Child or Adult? Inferring Smartphone Users' Age Group from Touch Measurements Alone

Radu-Daniel Vatavu[1], Lisa Anthony[2], Quincy Brown[3]

[1]University Stefan cel Mare of Suceava, Suceava, Romania
[2]Department of CISE, University of Florida, Gainesville, FL, USA
[3]Bowie State University, Bowie, MD, USA
vatavu@eed.usv.ro, lanthony@cise.ufl.edu, qbrown@bowiestate.edu

Abstract. We present a technique that classifies users' age group, *i.e.*, *child* or *adult*, from touch coordinates captured on touch-screen devices. Our technique delivered 86.5% accuracy (user-independent) on a dataset of 119 participants (89 children ages 3 to 6) when classifying each touch event one at a time and up to 99% accuracy when using a window of 7+ consecutive touches. Our results establish that it is possible to reliably classify a smartphone user *on the fly* as a *child* or an *adult* with high accuracy using only basic data about their touches, and will inform new, automatically adaptive interfaces for touch-screen devices.

Keywords: touch input, children, adults, age group, tap time, offset distance, touch accuracy, classifier, Bayes' rule, touch-screen, smartphone, experiment.

1 Introduction

As of 2013, 43% of adult Americans over age 16 own a tablet or e-reader and 56% of adult Americans own a smartphone [7,8], and touch-screen input has rapidly become the primary way many users interact with these mobile computing devices. This trend is especially true for children of pre-school and elementary school age (ages 3 to 10) [5]. Previous work has determined that touch and stroke gesture interaction behaviors differ between adults and children in ways that significantly impact gesture recognition and processing of users' touch input [2,3,6], *e.g.*, children's touches are more likely to be just outside the boundaries of their intended target [3] and their gestures are less likely to be recognized correctly than adults' gestures [2,3].

If touch-screen interfaces were able to determine on the fly whether the user is a child or an adult, applications could switch ad-hoc to processing or recognition algorithms tailored for one age group or the other. Such *adaptive touch interaction* mechanisms would assist users in their touch input, making them more accurate and more successful during touch-screen interaction. Unfortunately, research on age-detection classification is sparse, and what work exists so far focuses on gesture input [6], which provides a far richer source of classification features than simple taps; see for instance Blagojevic *et al*. [4] for an examination of 114 such features. On the other hand, touch input on its own is more challenging to interpret, because of the rudimentary features provided by current touch-screen hardware, typically limited to touch coordinates and timestamps and, in the best-case scenario, to rough approximations of touch pressure and area size.

We present in this work a technique for classifying users' age group, *i.e.*, *child* or *adult*, from touch input alone. In this work, a *child* is a person having 6 years at most; at this age, children's touch input behaviors are the most different from those of adults knowing that children's input performance and touch accuracy improve with age [3]. Our technique, using a Bayes' rule classifier, delivered 86.5% user-independent accuracy on a dataset of 119 participants (89 children ages 3 to 6) when classifying each touch event one at a time and 99% accuracy with a window of 7+ consecutive touch events. These results establish that it is possible to classify a user as a *child* or *adult* with high accuracy using only minimal data provided by all existing touch-screen hardware. Our contributions will inform the next generation of *adaptive user interfaces for touch-screen devices* that will be able to respond appropriately and *on the fly* to their users, without pre-configuration or the need of specialized sensors.

2 Touch input analysis to inform classifier design

We analyze in this section adults' and small children's touch input patterns using two measures that can be readily computed for any touch-screen device, *i.e.*, (1) the time between the moment when the user's finger touches the screen and the moment when the finger lifts off (TAP-TIME) and (2) the distance between the actual touch point where the user's finger was placed and the center of the touch target (OFFSET-DISTANCE); see Figure 1 for an illustration of these measures. We generate frequency distributions of adults' and children's tap times and offset distances (see Figure 2, next page) using the touch dataset of Vatavu *et a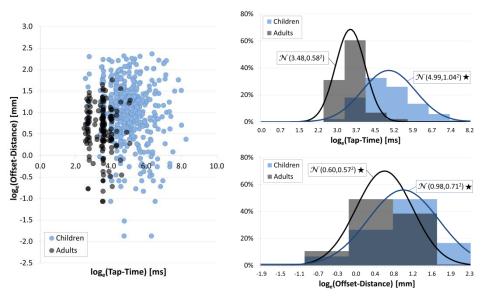l.* [9]. This dataset contains 587 samples from 119 participants (89 children with ages between 3 and 6 years old and 30 young adults)[1].

Fig. 1. TAP-TIME (τ) and OFFSET-DISTANCE (d) touch measures.

Figure 2 plots the *Child* and *Adult* classes in the TAP-TIME × OFFSET-DISTANCE space and the frequency distributions of the two measures. We found that all 4 distributions deviated significantly from normality, so we applied a log_e transform on our data. Follow-up Kolmogorov-Smirnov tests confirmed the normality of 3 out of the 4 log_e transformed distributions (*i.e.*, TAP-TIME for *Child*, OFFSET-DISTANCE for *Adult*, and OFFSET-DISTANCE for *Child*, which are marked with a star symbol ★ in Figure 2). The only distribution that still deviated significantly from normality was TAP-TIME for the *Adult* class ($D_{(149)} = 0.171, p < .01$). However, the effect size of the Kolmogorov-Smirnov test was considerably smaller than Cohen's suggested limit for small effect sizes ($r = .014 < .100$), while both skewness and kurtosis were close to zero ($s = -0.62$, CI$_{95\%}$= $[-0.81, -0.42]$ and $k = -0.76$, CI$_{95\%}$= $[-1.16, -0.36]$). These results support modeling log_e(TAP-TIME) and log_e(OFFSET-DISTANCE) data with normal distributions, which in turn leverage the option of designing a Bayes' rule classifier to discriminate between *Child* and *Adult* classes.

---

[1] The dataset of Vatavu *et al.* [9] is available to download at http://www.eed.usv.ro/~vatavu/ Note that, though there are more samples in the dataset from children than adults, all classification tests that we report in this paper were done with balanced data samples.

Fig. 2. Distribution of the *Child* and *Adult* classes in the TAP-TIME × OFFSET-DISTANCE space (left) and frequency histograms of each feature (right) with normal distributions superimposed. NOTE: we use log_e−transformed TAP-TIME and OFFSET-DISTANCE; a star symbol (★) indicates that the Kolmogorov-Smirnov test found no significant difference ($p > .01$) between our log_e-transformed data and a normal distribution $N(\mu, \sigma^2)$ with same mean and standard deviation.

Bayes' rule assigns a candidate measurement x to class j for which the *a posteriori* probability $p(class_j|x)$ is maximized [11] (p. 7):

$$p(class_j|x) = \frac{p(x|class_j) \cdot p(class_j)}{p(x)} \tag{1}$$

knowing the class-conditional density functions $p(x|class_j)$ (*i.e.*, the probability of a randomly chosen pattern x to lie with class j) and the *a priori* probabilities $p(class_j)$ (*i.e.*, how likely it is to see a pattern of each class); $p(x)$ is a normalizing factor that can be omitted when comparing *a posteriori* probabilities directly. For our problem, we have two classes only (*Child* and *Adult*), for which the *a priori* probabilities can be considered equal and thus can be omitted from eq. 1. Then, Bayes' rule becomes:

$$\text{Assign measurement } x \text{ to class} \begin{cases} Child & \text{if } p(x|Child) > p(x|Adult) \\ Adult & \text{otherwise} \end{cases} \tag{2}$$

With each measure and their combination, there are 3 possible classifier designs:

1. **Bayes' rule for TAP-TIME measurements** (*i.e.*, the TAP-TIME classifier). In this case, x is the log_e−transformed value of a TAP-TIME measurement collected from the user, $p(x|Child) = \frac{1}{\sqrt{2\pi}\sigma_c} exp\left\{-\frac{(x-\mu_c)^2}{2\sigma_c^2}\right\}$ and $p(x|Adult) = \frac{1}{\sqrt{2\pi}\sigma_a} exp\left\{-\frac{(x-\mu_a)^2}{2\sigma_a^2}\right\}$; μ_c and σ_c are the mean and standard deviation of TAP-TIME for the *Child* class, and μ_a and σ_a for *Adults*. For instance, $\mu_c = 4.99\ ms$ and $\sigma_c = 1.04\ ms$; see Figure 2, top-right (remember that these are log_e−transformed values).

2. **Bayes' rule for OFFSET-DISTANCE measurements** (OFFSET-DISTANCE classifier). In this case, x is the log_e−transformed value of an OFFSET-DISTANCE measurement,

$p(x|Child) = \frac{1}{\sqrt{2\pi}\sigma_c} exp\left\{-\frac{(x-\mu_c)^2}{2\sigma_c^2}\right\}$ and $p(x|Adult) = \frac{1}{\sqrt{2\pi}\sigma_a} exp\left\{-\frac{(x-\mu_a)^2}{2\sigma_a^2}\right\}$; μ_c and σ_c are the mean and standard deviation of OFFSET-DISTANCE for the *Child* class, and μ_a and σ_a for *Adults*. For instance, $\mu_a = 0.60\ ms$ and $\sigma_a = 0.57\ ms$; see Figure 2, bottom-right (remember that these are log_e-transformed values).

3. **Combined TIME & OFFSET-DISTANCE rule** (*i.e.*, the TIME-OFFSET classifier). In this case, x is a vector consisting of one TAP-TIME and one OFFSET-DISTANCE measurement, $p(x|Child) = \frac{1}{2\pi\sqrt{|\Sigma_c|}} exp\left\{-\frac{1}{2}(x-\mu_c)^T\Sigma_c^{-1}(x-\mu_c)\right\}$, $p(x|Adult) = \frac{1}{2\pi\sqrt{|\Sigma_a|}} exp\left\{-\frac{1}{2}(x-\mu_a)^T\Sigma_a^{-1}(x-\mu_a)\right\}$; μ_c is a vector containing the mean values of TAP-TIME and OFFSET-DISTANCE for the *Child* class and μ_a the mean vector for *Adults*; Σ_c and Σ_a represent the covariance matrices for the two classes. For instance, $\mu_c = [4.99, 0.98]$ and $\Sigma_c = \begin{bmatrix} 1.08 & -0.11 \\ -0.11 & 0.50 \end{bmatrix}$, where 1.08 is the variance (*i.e.*, square standard deviation) of TAP-TIME, 0.50 the variance of OFFSET-DISTANCE, and -0.11 the covariance between TAP-TIME and OFFSET-DISTANCE for the *Child* class (Figure 2, right).

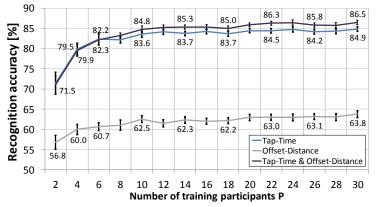## 3 Experiment #1: Classifying age group with one touch point only

We conducted a first experiment to evaluate the accuracy of the Bayes' rule classifiers to predict age group using measurements from *one touch point only*. We employed the touch dataset of Vatavu *et al.* [9], which contains 587 touch samples collected from 119 participants (89 children with ages between 3 and 6 years old and 30 young adults). The experiment was a within-subject design with two independent factors:

1. CLASSIFIER, nominal with 3 levels: TAP-TIME, OFFSET-DISTANCE, TIME-OFFSET.
2. The number of training participants P from which we estimated the mean and standard deviation values for TAP-TIME and OFFSET-DISTANCE employed by the class-conditional density functions $p(x|Child)$ and $p(x|Adult)$. We varied P between 2 and 30 participants for each class. (P=10 participants means that data from 5 children and 5 adults was used for training.) For each participant, a maximum of 5 samples are available in the dataset [9], making the size of the training set for estimating $p(x|Child)$ and $p(x|Adult)$ between 10 and 150.

We compute one dependent variable, the recognition ACCURACY of our classifiers, according to the following user-independent procedure from [10]: (1) for each P, we randomly select P children and P adults for training; (2) we randomly select 1 child and 1 adult for testing from the remaining participants, and we classify one randomly-selected touch point for each; (3) we repeat step 2 for 100 times for each training set, and we repeat step 1 for 100 times (*i.e.*, 100 different training sets). Overall, we report classification results from 15 (number of training participants P) × 100 (repetitions of each P) × 200 (classifications for each training set) = 300,000 classification trials. All tests are user-independent, so different data is used for the training and testing sets.

Friedman's test showed a significant effect of CLASSIFIER on recognition ACCURACY ($\chi^2_{(2,N=1500)} = 2304.714, p < .001$), and post-hoc Wilcoxon signed-rank tests showed significant differences for all classifier pairs ($p < .001$) with medium to large effect

sizes (r between .27 and .61); see Figure 3. The highest accuracy was delivered by the TIME-OFFSET classifier, which had an average performance of 83.9% and reached 86.5% with training data from 15 adults and 15 children. The TAP-TIME classifier came second with an average performance of 82.8% and maximum accuracy of 84.9%. The OFFSET-DISTANCE classifier exhibited the lowest performance with only 63.8% accuracy for the maximum number of P=30 training participants. Friedman's tests showed significant effects of P over all classifiers ($p < .001$), with ACCURACY increasing considerably for both TIME-OFFSET and TAP-TIME from P=2 to 8 participants (71% to 83%), after which it continued to increase slowly up to 86.5%.

Fig. 3. Recognition accuracy of Bayes' rule classifiers. NOTES: P=10 participants means that touch data from 5 children and 5 adults was used for training. Five touch measurements were used for each training participant. Error bars show 95% CIs.
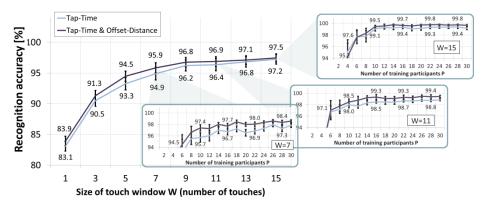

4 Experiment #2: Classifying age group with a touch window

The results from the previous section show good potential for two Bayes' rule classifiers to discriminate between *Child* and *Adult* classes, with accuracy up to 86.5%. In this section, we show how this accuracy rate can be much improved by employing a majority vote for which the classification decision is taken after analyzing several touch measurements in a row. A *touch window* represents a sequence of W consecutive touches. We classify each touch with one of our Bayes' rules and then count the number of *Child* and *Adult* votes. The majority count wins the classification and selecting an odd number for W guarantees no ties for our 2-class problem.

We conducted a second experiment to evaluate the accuracy of the Bayes' rule classifiers to predict users' age group using measurements from *touch windows*. The experiment was a within-subject design with three independent factors:
1. CLASSIFIER, nominal with 3 levels: TAP-TIME, OFFSET-DISTANCE, TIME-OFFSET.
2. The number of training participants $2 \leq P \leq 30$ from which we estimated the mean and standard deviation values for TAP-TIME and OFFSET-DISTANCE employed by the class-conditional density functions $p(x|Child)$ and $p(x|Adult)$.
3. The size of the touch window W with 1, 3, 5, 7, 9, 11, 13, and 15 touches.

We compute one dependent variable, the recognition ACCURACY of our classifiers, according to the following user-independent procedure [10]: (1) for each P value, we randomly select P children and P adults for training; (2) for each W, we randomly select W children touch samples and W adult touch samples from the remaining participants, and classify the two W-touch windows as *Child* or *Adult*; (3) we repeat step 2 for 100 times for each training set, and we repeat step 1 for 100 times (*i.e.*, 100 different training sets). Overall, we report results from 15 (number of training participants P) × 8 (sizes W of the touch window) × 100 (repetitions of each P) × 200 (classifications for each training set) = 2,400,000 classification trials. Because the dataset has a maximum of 5 touch samples per participant, we simulate larger windows by sampling across participants for this experiment. Since the training set shares no users in common with the testing set (user-independent), this approach is equivalent to off-the-shelf use cases.
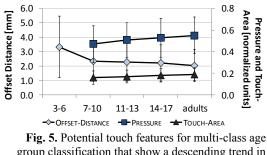

Fig. 4. Recognition accuracy of Bayes' rule classifiers on *touch windows*. The effect of number of participants P on accuracy shows an increasing trend overall; we highlight the effect of P for some of the touch windows (W=7, 11, and 15). For W=1, see Fig. 3. Error bars show 95% CIs.
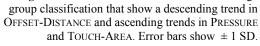
Friedman's test detected a significant effect of CLASSIFIER on ACCURACY ($\chi^2_{(2,N=12000)} = 18362.182, p < .001$) and post-hoc Wilcoxon signed-rank tests revealed significant differences between all classifier pairs (at $p < .001$) with medium to large effect sizes (r between .24 and .60). The OFFSET-DISTANCE classifier exhibited again the poorest performance (70.9% average accuracy), much lower than the other two classifiers (93.5% and 94.2%, respectively), so we disregard it from subsequent analysis. We found a significant effect of touch window size W on ACCURACY ($\chi^2_{(7,N=1500)} = 5975.910, p < .001$ for the TAP-TIME classification rule and $\chi^2_{(7,N=1500)} = 6280.743, p < .001$ for the TIME-OFFSET rule), with the average ACCURACY improving for both TAP-TIME and TIME-OFFSET from 83.1% and 83.9% for W=1 up to 97.2% and 97.5% for W=15; see Figure 4. For each touch window, performance improved with more training participants, *e.g.*, the TIME-OFFSET classifier delivered 97.1% accuracy with W=11 and P=4 participants (2 children + 2 adults), which increased to 99.4% when training data from P=30 participants (15 children + 15 adults) was used. Overall, the TIME-OFFSET classification rule delivered significantly better performance than TAP-TIME ($Z = -38.073, p < .001$), with a medium effect size ($r = .25$), while the actual average difference in accuracy was only 1.1%.

5 Discussion, conclusions, and future work

We showed that distance offset and tap time are enough to classify whether the user is a small *child* or an *adult* with reasonable accuracy (86.5%) using only one touch point and very high accuracy (>99%) with a window of 7+ touches. Although there are some limitations to our study, as we focused only on children ages 3 to 6 years, accuracy results are promising given the rudimentary data we worked with, *i.e.*, touch coordinates and timestamps alone. These results recommend our technique for discrimination of age groups at much finer granularity. However, we must leave such detailed investigations for future work, as they will most likely need examination of additional features to reach similar levels of accuracy. For example, Figure 5 shows two

Fig. 5. Potential touch features for multi-class age group classification that show a descending trend in OFFSET-DISTANCE and ascending trends in PRESSURE and TOUCH-AREA. Error bars show ± 1 SD.

such measures, PRESSURE and TOUCH-AREA for 5 age groups and 185 participants [3,9]. Because such measurements may not be available for all touch-screen hardware, fine age group discrimination may be restricted to some devices only. While we point the community to these opportunities, the exploration of such tradeoffs is future work.

Meanwhile, our technique can already be incorporated by practitioners into their designs. By placing our discoveries in the larger context set out by previous work on touch-screen interaction design for children [2,3,6,9], we can already anticipate multiple application opportunities made possible by our classification technique:

(1) *Adaptive widget layouts.* Since we know from the literature that children miss on-screen touch targets more often than adults [3], are less precise [9], and have more difficulty with smaller targets [3], children would benefit from an interface layout with larger widgets and more space in-between. Our technique makes it possible to detect that the user is a child and, therefore, the application can reconfigure the interface layout *on the fly*.

(2) *Intelligent widget activation.* When changing the widget layout is not desirable, our technique makes it possible to infer the user's intended target for near-misses by applying a distance offset filter less stringent for children than for adults, *e.g.*, the app may accept farther-away touches and still activate the target if it knows the user is a child. The layout and size of the widgets stay the same, but the application will apply different age-group-dependent processing. This mechanism will provide more flexibility for the application developer by allowing them to use the same interface for multiple target age groups. The increase in touch target acquisition flexibility while maintaining layout appearance is a desirable design option when children and adults use the interface collaboratively.

(3) *Adaptive activation of gesture processing techniques.* Knowing whether the user is a child or an adult can be used to call specialized event handlers in the touch

processing chain of the application. For instance, we know that gesture recognizers perform worse on children's gestures [3], so the application could benefit from the age group knowledge to know *when* to activate one recognizer or another [4,10]. Also, we know that visual feedback during gesture input is more important for children than for adults [2], and the app could use the age group data to increase or dim down visual feedback as necessary, which will provide more cues for children and more screen space for adults to display other app widgets and data.

(4) *Adaptive filtering of content and access to applications.* Knowing whether the user is a child or an adult can be used to automatically filter Internet content and to disable access to applications according to parental monitoring settings [1]. However, in contrast to existing approaches, our technique can inform the app *right away* (from the first touch or just a few touches) that the ownership of the device has changed and can enforce or disable parental settings automatically.

All these application opportunities can be easily incorporated into existing interfaces, increasing adaptiveness to users. We hope that this first examination of age group classification with basic, readily collectable touch data will empower the community with a new tool to make touch interfaces even more adaptable for users, and will inspire researchers to gain richer understanding of touch input for different age groups.

References

1. Álvarez, M., Torres, A., Rodríguez, E., Padilla, S., Rodrigo, M.J. 2013. Attitudes and parenting dimensions in parents' regulation of Internet use by primary and secondary school children. *Computers & Education 67*, 69-78
2. Anthony, L., Brown, Q., Nias, J., and Tate, B. 2013. Examining the need for visual feedback during gesture interaction on mobile touchscreen devices for kids. In *Proc. of IDC'2013.* ACM, New York, NY, USA, 157-164
3. Anthony, L., Brown, Q., Tate, B., Nias, J., Brewer, R., and Irwin, G. 2014. Designing smarter touch-based interfaces for educational contexts. *Personal and Ubiquitous Computing, 18*(6), 1471-1483
4. Blagojevic, R., Chang, S.H.H., Plimmer, B. 2010. The power of automatic feature selection: Rubine on steroids. In *Proc. of SBIM '10*, 79-86
5. Kang, C. 2013. Survey: For young children, mobile devices such as tablets, smart phones now a mainstay. The Washington Post
6. 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 *Proc. of SBIM '13*. ACM, New York, NY, USA, 33-42
7. Rainie, L., Smith, A. 2013. Tablet and e-reader ownership update. Pew Research Center
8. Smith, A. 2013. Smartphone ownership 2013. Pew Research Center
9. Vatavu, R.D., Cramariuc, G., Schipor, D.M. 2015. Touch interaction for children aged 3 to 6 years: Experimental findings and relationship to motor skills. *International Journal of Human-Computer Studies, 74*, 54-76
10. Vatavu, R.D., Anthony, L., Wobbrock, J.O. 2012. Gestures as point clouds: A $P recognizer for user interface prototypes. In *Proc. of ICMI '12*. ACM, NY, USA, 273-280
11. Webb, A. 2003. Statistical Pattern Recognition, 2nd Ed. John Wiley & Sons, UK.