for 6-year-olds, 90% for 7-year-olds, 85% for 8-year-olds, 88% for 9-year-olds, and 95% for 10-year-olds. The wide variation in accuracy indicates the importance of studying the interactions of individual ages of children based on cognitive and motor development.

IV. Gesture Recognition Techniques

As previously discussed, touchscreen gesture recognition accuracy rates are much lower for children than they are for adults, and even worse for younger children than older children (Anthony et al., 2015, 2013b, 2012a; Woodward et al., 2016). However, only a few recognition algorithms have been tested on children, and they are all template matching algorithms. There has been very little work on developing algorithms specifically targeted to recognizing children's gestures. Therefore, we describe the various recognizers that have been designed for adults with additional discussion on how these various types of recognizers may be useful in recognizing children's gestures.

Prior work in gesture recognition has seen the development of a large number of recognition algorithms. Most of these algorithms can be grouped into the following major categories: (A) template matching approaches, (B) feature based statistical classifiers, (C) hidden Markov models (HMMs), (D) neural networks, and (E) combinations of these methods. A short description of each of these types of recognition algorithms is provided here, with some examples of some recognizers in each category, followed by a discussion of their advantages and disadvantages, and promise for recognizing children's gestures. Table 3 summarizes the strengths and weaknesses of each of the categories of recognizer.

Table 4 shows a summary of all the recognizers discussed in this paper, including recognition results reported in the papers in which they were introduced. The column labeled "# Gestures" refers to the number of different types of gestures in the gesture set used for evaluation. A cell value of "NR" means not reported, indicating that the value was not reported in the original paper.

A. Template Matching Approaches

Template matching recognizers compare candidate gestures to preselected examples of the gestures and returning the closest match as the result. The members of the $-family ("dollar family") of recognizers, which includes the $1, $N, and $P recognizers (Anthony and Wobbrock, 2012, 2010; Vatavu et al., 2012; Wobbrock et al., 2007) are notable examples of this type of recognizer, and have been tested on children's gestures. These algorithms are popular due to being relatively easy to implement and highly accurate despite their simplicity (Taranta II and LaViola Jr., 2015).

The $1 recognizer (Wobbrock et al., 2007), the first of the three aforementioned $-family recognizers to be developed, recognizes unistroke gestures by scaling and resampling the points of each gesture uniformly, then finding the template which minimizes the distance between each corresponding pair of points in the candidate gesture and the template gesture. The $N recognizer (Anthony and Wobbrock, 2012, 2010), the second in the $-family, built on the limitations of the $1 recognizer, extending it to multistroke gestures. The $N

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Recognizer	Description	Accuracy	# Gestures	# Training Samples Used	# Users	Ind/Dep
Template Matchers						
GRANDMA (Rubine) (Rubine, 1991)	Early feature-based classifier; computes vector of 13 geometric features and compares values for candidate gestures to known gestures	92%	20	10/gesture	NR	Dep
GesEdit (Cho, 2006)	Feature-based classifier of Korean characters that outputs gesture type based on values of 9 different features of varying complexity	99.6%	11	N/A	20	Dep
GDE (Apte et al., 1993)	Multistroke geometric shape recognizer based on feature-based 'filters'	97.5%	6	N/A	10	Dep
Smithies (Smithies et al., 1999)	Recognizer employing approximately 50 features to recognize individual characters in a mathematical equation editor	N/A	N/A	N/A	9	Dep
Olsen et al. (Olsen et al., 2007)	Uses direction of substrokes as a feature to classify unistroke gestures	100%	9	204 total (23/gesture)	3	Dep
Blagojevich et al. (Blagojevic et al., 2010)	Introduced 114 features for Rubine and used machine learning to find best combination	96.9%	6	N/A	20	Dep
Feature Based Statistical Classifiers						
Lee et al. (Lee et al., 2007)	Graph-based symbol recognition algorithm comparing four matching algorithms	97%	23	14/gesture	9	Dep
Connell and Jain (Connell and Jain, 2001)	Template-based decision tree algorithm to classify handwritten characters	86.9%	36	17,928 total (498 /gesture)	NR	Dep
!FTL (Vanderdonckt et al., 2018)	Fast template matcher that matches triangles formed by vectors created based on the gesture's path	95.1%	14	30 /gesture	33	
Gestimator (Ye and Nurmi, 2015)	Gestures broken down into strokes which are then used in template matching	99.9%	18	5 /gesture	43	Dep
$1 (Wobbrock et al., 2007)	Unistroke point matching approach	99.5%	16	9 /gesture	10	Dep
$N (Anthony and Wobbrock, 2010)	Multistroke gesture matching	96.6%	16	9 /gesture	10	Dep
Protractor (Li, 2010)	Unistroke closed-form matching via minimum angular difference between gestures	99.6%	16	9 /gesture	10	Dep
$N-Protractor (Anthony and Wobbrock, 2012)	Closed form version of $N using Protractor's matching method to improve runtime	94.5%	16	15 /gesture	13	Dep
$P (Vatavu et al., 2012)	Point cloud based matching approach	99.4%	16	9 /gesture	20	Dep
$Q (Vatavu et al., 2018)	Fast point cloud matching approach designed for wearables	99.7%	16	10/gesture	20	Dep
1¢ (Herold and Stahovich, 2012)	Rotationally invariant approach building on $1	97%	10	14 /gesture	14	Dep
Penny Pincher (Taranta II and LaViola Jr., 2015)	Extension of $N matching vectors between adjacent points; improves runtime of $N	99.9%	18	10 /gesture	20	Ind
$3 (Kratz and Rohs, 2010)	Extension of $1 to 3-dimensional gestures	79.9%	10	5 /gesture	12	Dep
Hidden Markov Models						
Sezgin et al. (Sezgin and Davis, 2005)	Hidden Markov Model based approach to segmentation and recognition of symbols	96.5%	10	10 /gesture	10	Ind
Jiang and Sun (Jiang and Sun, 2005)	Hidden Markov Model based approach to sketch recognition	95%	23	14 /gesture	9	Ind
Li and Yeung (Xiaolin Li and Dit-Yan Yeung, 1997)	Hidden Markov Model based approach to recognizing segmented handwritten characters	91%	62	N/A	21	Ind
Anderson et al. (Anderson et al., 2004)	Hidden Markov Model based approach to recognizing touchscreen gestures	94.2%	12	60 /gesture	3	Ind

Table 4. A comparison of the recognizers presented in this survey.

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Neural Networks						
Lecun et al. (LeCun et al., 1990)	Neural network for recognizing handwriting digits	90%	10	984/gesture	NR	Ind
Singh and Amin (Sameer Singh, n.d.)	Neural network matching approach for recognizing symbols and sketches	86%	52	10 /gesture	NR	Ind
Lee (Lee, 1996)	Neural network approach for recognizing handwritten digits with high accuracy even for edge cases	97.1%	10	4,000 total (400 /gesture)	NR	Ind
Shrivastava and Sharma (Shrivastava and Sharma, 2012)	Neural network matching approach for recognizing characters	75%	20	11 /gesture	NR	Ind
Mixed Methods						
SHARK[2] (Kristensson and Zhai, 2004)	Graph-based template matching recognizer incorporating elements of feature-based classifiers	N/A	N/A	N/A	NR	Dep
SATIN (Hong and Landay, 2000)	Fusion of ten different classifiers which first beautify and correct perceived errors in gestures then recognize them	N/A	N/A	N/A	NR	Dep
Yin and Sun (Yin and Sun, 2005)	Multi-stroke template matcher based on minimal fitting error, supported by optimization via dynamic programming	98%	4	400 /gesture	4	Dep
Alimoglu and Alpaydin (Alimoglu and Alpaydin, 1997)	Recognizer capitalizing on various combinations of recognition approaches	99.3%	10	3,748 (approx. 37 /gesture)	44	Dep

Table 4 (continued). A comparison of the recognizers presented in this survey.

recognizer works by matching stroke sequences and ordering between the candidate and template gesture. Thus, the algorithm must consider every possible ordering and direction of strokes, leading to relatively long runtimes for gestures with many strokes. The $P recognizer (Vatavu et al., 2012) treats each multistroke gesture as a ‘cloud’ of points without regard to individual strokes. After scaling and rotating, the $P recognizer finds the best possible match between a candidate and template gesture, returning the template with the highest score as the result. $P is reported to deliver up to 99% accuracy on a corpus of 10 examples of each of 16 gesture types from 20 adults at a much lower computational cost than $N (Vatavu et al., 2012). However, $P does not consider the ordering of the strokes or the direction, which may be needed to recognize some types of gestures with high accuracy. In such cases, $N may achieve higher accuracy.

Vatavu et al. (Vatavu et al., 2012) provide a summary of the advantages and disadvantages of the $-family recognizers in the form of a “$-family cheat sheet.” $1 can only recognize unistroke gestures, whereas $N and $P can recognize both unistroke and multistroke gestures. All three of the recognizers have high accuracy (>98%). The algorithmic complexity, and thus the time taken to perform the recognition, is lowest for $1, followed by $P and then $N with the highest. Thus, $1 is ideal for unistroke gestures, and $P is the best choice for multistroke gestures. However, $P is not rotationally invariant, so it cannot distinguish between A and ∀, for example. The lack of rotational invariance in $P implies that it can only be used in contexts with predefined gesture sets that are guaranteed not to have rotational collisions between the gestures. In these cases, $N must be preferred over $P. Rotational invariance is important in recognizing children’s gestures, particularly due to their tendency to engage in “mirror writing” (Cornell, 1985), a common phenomenon in which children draw the intended gesture backwards or upside down.

The popularity of the $-family has led to a number of adaptations and improvements in other work. Herold and Stahovich (Herold and Stahovich, 2012) built on the $1 recognizer to build their 1¢ recognizer, which improves runtime by providing a one-dimensional representation of the gestures. Taranta and LaViola (Taranta II and LaViola Jr., 2015) introduced a $-family inspired multistroke recognizer called Penny Pincher that achieves high accuracy even in constrained timeframes. Penny Pincher operates by breaking gestures down into a series of two-dimensional vectors between pairs of adjacent points. In their $3 recognizer, Kratz and Rohls (Kratz and Rohs, 2010) built on the $1 recognizer to create a 3-dimensional recognizer by representing 3-dimensional gestures as continuous strokes and using a similar matching algorithm based on Euclidean distance.

Another template matching algorithm was presented by Connell and Jain (Connell and Jain, 2001). The algorithm operates by first reducing the gesture to a string based on the coordinates of the points in the strokes of the gesture, then performing a string-matching algorithm using a decision tree by calculating the distance between each pair of strings. The authors report a recognition rate of 86.9% accuracy on a corpus of approximately 18,000 gestures. In total, there were 36 classes of gestures collected from at least 21 users[2].

Vanderdonckt et al. created !FTL (Vanderdonckt et al., 2018), a fast template matcher based on a novel method of comparing gestures called local shape distance. In this method, gestures are broken down into triangles created by each three-point window along a resampled gesture’s path. The similarity of these triangles to templates is then compared to find the best match. The authors report approximately 95% accuracy on a set of 5,540 gestures taken from 14 gesture classes 33 participants.

[2] The exact number of users is not reported.

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children’s Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Gestimator (Ye and Nurmi, 2015) is a template matching algorithm that focuses on recognizing more complex gestures than other template matching approaches by first segmenting the gesture into its constituent strokes, then comparing combinations of individual strokes using a more traditional template matching approach. The authors report high accuracy: 99% with 5 training examples of 6 different gesture types from each of 13 participants, but the computational overhead is increased in adding segmentation.

Lee et al. (Lee et al., 2007) presented a graph-based template matching approach to symbol recognition that examined four different matching techniques, all of which obtained over 97% matching accuracy on a corpus of 15 examples of each of 23 gesture types collected from each of 9 participants.

Template matching algorithms are usually quite simple in their implementation, making them ideal for novice programmers who wish to quickly add gesture recognition to their user interface prototypes (Vatavu et al., 2012; Wobbrock et al., 2007). Despite their simplicity, template matching algorithms can achieve very high accuracy on adults' gestures, up to 99% in some cases, with sufficient training examples. For example, $P can reach 99% overall with 5 training examples (Vatavu et al., 2012). As prior work has shown (Anthony et al., 2012a; Woodward et al., 2016), template matchers are generally not well suited for recognizing children's gestures. Furthermore, template matching approaches are intended to be quick, easy substitutes for more complex recognizers when prototyping. They are not intended to be state of the art recognizers, so we expect recognition rates for other types of recognizers will be higher for children. We suggest the poor performance of template matchers is primarily because they rely on consistency among gestures of the same type.

B. *Feature Based Statistical Classifiers*

Feature based statistical classifiers employ a vector of features (that is, a group of metrics or measurements that are calculated on the gesture) to quantify the gesture and to use in classification. These features can take on any kind of value, but are usually numeric since they are calculations based on geometric properties. These values are then compared to a predefined threshold, and the recognizer returns the best match as the result. These recognizers are relatively easy to implement, but the complexity depends on the features used.

One of the earliest and most well-known feature-based recognition algorithms was described by Rubine in 1991 (Rubine, 1991). The recognizer computes 13 geometric features, such as the sine and cosine of the initial angle and the total length of the gesture and stores them as a vector. This vector is used to compare the candidate gestures using a linear discriminator. The candidate gesture with the most similar features to the test gesture is chosen as the result. Rubine's algorithm has been used as the basis for a number of gesture-based interfaces, such as Garnet (Myers et al., 1990), Amulet (Myers et al., 1995), and gdt (Long et al., 1999). Cho (Cho, 2006) used a technique similar to Rubine in which 9 different features are used to classify hand-drawn Korean characters with over 99% accuracy. Apte et al's (Apte et al., 1993) GDE recognizer used various features as 'filters' to recognize multistroke geometric shapes, but suffered from the inability to recognize those same shapes when drawn in a single stroke. Smithies (Smithies et al., 1999) used a feature vector of approximately 50 dimensions to recognize handwritten characters in a math equation editor. Blagojevic et al. (Blagojevic et al., 2010) developed a set of over 100 geometric features for use in classifying gestures and sketches, then used machine learning techniques to determine which combination of features resulted in the highest accuracy rate. The authors then altered Rubine's algorithm to use the selected features rather than Rubine's original features. The majority of the most effective features selected by the machine learning algorithm were related to either the curvature of the gesture or its size. The preference for these features indicates they are highly related to recognition, which could prove useful in future work on designing new recognition algorithms. In the domain of children's gestures specifically, Shaw and Anthony (Shaw, 2017; Shaw and Anthony, 2016a, 2016b) found that features related to curvature were the most affected by age, with younger children having a very large variation in curvature.

In Olsen et al.'s (Olsen et al., 2007) feature-based classifier, the extracted feature vector is a representation of the angle at which each substroke of a unistroke gesture is drawn. The feature vector is divided into six components, with the first representing the number of substrokes whose angle is between 0° and 30° relative to the horizontal, the second the number of substrokes between 30° and 60°, and so on. Euclidean distance matching is then used to recognize results by comparing feature vectors of the gestures. The authors report 100% recognition accuracy with a set of 9 different gesture types over a total of 204 gestures collected from 3 users.

Feature-based classifiers benefit from their relative simplicity compared to other forms of recognizers, allowing them to achieve very fast runtimes when the number of features computed is low. However, feature-based statistical classifiers are limited in that they make the assumption that gestures can be described by a mathematical formula derived from the features chosen for recognition. In cases where the gestures fit the model very nicely, high recognition rates are likely, but when they do not fit the model the recognition rates will likely be much lower. In previous work analyzing gesture articulation features of children's gestures, high levels of variance have been found (Shaw and Anthony, 2016a). This high level of variance indicates that the models used by these feature-based statistical classifiers will probably not fit children's gestures well, so we suggest that feature-based models will likely not be able to recognize children's gestures with high accuracy.

C. *Hidden Markov Models (HMM)*

While both template matchers and feature-based statistical classifiers are relatively simple for developers to implement, we hypothesize that machine learning approaches can achieve much higher accuracy for children's gestures. However, these machine learning approaches can be prohibitive since they require much larger datasets than the previously discussed methods. Two of the most common machine learning

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

approaches are Hidden Markov Model recognizers and Neural Networks.

Hidden Markov Models (HMMs) can be used to recognize gestures by breaking the gesture into a series of points or strokes that are input to the HMM, then used to determine a recognition result based on prior training examples. HMMs describe probabilistic processes in which there are unseen (hidden) states. The HMM uses these states to compute the most likely sequence of inputs (Baum and Petrie, 1966). As with general Markov models, the transition from one state to the next depends on only the current state. The HMM uses a dynamic programming (Bellman, 1952) algorithm to determine the recognition result by matching the observed sequence of inputs to the best fit among training data. This method of sequential states is useful in gesture and handwriting recognition, wherein each stroke (or letter in the case of handwriting) can be treated as a separate state. HMMs are particularly well-suited for gesture recognition because they represent a statistical model of spatio-temporal data that can handle variations in the articulation of the gesture.

Li and Leung (Xiaolin Li and Dit-Yan Yeung, 1997) presented an HMM-based recognizer that classifies constituent strokes of a gesture based on their position on the canvas, which serve as the states for the model. They report a recognition accuracy of 91% over a gesture set including 62 classes letters and numbers produced by 21 people. Sezgin and Davis (Sezgin and Davis, 2005) introduced an HMM-based recognition algorithm for sketches in which, as with the previous recognizer, each stroke of the sketch is treated as a state. In Sezgin and Davis's (Sezgin and Davis, 2005) approach, however, a separate HMM represents each of the gesture types in the set. The algorithm then matches the observed sequence of states to the HMM that it most closely resembles. The authors report 96.5% after training HMMs on 10 different gesture types with a total of 6 examples each, from a total of 10 users. Jiang and Sun's (Jiang and Sun, 2005) stroke-based HMM approach resulted in 95% recognition after being trained on 14,611 gestures across 9 gesture types from 2 users. In contrast to other HMM-based recognizers, Anderson et al. (Anderson et al., 2004) used individual points of gestures as the states in their HMM-based recognizer. The authors report an overall accuracy of 94.18% over 10 samples of 11 different gesture categories from a total of 3 participants.

HMMs have also been used in handwriting at the level of recognizing words rather than characters. An example of an early system of this type is that of Chen (Chen, 1994), who used an HMM-based approach to recognize handwriting by first segmenting the word into individual letters then using each letter as a state in the model. The author reports 43.6% accuracy over 1,563 words from an unspecified number of users. The ability of HMM-based algorithms to recognize gestures from large corpuses has made them a popular choice for recognizing handwriting in many languages (Ahmed and Azeem, 2011; El-hajj et al., 2005; Hu et al., 2000; Hu and Brown, 1996; Mohamad and Likforman-sulem, 2009; Nakai et al., 2001; Roy et al., 2015). Further information on the use of HMMs for handwriting recognition can be found in Plötz and Fink's (Plötz and Fink, 2009) survey of the topic.

HMMs have the advantage that they can achieve high accuracy even for very large gesture corpuses and sets. However, they can have high overhead due to the segmentation into states required to create and train the model, and training can require a large corpus of data, making their use impractical in some cases. In contrast to feature-based statistical classifiers and template matchers, HMM-based recognizers generally require that a developer have some knowledge of machine learning, thus making them less accessible to novice developers. However, we hypothesize that future work using HMM-based recognition of children's gestures may be more accurate than that of template matchers or feature-based classifiers, especially since HMMs can account for some variations in the way the gesture is articulated.

D. Neural Networks

Another common machine learning based approach to recognizing gestures is through the use of neural networks. Neural networks consist of a number of nodes (or neurons) arranged in layers such that the output of one layer is received as input in the next layer, and so on until the final layer outputs a classification decision (Mitchell, 1997). While being trained, the network compares its output to the correct output and back-propagates findings to the previous levels, allowing them to adjust the weights of the inputs used in their calculations, thereby improving their accuracy. Each node accepts numerical input values and then produces numerical output values. This can be applied to gesture recognition by feeding features calculated on the gesture in as input to the first layer of nodes. For example, a method used in recognizing handwritten digits is to convert each digit to an image, then treat it as an array of pixels where the RGB value of each pixel is sent to a different input layer. The first layer of nodes then performs computations on these values to feed to the next layer, and so on until a recognition result is reached (Mitchell, 1997). Neural networks are well suited for achieving high accuracy rates in identifying gestures from very large sets with over 100 gesture types (Zhang et al., 2018). Neural networks can be used to recognize gestures by supplying the network with a large body of correctly identified gestures on which to train, as well as the features that should be used to determine the result. The neural network can then create a model based on the input and selected features that will classify further examples of the gestures.

An early neural network-based recognizer was that of LeCun et al. (LeCun et al., 1990), which showed that back-propagation could be used to recognize handwritten digits from a large number of users. The authors report 90% recognition accuracy on a set of over 10,000 digits from multiple users[3]. Singh and Amin (Singh and Amin, 2014) used a neural network algorithm to recognize hand-printed characters, achieving 86% accuracy in recognizing characters from a set of 52 different characters from 21 writers by extracting primitive features such as straight

[3] The exact number of digits is not reported.

Please cite the definitive version of this paper: Alex Shaw, Jaime Ruiz, Lisa Anthony, A Survey on Applying Automated Recognition of Touchscreen Stroke Gestures to Children's Input, *Interacting with Computers*, Volume 32, Issue 5-6, September-November 2020, p.524–547, https://doi.org/10.1093/iwc/iwab009

Feature	Description	Feature	Description
R_1	Cosine of initial gesture angle	L_1 (Aspect)	Absolute value of (45 – R_4)
R_2	Sine of initial gesture angle	L_2 (Curviness)	Sum of all angles less than 19°
R_3	Length of bounding box	L_3	R_9 / R_8
R_4	Angle of bounding box	L_4 (Density 1)	R_8 / R_5
R_5	Distance between first and last point	L_5 (Density 2)	R_8 / R_3
R_6	Cosine of angle between first and last point	L_6 (Openness)	R_5 / R_3
R_7	Sine of angle between first and last point	L_7	Area of bounding box
R_8	Total gesture length	L_8	log(L_7)
R_9	Total angle traversed	L_9	R_9 / R_{10}
R_{10}	Sum of absolute value of angles	L_{10}	log(R_8)
R_{11}	Sum of squares of angle at each point	L_{11}	log(R_{11})
R_{12}	Max speed of gesture squared		
R_{13}	Duration of gesture		

Table 5. The 13 features employed in Rubine's (Rubine, 1991) recognizer (left) and the 11 features employed in Long et al.'s (Long et al., 2000) analysis (right). These features have been widely built on in creating new feature sets for new recognizers.

lines, curves, and loops. Lee et al. (Lee, 1996) described a neural network recognizer that recognizes single numerals with up to 99% accuracy on a set of 22,000 gestures of 22 different types from 9 participants. Another neural network recognition algorithm, though not used for handwriting/gesture recognition, was presented by Shrivastava and Sharma (Shrivastava and Sharma, 2012), which classified 360 computer-generated characters from 20 different fonts with high accuracy (up to 97%). As with Singh and Amin's approach, Shrivastava and Sharma's algorithm used an image-based method that extracts features of the gestures such as vertical and horizontal symmetry.

Neural networks are beneficial in that, when they are applied to large gesture sets, they can obtain high accuracy. However, one of the main drawbacks to using neural networks, beyond the technical knowledge required to implement them, is the large amount of data required to train the recognizer. This presents a challenge in the domain of recognizing children's gestures due to the lack of publicly available data and the difficulty involved in collecting new data from children (Punch, 2002). However, given sufficient data, we hypothesize that neural networks will be able to achieve higher recognition than simpler methods due to their ability to adapt the model based on the variability seen in children's gestures.

E. Mixed Methods

Several recognition algorithms employ a combination of the above techniques, allowing them to capitalize on the strong points of each, but also suffering from the overhead incurred by combining them. A notable example of a combined method is Kristensson and Zhai's SHARK[2] recognizer (Kristensson and Zhai, 2004; Zhai et al., 2012) for shorthand writing in pen-based computers, which uses a recognition pipeline including *template pruning*, that is, removing unneeded candidate gestures from the set of templates, and geometric analyses based on the shape of gestures produced when interacting with a digital keyboard. While the authors do not perform a recognition study, they do report a study showing that users are able to achieve fast word-entry rates using their system. Hong and Landay's SATIN (Hong and Landay, 2000) also employs this mixed methods paradigm, using ten different interpreters in concert to produce a recognition result. The authors illustrate the use of their system in a sketch-based application for drawing circuits. Because the focus of the work is on the design of the system, the authors do not perform a recognition study. Yin and Sun (Yin and Sun, 2005) created a novel multi-stroke recognition algorithm using a template matching method based on minimum fitting supported by optimization via dynamic programming. The authors report 98% accuracy on a set of 100 examples of each of 4 gesture types produced by 4 people. Alimoglu and Alpaydin (Alimoglu and Alpaydin, 1997) described a number of methods for incorporating multiple classifiers into a single recognizer, reporting improved rates for dynamic (mixed) recognizers over static, consistent with findings in other similar work (Gader et al., 1996; Verma et al., 2001). The overall accuracy rate reported is 99.3% after training on 3,748 gestures from a total of 10 users.

The advantages and disadvantages of using mixed methods for gesture recognition accuracy are largely determined by the types of recognizers included. Mixed methods allow an algorithm to capitalize on the parts of an algorithm that perform best depending on the type of input. However, the need to combine various recognizers adds additional overhead. The format of the input and output for each recognizer may also be different, so the developer may have the additional burden of converting the data between formats. However, because they could be configured to combine the aspects of the various types of recognizers that are best able to handle children's gestures, we hypothesize that a mixed method recognizer may be likely to be able to achieve higher accuracy than any of the other categories.

V. Gesture Articulation Features and Analyses

As discussed earlier in this paper, it is still not fully clear what specific behaviors cause recognition to be poor for children's gestures. To help further this understanding, prior work has employed analysis techniques to children's gestures that had previously been used to gain a deeper understanding of adults' gestures. One method of exploring these gestures in more detail that has been employed by several studies has been the examination of articulation features of gestures produced by various types of users. Gesture features generally refer to some measure calculated on the gesture based on its geometric properties. Gesture *articulation* features are specifically designed to capture something about how the user has produced the gesture (Anthony et al., 2013b; Long et al., 1999). Gesture features can be directly employed by recognizers (e.g., feature based statistical recognizers), or indirectly affect recognition (e.g., in template matchers or other approaches), so it is important to study them to understand how to improve recognition rates. A thorough analysis might also help researchers understand what articulation patterns specifically cause recognition rates to be lower for children than adults. Features alone probably cannot fully explain recognition differences, but they do offer a starting point for further investigation.

Most work in the area of gesture articulation feature analysis has focused on adults (Blagojevic et al., 2010; Rubine, 1991; Vatavu et al., 2013a), though there have been a few exceptions, outlined in this section. There remains, however, a need for a more comprehensive study of gestures produced by all ages, particularly of very young children (e.g., younger than age 5), to provide a fuller characterization of children's gesture articulation patterns. There are, unsurprisingly, a large number of features that have been employed to analyze and characterize adults' touchscreen gestures. Three of the most popular sets of features include Rubine's (Rubine, 1991), which were used for a feature based recognizer; Long et al.'s (Long et al., 2000) features, used for studying adults' gestures; and Vatavu et al.'s relative accuracy measures (Vatavu et al., 2014), which provide a measure of consistency between two gestures of the same type. Many of the features in the latter two sets were adapted from Rubine's features. Due to their importance in the literature, we provide a brief description of the features introduced by Rubine (Rubine, 1991) and Long et al. (Long et al., 2000) in Table 5, and Vatavu et al.'s (Vatavu et al., 2013a) relative accuracy measures are listed in Table 6.

A. Children's Gestures

As previously mentioned, studies of the articulation features of children's touchscreen gestures have been quite limited. Anthony et al. (Anthony et al., 2015) computed several "simple" and "complex" features of gestures of 10- to 17-year-olds, as well as adults, examining the impact of age and visual feedback. The simple features were taken from previous work examining children's gestures (Anthony et al., 2013b; Brown et al., 2013) and included, for example, number of strokes, gesture area, and gesture speed. Some of the complex features, which had previously been used only to examine adults' gestures (Rubine, 1991; Vatavu et al., 2013a), included gesture sharpness and gesture curviness. Age group did not have a significant effect on any of the nine simple (absolute) features and had a significant effect on only one of the seven complex (relative accuracy) features, according to a one-way ANOVA. Another study aiming to improve characterizations of children's touchscreen gestures is that of Shaw and Anthony (Shaw and Anthony, 2016a), who analyzed gestures from children ages 5 to 10 years old, as well as adults. In the study, the authors calculated the values of 10 simple features calculated on a single gesture (Anthony et al., 2013b), such as the total length of the gesture and the area of the bounding box of the gesture, and 12 complex features (Vatavu et al., 2013a). In contrast to Anthony et al.'s (Anthony et al., 2015) analysis, Shaw and Anthony (Shaw and Anthony, 2016a) found a significant effect of age group on six of the 10 simple features and all 12 of the relative accuracy features. Furthermore, many of the features showed the same general trend of increasing with age, like recognition rates, helping to paint a partial picture of why recognition rates are so much lower for younger children than for older children and adults. For example, consider the size error feature, defined as the average distance between corresponding points of two gestures of the same type, which follows a similar pattern by age as recognition rates for the same

Feature	**Description**
Shape Error	The average deviation between two gestures based on Euclidean distance
Shape Variability	The standard deviation of the distances between the points of two gestures
Length Error	A measure of the inconsistency of lengths of strokes between two gestures
Size Error	A measure of the inconsistency between the areas of the bounding boxes of two gestures
Bending Error	The average of differences between corresponding turning angles of two gestures
Bending Variability	The standard deviation of differences between corresponding turning angles of two gestures of the same type
Time Error	The difference in the amount of time taken to articulate two gestures
Time Variability	The standard deviation of the differences of the timestamps of each individual point in a gesture
R_9	The standard deviation of the differences of the timestamps of each individual point in a gesture
Speed Error	The difference in speed of production of two gestures
Speed Variability	The standard deviation of differences in the speed of production of two gestures
Stroke Count Error	The difference in number of strokes of two gestures of the same type
Stroke Ordering Error	A measure of the inconsistency in the order that different strokes of a gesture are drawn

Table 6. The relative accuracy features introduced by Vatavu et al. (Vatavu et al., 2013a).

gestures. Shaw and Anthony (Shaw and Anthony, 2016a) show that younger children's gestures have a larger variation in size than older children, indicating less developed control in their articulation. This difference in size could, for example, cause recognition errors in recognizers that use size-based features as part of the recognition process. These articulation features help partially explain why older children's gestures are better recognized than those of younger children.

The analyses presented in this section provide a starting point for understanding the articulation patterns exhibited by children, but further work is needed. Anthony et al. (Anthony et al., 2015) found that age group had a significant effect on only one feature in their study of 10- to 17-year-olds, while Shaw and Anthony's (Shaw and Anthony, 2016a) analysis found a significant impact of age on most of the features in their study of 5- to 10-year-olds, showing that younger children are more inconsistent in their gesturing patterns than older children and adults. Further work should look at new features that could potentially aid in understanding children's gestures in the 5- to 10-year-old age group. Finally, children as young as two years old are using touchscreen devices (Ahearne et al., 2016), but there has thus far been no in-depth characterization of the articulation features in the 2- to 5-year-old age range. Support for children with motor impairments affecting dexterity and fine motor skills might also be improved by analyzing their gesture features. Thus, future work should address this area to help improve understanding of children's touchscreen gesture articulation in order to improve recognition of their gestures.

VI. Summary and Future Work

In this paper, we have identified several open questions in gesture recognition. We now provide a summary of areas that we have described in this paper as well as areas for future work to improve the state of gesture interactions for children.

Our survey of existing work in gesture recognition for children has led us to several conclusions. Prior work has clearly not been sufficient to give a full idea of how existing recognizers perform on recognizing children's gestures since only simple template matchers have been tested. Past work describes behavior in terms of articulation features that may affect recognition, but there has been no mapping between children's gesture articulation behavior and recognition rates. Furthermore, the articulation features that have been used to examine children's gestures have been the same as those used for adults' gestures, but new articulation features based on children's gesture interactions may help developers better understand why recognition rates are poor.

An obvious next step in gesture interaction for children is to compare non-template based recognizers' performance in recognizing children's gestures. Table 4 of this paper gives a sample of recognizers that could be evaluated on children's gestures to improve the state of our understanding.

Based on our survey, we identify several areas for future work in gesture recognition for children, which we discuss here.

Characterizing young children's gestures. While there have been a number of studies characterizing children's gestures, most have focused on children ages 5 and older. Thus, little is known about the articulation and interaction patterns of children of younger ages. However, children are using touchscreens at these young ages (Common Sense Media, 2013), and are undergoing important developmental and physiological changes that could impact their touchscreen interactions in interesting and important ways. For example, young children in preschool are developing their writing skills and are developing their motor skills, which could impact gesture recognition. Improved characterizations of children's gestures can help us understand the specific aspects of the gestures that cause them to be poorly recognized, allowing developers to improve recognition by accounting for these aspects. For example, an application could apply a correction as a preprocessing step or ignore certain parts of the gesture that exhibit high variance.

Designing child-centered gesture sets. Traditional gesture sets have been designed for adults, but we have summarized prior work that has illustrated many ways in which children's gestures are different. Elicitation and guessability studies for children using touchscreen devices could provide useful information for designing gesture-based applications for children, especially since children are more likely to try novel new gestures when interacting with touchscreen devices (Rust et al., 2014; Soni et al., 2019b). Considering children's motor control development can also inform the design of better gesture sets. Furthermore, improved gesture set design could lead directly into improved recognition due to higher suitability for the target age groups. A new gesture set could be designed specifically to help children practice types of gestures that often cause them trouble.

New metrics for understanding children's gestures. A potential area for further exploration is the use of articulation features to help identify the specific behaviors that lead to poor recognition. For example, children often have trouble joining the ends of their strokes together, leading to large distances between endpoints that are intended to meet. A new articulation feature could be designed to quantify this behavior so that its relationships to recognition rates could be analyzed. In fact, this could be done using various recognizers to help establish which behaviors are correlated to their performance.

Improving recognition rates for children. As we have established, children's gestures are not recognized as well as those of adults. Thus, a major open area for future work is improving recognition rates. Existing gesture recognition algorithms have been designed to recognize adults' gestures, but a child-centered recognizer could offer the opportunity to specifically address common patterns in children's gestures. Most work on recognizing children's gestures so far has used template matching approaches like $N and $P (Anthony and Wobbrock, 2012; Vatavu et al., 2012), which may be more susceptible to failure given the inconsistencies children exhibit. However, these recognizers were not intended to be state-of-the-art approaches to recognition, but rather to be simple, easy to implement algorithms for rapid prototyping. Thus, we encourage future work on other recognition approaches to understand how well state-of-the-art approaches can perform. Future work can examine other recognition approaches, and new mixed methods for recognizing children's gestures.

Another possible method of improving recognition accuracy for children's gestures may be to employ some sort of preprocessing step in which the raw gesture data is transformed in some way. For example, consider the process of *beautification* (Julia and Faure, 1995). Beautification transforms raw input data into more aesthetically pleasing ("beautified") input based on the geometric properties of the gesture. A beautification algorithm might, for example, examine each stroke and determine whether it most resembles a line or a curve, then transform the data to the correct shape. After each stroke is reshaped, the algorithm examines the relationships between the strokes to determine whether the strokes should be joined. The final, "beautified" gesture consists of perfect lines and curves and is more likely to resemble the canonical form of the gesture. This process may be more difficult for children's gestures than for those of adults, but it may present an opportunity for obtaining significantly improved recognition results without development of a new recognizer.

Supporting improved educational technology. Educational technology and intelligent tutoring systems are increasingly using gesture-based interactions and recognition. The widespread nature of touchscreen devices means that educational systems have the potential to offer meaningful instruction to children who may not have easy access to educational resources or personalized instruction. For example, a highly adaptive touchscreen tutoring system may be able to provide feedback to children when a personal tutor is not available. The tactile nature of touchscreen devices can serve as a natural benefit for children developing their motor skills (Lovato and Waxman, 2016), which educational systems can use to their advantage. A key area of research beyond improving recognition is to examine the best ways to leverage recognition systems to achieve the goals of improving educational experiences.

Gesture interactions are an important component of touchscreen usage, and children's gesture interactions present a particular challenge in that they are much less consistent than those of adults. Recent advances in our understanding of both children's and adults' gestures have significantly improved the state of touchscreen gesture interaction, but the low recognition rates for younger children's gestures indicate a need for further examination of their gestures. Due to the small body of work on the recognition of children's gestures, we have surveyed existing work on both children and adult's gesture interactions with the intent to help motivate future areas of work that may lead to improved understanding of children's behavior when making gestures. This improved understanding can, in turn, help inform future work on recognizing children's gestures. Better recognition rates for children's gesture interactions could have a number of benefits, especially pertaining to systems using intelligent gesture-based interfaces to support learning.

References

Abdul Aziz, N.A., Batmaz, F., Stone, R., Chung, P.W.H., 2013. Selection of touch gestures for children's applications, in: Proceedings of the Science and Information Conference. pp. 721–726. https://doi.org/10.14569/IJACSA.2014.050415

Ahearne, C., Dilworth, S., Rollings, R., Livingstone, V., Murray, D., 2016. Touch-screen technology usage in toddlers. Arch. Dis. Child. 101, 181.

Ahmed, H., Azeem, S.A., 2011. On-line Arabic handwriting recognition system based on HMM, in: Proceedings of the International Conference on Document Analysis and Recognition. pp. 1324–1328. https://doi.org/10.1109/ICDAR.2011.266

Alimoglu, F., Alpaydin, E., 1997. Combining multiple representations and classifiers for pen-based handwritten digit recognition, in: Proceedings of the Fourth International Conference on Document Analysis and Recognition. IEEE Computing Society, pp. 637–640. https://doi.org/10.1109/ICDAR.1997.620583

Anderson, D., Bailey, C., Skubic, M., 2004. Hidden Markov Model Symbol Recognition for Sketch-Based Interfaces, in: AAAI Fall Symposium. AAAI Press, Menlo Park, CA, pp. 15–21.

Andrén, M., 2011. The Organization of Children's Pointing Stroke Endpoints, in: Integrating Gestures. pp. 153–162.

Anthony, L., 2019. Physical dimensions of children's touchscreen interactions: Lessons from five years of study on the MTAGIC project. Int. J. Hum. Comput. Stud. 128, 1–16. https://doi.org/https://doi.org/10.1016/j.ijhcs.2019.02.005

Anthony, L., Brown, Q., Nias, J., Tate, B., 2015. Children (and Adults) Benefit From Visual Feedback During Gesture Interaction on Mobile Touchscreen Devices. Int. J. Child-Computer Interact. 6, 17–27.

Anthony, L., Brown, Q., Nias, J., Tate, B., 2013a. Examining the need for visual feedback during gesture interaction on mobile touchscreen devices for kids, in: Proceedings of the International Conference on Interaction Design and Children (IDC '13). ACM Press, New York, New York, USA, New York, USA, pp. 157–164. https://doi.org/10.1145/2485760.2485775

Anthony, L., Brown, Q., Nias, J., Tate, B., Mohan, S., 2012a. Interaction and recognition challenges in interpreting children's touch and gesture input on mobile devices, in: Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '12). ACM Press, New York, New York, USA, pp. 225–234. https://doi.org/10.1145/2396636.2396671

Anthony, L., Vatavu, R.-D., Wobbrock, J.O., 2013b. Understanding the Consistency of Users' Pen and Finger Stroke Gesture Articulation, in: Proceedings of Graphics Interface (GI '13). Canadian Information Processing Society, pp. 87–94.

Anthony, L., Wobbrock, J.O., 2012. $ N-Protractor: A Fast and Accurate Multistroke Recognizer, in: Proceedings of Graphics Interface (GI '12). Canadian Information Processing Society, pp. 117–120.

Anthony, L., Wobbrock, J.O., 2010. A lightweight multistroke recognizer for user interface prototypes, in: Proceedings of Graphics Interface (GI ’10), GI ’10. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, pp. 245–252.

Anthony, L., Yang, J., Koedinger, K.R., 2012b. A paradigm for handwriting-based intelligent tutors. Int. J. Hum. Comput. Stud. 70, 866–887. https://doi.org/10.1016/j.ijhcs.2012.04.003

Appert, C., Zhai, S., 2009. Using Strokes as Command Shortcuts : Cognitive Benefits and Toolkit Support. Proc. Int. Conf. Hum. Factors Comput. Syst. 2289–2298. https://doi.org/10.1145/1518701.1519052

Apps in My Pocket Ltd, 2018. abc PocketPhonics: letter sounds & writing + first words.

Apte, A., Vo, V., Kimura, T.D., 1993. Recognizing multistroke geometric shapes, Proceedings of the 6th annual ACM symposium on User interface software and technology - UIST ’93. https://doi.org/10.1145/168642.168654

Arif, A.S., Sylla, C., 2013. A comparative evaluation of touch and pen gestures for adult and child users, in: Proceedings of the International Conference on Interaction Design and Children (IDC ’13). ACM Press, New York, New York, USA, pp. 392–395. https://doi.org/10.1145/2485760.2485804

Baum, L.E., Petrie, T., 1966. Statistical Inference for Probabilistic Functions of Finite State Markov Chains. Ann. Math. Stat. 37, 1554–1563. https://doi.org/10.1214/aoms/1177699147

Bellman, R., 1952. On the Theory of Dynamic Programming. Proc. Natl. Acad. Sci. U. S. A. 38, 716–719.

Blagojevic, R., Chang, S.H.H., Plimmer, B., 2010. The power of automatic feature selection: rubine on steroids. Jt. Sess. Seventh Sketch-Based Interfaces Model. Work. Eighth Symp. Non-Photorealistic Animat. Render. 79–86.

Brewer, R., Anthony, L., Brown, Q., Irwin, G., Nias, J., Tate, B., 2013. Using gamification to motivate children to complete empirical studies in lab environments, in: Proceedings of the International Conference on Interaction Design and Children (IDC ’13). ACM Press, New York, New York, USA, New York, USA, pp. 388–391. https://doi.org/10.1145/2485760.2485816

Brown, Q., Anthony, L., 2012. Toward comparing the touchscreen interaction patterns of kids and adults, in: Proceedings of the SIGCHI Workshop on Educational Software, Interfaces and Technology. p. 4pp.

Brown, Q., Anthony, L., Nias, J., Tate, B., Brewer, R., Irwin, G., 2013. Towards Designing Adaptive Touch-Based Interfaces, in: Proceedings of the ACM SIGCHI Mobile Accessibility Workshop. p. 4pp.

Chen, M.-Y., 1994. Off-Line Handwritten Word Recognition Using a Hidden Markov Model Type Stochastic Network. IEEE Trans. Pattern Anal. Mach. Intell. 16, 481–496. https://doi.org/10.1109/34.291449

Cheng, H., Member, S., Yang, L., Liu, Z., 2016. Survey on 3D Hand Gesture Recognition. IEEE Trans. Circuits Syst. Video Technol. 26, 1659–1673.

Cheng, Y., Ji, X., Li, X., Zhang, T., Malebary, S., Qu, X., Xu, W., 2020. Identifying Child Users via Touchscreen Interactions. ACM Trans. Sens. Networks 16. https://doi.org/10.1145/3403574

Cho, M.G., 2006. A new gesture recognition algorithm and segmentation method of Korean scripts for gesture-allowed ink editor. Inf. Sci. (Ny). 176, 1290–1303. https://doi.org/10.1016/j.ins.2005.04.006

Common Sense Media, 2017. The Common Sense Census: Media Use by Kids Age Zero to Eight 2017.

Common Sense Media, 2013. Zero to Eight: Children’s Media Use in America 2013 [WWW Document]. URL https://www.commonsensemedia.org/research/zero-to-eight-childrens-media-use-in-america-2013

Connell, S., Kuo, P.-Y., Liu, L., Piper, A.M., 2013. A Wizard-of-Oz elicitation study examining child-defined gestures with a whole-body interface, in: Proceedings of the International Conference on Interaction Design and Children. pp. 277–280. https://doi.org/10.1145/2485760.2485823

Connell, S.D., Jain, A.K., 2001. Template-based online character recognition. Pattern Recognit. 34, 1–14. https://doi.org/10.1016/S0031-3203(99)00197-1

Cornell, J.M., 1985. Spontaneous Mirror-Writing in Children. Can. J. Psychol. 39, 174–179. https://doi.org/10.1037/h0080122

Crescenzi Lanna, L., Grané Oro, M., 2019. Touch gesture performed by children under 3 years old when drawing and coloring on a tablet. Int. J. Hum. Comput. Stud. 124, 1–12. https://doi.org/https://doi.org/10.1016/j.ijhcs.2018.11.008

El-hajj, R., Likforman-sulem, L., Mokbel, C., 2005. Arabic Handwriting Recognition Using Baseline Dependant Features and Hidden Markov Modeling, in: Proceedings of the International Conference on Document Analysis and Recognition (ICDAR ’05). IEEE, pp. 6–10.

Findlater, L., Froehlich, J., Fattal, K., Wobbrock, J.O., Dastyar, T., 2013. Age - Related Differences in Performance with Touchscreens Compared to Traditional Mouse Input, in: Proceedings of the International Conference on Human Factors in Computing (CHI ’12). pp. 1–4. https://doi.org/10.1145/2470654.2470703

Flewitt, R., Messer, D., Kucirkova, N., 2015. New directions for early literacy in a digital age: The iPad. J. Early Child. Lit. 15, 289–310. https://doi.org/10.1177/1468798414533560

Gader, P.D., Mohamed, M.A., Keller, J.M., 1996. Fusion of handwritten word classifiers. Pattern Recognit. Lett. 17, 577–584. https://doi.org/10.1016/0167-8655(96)00021-9

Gardner, H., 1985. Frames of Mind: The Theory of Multiple Intelligences. Theory Mult. Intell. https://doi.org/10.2307/3324261

Gathercole, S.E., 1999. Cognitive Approaches to the Development of Short-Term Memory. Trends Cogn. Sci. 3, 410–419.

Herold, J., Stahovich, T.F., 2012. The 1¢ Recognizer: a fast, accurate, and easy-to-implement handwritten gesture recognition technique, in: Proceedings of the International Symposium on Sketch-Based Interfaces

and Modeling. Eurographics Association, pp. 39–46.
Hiniker, A., Sobel, K., Hong, S.R., Suh, H., Irish, I., Kim, D., Kientz, J.A., 2015. Touchscreen Prompts for Preschoolers: Designing Developmentally Appropriate Techniques for Teaching Young Children to Perform Gestures, in: Proceedings of the International Conference on Interaction Design and Children. ACM Press, pp. 109–118. https://doi.org/10.1145/2771839.2771851
Hong, J.I., Landay, J.A., 2000. SATIN: a toolkit for informal ink-based applications, in: Proceedings of the International Symposium on User Interface Software and Technology. ACM Press, New York, New York, USA, pp. 63–72. https://doi.org/10.1145/354401.354412
Hse, H.H., Richard Newton, A., 2005. Recognition and beautification of multi-stroke symbols in digital ink. Comput. Graph. 29, 533–546. https://doi.org/10.1016/j.cag.2005.05.006
Hu, J., Brown, M.K., 1996. HMM Based On-Line Handwriting Recognition. IEEE Trans. Pattern Anal. Mach. Intell. 18, 1039–1045.
Hu, J., Lim, S.G., Brown, M.K., 2000. Writer independent on-line handwriting recognition using an HMM approach. Pattern Recognit. 33, 133–147. https://doi.org/https://doi.org/10.1016/S0031-3203(99)00043-6
Huber, B., Tarasuik, J., Antoniou, M.N., Garrett, C., Bowe, S.J., Kaufman, J., 2016. Young children's transfer of learning from a touchscreen device. Comput. Human Behav. 56, 56–64. https://doi.org/https://doi.org/10.1016/j.chb.2015.11.010
Imangi, 2011. Temple Run.
Jego, J.F., Paljic, A., Fuchs, P., 2013. User-defined gestural interaction: A study on gesture memorization, in: Proceedings of the IEEE Symposium on 3D User Interfaces. pp. 7–10. https://doi.org/10.1109/3DUI.2013.6550189
Jiang, W., Sun, Z., 2005. HMM-Based Online Multi-Stroke Sketch Recognition, in: Proceedings of the Fourth International Conference on Machine Learning and Cybernetics (ICMLC '05). pp. 18–21.
Johnson, B., Turner, L.A., 2002. Data Collection Strategies in Mixed Methods Research, in: Handbook of Mixed Methods in Social & Behavioral Research. pp. 297–319.
Julia, L., Faure, C., 1995. Pattern recognition and beautification for a pen based interface, in: Proceedings of the International Conference on Document Analysis and Recognition. pp. 58–63.
Kakebeeke, T.H., Caflisch, J., Chaouch, A., Rousson, V., Largo, R.H., Jenni, O.G., 2013. Neuromotor development in children. Part 3: motor performance in 3- to 5-year-olds. Dev. Med. Child Neurol. 55, 248–256. https://doi.org/10.1111/dmcn.12034
Kim, H., Taele, P., Seo, J., Liew, J., Hammond, T., 2016. EasySketch2 : A novel sketch-based interface for improving children's fine motor skills and school readiness, in: Association, E. (Ed.), Proceedings of the International Symposium on Sketch-Based Interfaces and Modeling. pp. 69–78.
Kim, H., Taele, P., Valentine, S., McTigue, E., Hammond, T., 2013. 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 (SBIM '13). ACM Press, New York, New York, USA, pp. 33–42. https://doi.org/10.1145/2487381.2487389
King.com Ltd., 2017. Candy Crush Saga.
Konstantinos G. Derpanis, 2004. A Review of Vision-Based Hand Gestures.
Kratz, S., Rohs, M., 2010. The $3 Recognizer: Simple 3D Gesture Recognition on Mobile Devices, in: Proceedings of the International Conference on Intelligent User Interfaces. pp. 419–420. https://doi.org/10.1145/1719970.1720051
Kristensson, P.-O., Zhai, S., 2004. SHARK^2: A Large Vocabulary Shorthand Writing System for Pen-Based Computers. Proc. 17th Annu. ACM Symp. User interface Softw. Technol. - UIST '04 6, 43–52. https://doi.org/10.1145/1029632.1029640
Lauricella, A.R., Pempek, T.A., Barr, R., Calvert, S.L., 2010. Contingent computer interactions for young children's object retrieval success. J. Appl. Dev. Psychol. 31, 362–369. https://doi.org/https://doi.org/10.1016/j.appdev.2010.06.002
LeCun, Y., Boser, B.E., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W.E., Jackel, L.D., 1990. Handwritten Digit Recognition with a Back-Propagation Network. Adv. Neural Inf. Process. Syst. 2, 396–404.
Lee, S.W., 1996. Off-line recognition of totally unconstrained handwritten numerals using multilayer cluster neural network. IEEE Trans. Pattern Anal. Mach. Intell. 18, 648–652. https://doi.org/10.1109/34.506416
Lee, W., Burak Kara, L., Stahovich, T.F., 2007. An efficient graph-based recognizer for hand-drawn symbols. Comput. Graph. 31, 554–567. https://doi.org/10.1016/j.cag.2007.04.007
Li, Y., 2010. Protractor: a fast and accurate gesture recognizer. Proc. 28th Int. Conf. Hum. factors Comput. Syst. 2169–2172. https://doi.org/10.1145/1753326.1753654
Long, A.C., Landay, J.A., Rowe, L.A., Michiels, J., 2000. Visual similarity of pen gestures, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '00. ACM Press, New York, New York, USA, pp. 360–367. https://doi.org/10.1145/332040.332458
Long, A.C., Landay, J. a, Rowe, L. a, 1999. Implications For a Gesture Design Tool. Proc. Int. Conf. Hum. Factors Comput. Syst. 40–47. https://doi.org/10.1145/302979.302985
Lovato, S.B., Waxman, S.R., 2016. Young Children Learning from Touch Screens: Taking a Wider View. Front. Psychol. 7, 1078. https://doi.org/10.3389/fpsyg.2016.01078
Malmi, L., Korhonen, A., 2004. Automatic feedback and resubmissions as learning aid, in: Proceedings of the IEEE International Conference on Advanced Learning Technologies (ICALT '04). IEEE Computing Society, pp. 186–190.

https://doi.org/10.1109/ICALT.2004.1357400
Mauney, D., Howarth, J., Wirtanen, A., Capra, M., 2010. Cultural Similarities and Differences in User-Defined Gestures for Touchscreen User Interfaces. Ext. Abstr. 28th Int. Conf. Hum. Factors Comput. Syst. 4015–4020. https://doi.org/10.1145/1753846.1754095
McKnight, L., Fitton, D., 2010. Touch-screen technology for children, in: Proceedings of the International Conference on Interaction Design and Children. ACM Press, New York, New York, USA, pp. 238–241. https://doi.org/10.1145/1810543.1810580
Mitchell, T.M., 1997. Machine Learning, McGraw-Hill. https://doi.org/10.1145/242224.242229
Mitra, S., Acharya, T., 2007. Gesture Recognition : A Survey. IEEE Trans. Syst. Man. Cybern. 37, 311–324.
Mohamad, R.A., Likforman-sulem, L., 2009. Combining Slanted-Frame Classifiers for Improved HMM-Based Arabic Handwriting Recognition. IEEE Trans. Pattern Anal. Mach. Intell. 31, 1165–1177.
Morris, M.R., Danielescu, A., Drucker, S., Fisher, D., Lee, B., Schraefel, C., Wobbrock, J.O., 2014. Reducing Legacy Bias in Gesture Elicitation Studies. Interactions 21, 40–45. https://doi.org/10.1145/2591689
Morris, M.R., Wobbrock, J.O., Wilson, A.D., 2010. Understanding users' preferences for surface gestures, in: Proceedings of Graphics Interface 2010. pp. 261–268. https://doi.org/10.1016/j.actamat.2009.07.058
Murthy, G.R.S., Jadon, R.S., 2009. A Review of Vision Based Hand Gestures Recognition, International Journal of Information Technology and Knowledge Management.
Myers, B.A., Giuse, D.A., Dannenberg, R.B., Vander Zanden, B., Kosbie, D.S., Pervin, E., Mickish, A., Marchal, P., 1990. Garnet: Comprehensive Support for Graphical, Highly Interactive User Interfaces. Computer (Long. Beach. Calif). 23, 71–85. https://doi.org/10.1109/2.60882
Myers, B.A., McDaniel, R., Michish, A., Klimovitski, A., 1995. The Design for the Amulet User Interface Toolkit. Hum. Comput. J. 9.
Nacenta, M.A., Kamber, Y., Qiang, Y., Kristensson, P.O., 2013. Memorability of pre-designed and user-defined gesture sets, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13. ACM Press, New York, New York, USA, p. 1099. https://doi.org/10.1145/2470654.2466142
Nacher, V., Jaen, J., Navarro, E., Catala, A., Gonzalez, P., 2015. Multi-touch gestures for pre-kindergarten children. Int. J. Hum. Comput. Stud. 73, 37–51. https://doi.org/10.1016/j.ijhcs.2014.08.004
Nakai, M., Akira, N., Shimodaira, H., Sagayama, S., 2001. Substroke Approach to HMM-based On-line Kanji Handwriting Recognition, in: Proceedings of the International Conference on Document Analysis and Recognition (ICDAR '01). IEEE Computing Society, pp. 491–495.
Nathan, M.J., Alibali, M.W., Koedinger, K.R., 2003. Expert Blind Spot: Where Content Knowledge & Pedagogical Content Knowledge Collide.
Newell, K.M., 1991. Motor skill acquisition. Annu. Rev. Psychol. 42, 213–237. https://doi.org/10.1146/annurev.ps.42.020191.001241
Nguyen, T., Roy, A., Memon, N., 2019. Kid on the phone! Toward automatic detection of children on mobile devices. Comput. Secur. 84, 334–348. https://doi.org/https://doi.org/10.1016/j.cose.2019.04.001
Nitrome, 2015. Magic Touch: Wizard for Hire.
Olsen, L., Samavati, F.F., Sousa, M.C., 2007. Fast Stroke Matching by Angle Quantization. Int. Conf. Immersive Telecommun. https://doi.org/10.4108/ICST.IMMERSCOM2007.2114
Olsen, L., Samavati, F.F., Sousa, M.C., Jorge, J.A., 2009. Sketch-based modeling: A survey. Comput. Graph. 33, 85–103. https://doi.org/https://doi.org/10.1016/j.cag.2008.09.013
Piaget, J., 1983. Piaget's Theory, in: Mussen, P. (Ed.), Handbook of Child Psychology. Wiley & Sons, New York, NY, USA.
Plötz, T., Fink, G.A., 2009. Markov models for offline handwriting recognition: a survey. Int. J. Doc. Anal. Recognit. 12, 269. https://doi.org/10.1007/s10032-009-0098-4
Punch, S., 2002. Research with Children: The same or different from research with adults? Childhood 9, 321–341.
Ravindran, M., 2010. Survey on Various Gesture Recognition Techniques for Interfacing Machines Based on Ambient Intelligence. Int. J. Comput. Sci. Eng. Surv. 1, 31–42.
Roy, P.P., Bhunia, A.K., Das, A., Dey, P., Pal, U., 2015. HMM-based Indic handwritten word recognition using zone segmentation. Pattern Recognit. https://doi.org/10.1016/j.patcog.2016.04.012
Rubine, D., 1991. Specifying gestures by example. ACM SIGGRAPH Comput. Graph. 25, 329–337. https://doi.org/10.1145/127719.122753
Rust, K., Malu, M., Anthony, L., Findlater, L., 2014. Understanding childdefined gestures and children's mental models for touchscreen tabletop interaction. Proc. Internatioanl Conf. Interact. Des. Child. (IDC '14) 201–204. https://doi.org/10.1145/2593968.2610452
Sameer Singh, A.A., n.d. Neural Network Recognition of Hand-Printed Characters.
Seifert, K., Hoffnung, R.J., 1987. Child and Adolescent Development. Houghton Mifflin, Boston.
Sezgin, T.M., Davis, R., 2005. HMM-based efficient sketch recognition, in: Proceedings of the International Conference on Intelligent User Interfaces. ACM Press, New York, New York, USA, pp. 281–283. https://doi.org/10.1145/1040830.1040899
Shaw, A., 2017. Human-Centered Recognition of Children's Touchscreen Gestures, in: Proceedings of the 19th ACM International Conference on Multimodal Interaction, ICMI '17. Association for Computing Machinery, New York, NY, USA, pp. 638–642. https://doi.org/10.1145/3136755.3137033
Shaw, A., Anthony, L., 2016a. Analyzing the Articulation Features of Children's Touchscreen Gestures, in: Proceedings of the International Conference on Multimodal Interaction (ICMI '16). ACM Press, pp. 333–340. https://doi.org/10.1145/2993148.2993179

Shaw, A., Anthony, L., 2016b. Toward a Systematic Understanding of Children's Touchscreen Gestures, in: Extended Abstracts of the ACM SIGCHI Conference on Human Factors in Computing Systems. p. 1752 - 1759.

Shaw, A., Anthony, L., Ruiz, J., 2017. Comparing Human and Machine Recognition of Children's Touchscreen Stroke Gestures, in: Proceedings of the ACM International Conference on Multimodal Interaction (ICMI '17). ACM Press, New York, New York, USA, pp. 32 - 40. https://doi.org/10.1145/3136755.3136810

Shrivastava, V., Sharma, N., 2012. Artificial Neural Network Based Optical Character Recognition. Signal Image Process. An Int. J. 3, 73–80.

Shuler, C., 2009. Pockets of Potential: Using Mobile Technologies to Promote Children's Learning. Joan Ganz Cooney Center at Sesame Workshop, New York, NY.

Singh, S., Amin, A., 2014. Neural Network Recognition of Hand-printed Characters. Neural Comput. Appl. 8, 67–76. https://doi.org/10.1007/s005210050008

Smithies, S., Novins, K., Arvo, J., 1999. A handwriting-based equation editor, in: Proceedings of Graphics Interface. pp. 84–91.

Soni, N., Aloba, A., Morga, K.S., Wisniewski, P.J., Anthony, L., 2019a. A Framework of Touchscreen Interaction Design Recommendations for Children (TIDRC): Characterizing the Gap between Research Evidence and Design Practice, in: Proceedings of the International Conference on Interaction Design and Children (IDC '19). Boise, ID, p. To appear.

Soni, N., Gleaves, S., Neff, H., Morrison-Smith, S., Esmaeili, S., Mayne, I., Bapat, S., Schuman, C., Stofer, K.A., Anthony, L., 2019b. Do User-Defined Gestures for Flatscreens Generalize to Interactive Spherical Displays for Adults and Children?, in: Proceedings of the International Symposium on Pervasive Displays (PerDis '19). p. to appaer.

Taranta II, E.M., LaViola Jr., J.J., 2015. Penny pincher: a blazing fast, highly accurate $-family recognizer 195–202.

Toast Games, 2016. We Are Magic.

Valentine, S., Vides, F., Lucchese, G., Turner, D., 2012. Mechanix: A Sketch-Based Tutoring System for Statics Courses. 24th Annu. Conf. Innov. Appl. Artif. Intell. 2253–2260.

Vanderdonckt, J., Roselli, P., Pérez-Medina, J.L., 2018. !!FTL, an Articulation-Invariant Stroke Gesture Recognizer with Controllable Position, Scale, and Rotation Invariances, in: Proceedings of the 20th ACM International Conference on Multimodal Interaction, ICMI '18. Association for Computing Machinery, New York, NY, USA, pp. 125–134. https://doi.org/10.1145/3242969.3243032

Vatavu, R.-D., Anthony, L., Wobbrock, J.O., 2018. $Q: A Super-Quick, Articulation-Invariant Stroke-Gesture Recognizer for Low-Resource Devices, in: Proceedings of the International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18). ACM, New York, NY, USA, pp. 623–635.

Vatavu, R.-D., Anthony, L., Wobbrock, J.O., 2014. Gesture Heatmaps: Understanding Gesture Performance with Colorful Visualizations, in: Proceedings of the ACM International Conference on Multimodal Interaction. ACM Press, New York, New York, USA, pp. 172–179. https://doi.org/10.1145/2663204.2663256

Vatavu, R.-D., Anthony, L., Wobbrock, J.O., 2013a. Relative accuracy measures for stroke gestures, in: Proceedings of the ACM International Conference on Multimodal Interaction. ACM Press, New York, New York, USA, pp. 279–286. https://doi.org/10.1145/2522848.2522875

Vatavu, R.-D., Anthony, L., Wobbrock, J.O., 2012. Gestures as point clouds: a $P recognizer for user interface prototypes, in: Proceedings of the ACM International Conference on Multimodal Interaction (ICMI '12). ACM Press, New York, New York, USA, pp. 273–280. https://doi.org/10.1145/2388676.2388732

Vatavu, R.-D., Casiez, G., Grisoni, L., 2013b. Small, medium, or large?: estimating the user-perceived scale of stroke gestures. Proc. Int. Conf. Hum. Factors Comput. Syst. 277–280. https://doi.org/10.1145/2470654.2470692

Vatavu, R.D., Anthony, L., Brown, Q., 2015a. Child or adult? Inferring Smartphone users' age group from touch measurements alone, in: INTERACT. pp. 1–9. https://doi.org/10.1007/978-3-319-22723-8_1

Vatavu, R.D., Cramariuc, G., Schipor, D.M., 2015b. Touch interaction for children aged 3 to 6 years: Experimental findings and relationship to motor skills. Int. J. Hum. Comput. Stud. 74, 54–76. https://doi.org/10.1016/j.ijhcs.2014.10.007

Verma, B., Gader, P., Chen, W.-T., 2001. Fusion of multiple handwritten word recognition techniques. Pattern Recognit. Lett. 22, 991–998. https://doi.org/10.1016/S0167-8655(01)00046-0

Willems, D., Niels, R., van Gerven, M., Vuurpijl, L., 2009. Iconic and multi-stroke gesture recognition. Pattern Recognit. 42, 3303–3312. https://doi.org/10.1016/j.patcog.2009.01.030

Williford, B., Runyon, M., Malla, A.H., Li, W., Linsey, J., Hammond, T., 2017. ZenSketch: A Sketch-based Game For Improving Line Work, in: Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play (CHI PLAY EA '17), CHI PLAY '17 Extended Abstracts. ACM, New York, NY, USA, pp. 591–598. https://doi.org/10.1145/3130859.3130861

Wobbrock, J.O., Aung, H.H., Rothrock, B., Myers, B.A., 2005. Maximizing the guessability of symbolic input, in: Proceedings of CHI Extended Abstracts on Human Factors in Computing Systems (CHI EA '05). ACM Press, New York, New York, USA, p. 1869. https://doi.org/10.1145/1056808.1057043

Wobbrock, J.O., Morris, M.R., Wilson, A.D., 2009. User-defined gestures for surface computing. Proc. Int. Conf. Hum. Factors Comput. Syst. (CHI 09) 1083. https://doi.org/10.1145/1518701.1518866

Wobbrock, J.O., Wilson, A.D., Li, Y., 2007. Gestures without libraries, toolkits or training: a $1 recognizer for user interface prototypes, in: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '07). ACM Press, New York, New York, USA, New York, USA, pp. 159–168. https://doi.org/10.1145/1294211.1294238

Woodward, J., Shaw, A., Aloba, A., Jain, A., Ruiz, J., Anthony, L., 2017. Tablets, tabletops, and smartphones: cross-platform comparisons of children's touchscreen interactions, in: Proceedings of the International Conference on Multimodal Interaction (ICMI '17). ACM Press, pp. 5–14.

Woodward, J., Shaw, A., Luc, A., Craig, B., Das, J., Hall, P., Hollay, A., Irwin, G., Sikich, D., Brown, Q., Anthony, L., 2016. Characterizing How Interface Complexity Affects Children's Touchscreen Interactions, in: Proceedings of the ACM International Conference on Human Factors in Computing Systems (CHI '16). ACM Press, San Jose, CA, USA, CA, USA, pp. 1921–1933. https://doi.org/10.1145/2858036.2858200

Xiaolin Li, Dit-Yan Yeung, 1997. On-line handwritten alphanumeric character recognition using dominant points in strokes. Pattern Recognit. 30, 31–44. https://doi.org/10.1016/S0031-3203(96)00052-0

Ye, Y., Nurmi, P., 2015. Gestimator: Shape and Stroke Similarity Based Gesture Recognition 219–226. https://doi.org/10.1145/2818346.2820734

Yin, J., Sun, Z., 2005. An Online Multi-stroke Sketch Recognition Method, in: Affective Computing and Intelligent Interaction. Springer, pp. 803–810.

Zhai, S., Kristensson, P.O., Appert, C., Andersen, T.H., Cao, X., 2012. Foundational Issues in Touch-Surface Stroke Gesture Design — An Integrative Review. Found. Trends® Human–Computer Interact. 5, 97–205. https://doi.org/10.1561/1100000012

Zhang, Q., Lee, K.-C., Bao, H., You, Y., Li, W., Guo, D., 2018. Large scale classification in deep neural network with Label Mapping. CoRR abs/1806.02507.