

***PhonePoint Pen:***  
**Using Mobile Phones to Write in Air**

**Sandip Agrawal**

**Advisor: Romit Roy Choudhury**



SUBMITTED FOR GRADUATION  
WITH DEPARTMENTAL DISTINCTION IN  
ELECTRICAL AND COMPUTER ENGINEERING

(SUBMITTED ON APRIL 14, 2010)

## Acknowledgement

I would like to take this opportunity to express my gratitude to my advisor, Prof. Romit Roy Choudhury, for all his advice and guidance throughout the project. This work would not have been possible without his support and encouragement. Special thanks to my colleagues, Ionut Constandache and Shravan Gaonkar, for their immense help in completing the project. Furthermore, I would also like to thank everyone in SyNRG group, including Souvik Sen, Naveen Shantapuri, Justin Manweiler, Martin Azizyan, and Xuan Bao, for making this a memorable experience. I am also grateful to Dean Absher for providing me the opportunity to get the research experience through the Pratt Fellow Program. Finally, I would like to dedicate this thesis to my parents. I am utterly grateful for all their love, support and encouragement throughout. I love you, dad and mom!

## Abstract

Numerous sensors in modern mobile phones enable a range of people-centric applications. This paper envisions a system called *PhonePoint Pen* that uses the in-built accelerometer in mobile phones to recognize human writing. By holding the phone like a pen, a user should be able to write short messages or draw simple diagrams in the air. The acceleration due to hand gestures can be translated into geometric strokes, and recognized as characters. We prototype the *PhonePoint Pen* on the Nokia N95 platform, and evaluate it through real users. Results show that English characters can be identified with an average accuracy of 91.9%, if the users conform to a few reasonable constraints. Future work is focussed on refining the prototype, with the goal of offering a new user experience that complements keyboards and touch-screens.

## 1 Introduction

Imagine the following scenario. While driving to office, Bob stops at a traffic light. As he mentally sifts through his tasks for the day, he remembers that he needs to call his friend, Alice, very soon. Since Bob tends to forget his personal commitments, he decides to make a note of this task. Therefore, while keeping his gaze on the traffic lights, he draws out the phone from his pocket, and by holding it like a pen, he writes “ALICE” in the air. He also gestures a check-mark to email the written note to himself. He does not look at any of these hand-gestures he makes. Once in his office, he finds an email in his mailbox that reads “PhonePoint Pen – ALICE”. Bob calls Alice, talks to her, and deletes the email. The figure below shows the output of air-writing ALICE using the PhonePoint Pen.



The above is a fictional scenario, however, representative of broad possibilities in the area of human-sensor interaction. A particular possibility pertains to a sensor-assisted input technology that can easily “note down” short pieces of information. Although existing technologies have made valuable advances to meet these needs, the quality of user-experience could still improve. We discuss some avenues of improvement, and motivate the potential of PhonePoint Pens.

Typing an SMS, while popular among the youth, has been unpopular among a moderate section of society. Studies report user dissatisfaction with mobile phone typing [17, 6, 5]. The major sources of discomfort arise from small key sizes, short inter-key spacings, and the need for multi-tapping in some phone keyboards. With increasingly smaller phones, keyboard sizes may decrease, exacerbating the difficulty in physical typing. The difficulties are likely to be pronounced for elderly people and motor impaired patients.

Even if future keyboards [15] improve the typing experience, some problems may still persist. While walking, or with one hand occupied, typing in information may be inconvenient. Using the mobile phone accelerometer to capture hand gestures, and carefully laying them out in text or image, can improve the user experience. The ability to write without having to look at the phone keypad may offer an added advantage.

One may argue that voice recorder applications on mobile phones may be an easy way to input short pieces of information. However, searching and editing voice-recorded content is difficult (unless processed through a separate speech-to-text software). Further, playing aloud the voice messages can be inconvenient and time-consuming. Writing in air, and converting them to typed text, may alleviate these problems.

Current approaches are largely ad hoc. People use whatever is quickly reachable, including pen-and-paper, sticky notes, one's own palm, etc. None of these scale because they are not always handy, and more importantly, not always connected to the Internet. Thus, hastily noted information gets scattered, making information organization and retrieval hard.

This paper proposes to use the in-built accelerometer in modern mobile phones as an easy and ubiquitous way of capturing (short) written information. The problem definition bears similarity to known problems in gesture recognition. However, as we will see later, recognizing actual alphabets in air (using the phone processor, a noisy accelerometer, and no software training), raises a number of new challenges. For instance, as a part of writing the alphabet "A" on paper, one must write "/" first, lift and reposition the pen on the paper, and then write the "—". When writing in air, the phone cannot easily say which part of the hand-movement is intended to be the "re-positioning" of the pen. The problem is further complicated by the inherent noise in mobile phone accelerometers, the user's involuntary wrist-rotation, and practical difficulties in deriving displacement from noisy acceleration, etc. The PhonePoint Pen addresses majority of these challenges by treating the accelerometer readings as a digital signal, and successively refining it through simple numerical and signal processing algorithms. Simplicity is important to ensure that the operations can be performed on the phone processor. Once individual geometric movements have been tracked, their sequence of occurrence is matched against a decision tree (a simple grammar). The matching operation yields the English alphabet.

Our current prototype requires the user to get used to a few soft constraints. Mainly, we advised users to pretend that they are writing on an imaginary blackboard – this reduced the user's wrist and elbow rotation while gesturing a stroke. Users were also requested to write 12-inch sized capital letters at a moderate speed (not too fast), to ensure that the sampling frequency of the accelerometer was adequate to capture the motions. Users who performed the tests after 3 to 4 minutes of rehearsing, achieved an average accuracy of 91.9% with English alphabets. The geometric representation of the characters (shown as 2D images of the actual hand-writing) were legible in around 83% of the cases. Survey responses from randomly picked student users, as well as from real patients and doctors in the university hospital, were positive. The absence of visual feedback while writing did

not appear to be a concern at all, while the energy consumption from air-writing proved to be marginal. While further research is necessary to attain a commercial-grade input system, we believe that our prototype presents promise of viability in the near future. Our contributions are as follows.

- **We explore the viability of using the mobile phone accelerometer to write in the air.** While a number of gesture recognition schemes already exist [7, 8, 3, 10], the ability to write English alphabets (and draw simple diagrams) present distinct challenges.
- **We characterize the nature of the challenges and propose a multi-phase approach to recognize alphabets and words.** Our algorithms are deliberately simple for on-phone real-time operation. We also develop a customized spell checker that corrects motion-related errors, such as between  $D$  and  $P$ .
- **We implement PhonePoint Pen (P3) on Nokia N95s, and test it with 10 student users and 5 hospital patients.** All but one student user was able to write with high accuracy, while their geometric outputs were quite legible. Accuracy with real patients was poor (mostly due to usability problems), however, the feedback from medical practitioners were unanimously encouraging.

The PhonePoint Pen demo video, and other related information, is available at:

<http://synrg.ee.duke.edu/media.htm>

The rest of the paper is organized as follows. Section 2 discusses the potential use-cases for the PhonePoint Pen. The core design challenges are discussed in Section 3, followed by the system design and algorithms in Section 4. The implementation and evaluation is presented in Section 5. Section 6 discusses some of the remaining research challenges, and the related work is visited in Section 7. Finally, the paper closes with a summary in Section 8.


## 2 Use Cases

We begin with a discussion of possible use cases for P3. These are not necessarily to express the utility of current system; instead they are a vision of the future. However, where applicable, we show examples from our current system.

*Assistive Communications for Impaired Patients:* The Speech Pathology and Surgery division of our University Medical School expressed keen interest in using the PhonePoint Pen as an assistive technology for impaired patients. Several patients suffer from inherent speech impairments, or experience similar conditions after surgeries. War veterans often loose fingers or limbs, while others lack finger-dexterity for typing on keypads. Yet, these patients are often capable of broad (one-handed) gestures, such as in sign-languages. P3

can be of assistance to such patients. It can permit some level of impromptu communication between a speech/hearing-impaired patient and someone who does not understand sign-languages. We have performed 15-minute experiments with 5 real patients at the Hospital, and discussed the applicability of the system with surgeons and healthcare professionals. We discuss the experiences and outcomes in Section 5.4.

*Equations and Sketching:* One of the P3 test users suggested the possibility of quickly writing equations in the air. Equations are difficult to write with regular phone keyboards, and P3 may be convenient. Other use-cases involve taking mental notes or sketching simple diagrams. One may sketch driving directions, or draw out a desired food item (e.g., fish) in a foreign country’s restaurant (figures below show a  $\sqrt{2}$  and a fish).



*Emergency Operations and First Responders:* Emergency scenarios are often unsuitable for typing, or even talking on the phone, because the observer may be engaged in looking at the events around her. P3 allows for taking notes (with one hand) without requiring the user’s visual attention. The ability to look around and gesture at the same time may be useful in these situations.

While the above use-cases are specific to phone-based systems, one may envision more creative applications of air-writing (e.g., writing “CNN” in the air may switch the TV to the CNN news channel). We aim to design PhonePoint Pen with the aim of enabling such people-centric applications.

### 3 Core Challenges

Existing systems, such as the Wii remote [13] and others [21, 22, 8, 3, 10], have the ability to identify hand gestures with good accuracy. These gestures have been utilized to recognize a few numeric digits [21, 22]. Moreover, several of these systems are more resourceful in sensor hardware, including gyroscopes, webcams, and more CPU power [2, 12, 4, 16]. Writing English alphabets/words in real-time with commodity phones has been an unexplored problem. To this end, we discuss core research challenges in P3, and propose our approaches alongside each of them. We then assemble these building blocks into a functional prototype.

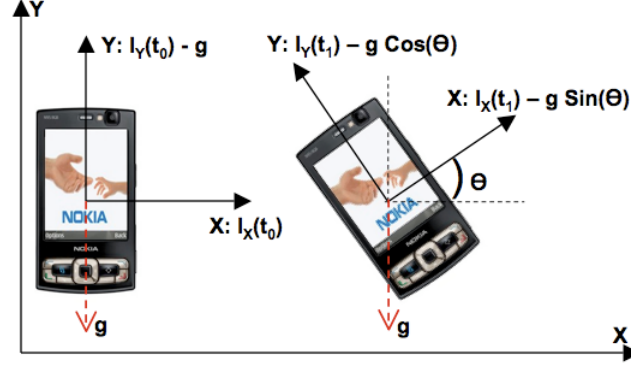


Figure 1: Earth's gravity projected on the XY axes; the axes are a function of the phone's orientation.

## (1) Filtering Rotation without Gyroscope

**Issue:** Nokia N95 phones are equipped with a 3-axis accelerometer that detects acceleration in the X, Y, and Z directions. Figure 2(b) shows an example of raw accelerometer readings on each of the 3 axes. The accelerometers measure linear movement along each axis, but cannot detect rotation. Hence, if the human grip rotates while writing, the reference frame of acceleration gets affected. Existing devices like “Wii Motion Plus” and Airmouse employ a gyroscope to discriminate rotation [14, 12]. In the absence of gyroscopes in phones, compensating for hand rotation is a problem.

**Proposed Approach:** We begin with a brief functional explanation of the gyroscope. Consider the position of a gyroscope-enabled phone (GEP) at time  $t = t_0$  in 2D space (shown in the left side of Figure 1). At this initial position, the figure shows that the GEP's axes are aligned with the earth's reference axes (i.e., gravity is exactly in the negative Y direction). The accelerometer reading at this position is  $\langle I_x(t_0), I_y(t_0) - g \rangle$ , where  $I_x(t_0)$  and  $I_y(t_0)$  are the instantaneous accelerations along the  $x$  and  $y$  axes at time  $t_0$  respectively, and  $g$  is gravity. Now, the phone may rotate at the same physical position at time  $t_1$  also shown in Figure 1 (right). The phone now makes an angle  $\theta$  with the earth's reference frame, and the accelerometer readings are  $\langle I_x(t_1) - g \sin(\theta), I_y(t_1) - g \cos(\theta) \rangle$ . However, it is possible that the phone moved along the XY plane in a manner that induced the same acceleration as caused by the rotation. This leads to an ambiguity that gyroscopes and accelerometers can together resolve (using angular velocity detection in gyroscopes). However, based on the accelerometer readings alone, linear movements and rotation cannot be easily discriminated.

This is a difficult problem, and we cope with this by imposing a soft constraint on the user. (i) The simpler approach is to pretend that one of the corners of the phone is the pen tip, and to hold it in a non-rotating grip (shown in Figure 2(a)). Some users also found it easier to hold it like a white-board eraser – this grip also reduced wrist-rotation. (ii) Alternatively, while writing an alphabet, users may briefly pause between two




Figure 2: (a) Pretending the phone’s corner to be the pen-tip reduces rotation. (b) Raw accelerometer data while drawing a rectangle (note gravity on the Z axis). (c) Moving average computation.




Figure 3: (a) Final processed acceleration readings (b)Computing velocity as an intermediate step towards measuring displacement. (c) The approximate rectangle as the final output.

“strokes”. The pause is often natural because the user changes the direction of movement (from one stroke to another). For example, while writing an “A”, the pause after stroke “/” and before the starting of stroke “\” can be exploited. An accelerometer snapshot at this paused instant can identify the components of gravity on each axes, and hence, the angular orientation  $\theta$  can be determined. Knowing  $\theta$ , the phone’s subsequent movement can be derived. Of course, we assume that the phone rotates only in-between two strokes, and not within any given stroke. The assumption gets violated in some cases, reducing P3’s character recognition accuracy.

## (2) Suppressing Background Vibration

**Issue:** Accelerometers are sensitive to small vibrations. Figure 2(b) reports acceleration readings as the user draws a rectangle using 4 strokes (around 350 units on the Z-axis is due to earth’s gravity). A significant amount of jitter is caused by natural hand vibrations. Furthermore, the accelerometer itself has measurement errors. It is necessary to suppress this background vibration (noise), in order to extract jitter-free pen gestures.



**Proposed Approach:** To cope with vibrational noise, we apply two noise-reduction steps (the acceleration is treated as a discrete sequence of signal samples). First, we smooth the accelerometer readings by applying a moving average over the last  $n$  readings (in our current prototype,  $n=7$ ). The results are presented in Figure 2(c). Next, we label all acceleration samples less than  $0.5m/s^2$  as noise. We chose this threshold based on the average vibration caused when the phone was held stationary. All noise-labeled samples are suppressed (i.e., set to 0). Figure 3(a) shows the combined effect of smoothing and noise suppression.

### (3) Computing Displacement of Phone

**Issue:** The phone’s physical displacement is necessary to estimate the size of the air-written character, as well as their relative positions (such as in equations, figures, etc.). The displacement  $\delta$  is essentially computed as  $\delta = \int (\int a dt) dt$ , where  $a$  is the instantaneous acceleration. In other words, the algorithm first computes the velocity (the integration of acceleration), followed by the displacement (the integration of velocity). Noise in the acceleration readings will reflect on the velocity computation, and will get magnified in the computation of displacement. For instance, an erroneous short positive impulse in the accelerometer (i.e., acceleration becoming positive and then returning to zero), results in a positive velocity. Unless an identical negative impulse compensates for the positive impulse, the phone would appear to be in a state of continuous velocity. When this velocity is integrated, the displacement error will be large.

**Proposed Approach:** In order to reduce the velocity-drift errors, we look at consecutive accelerometer readings labelled as noise in the previous step. We reset the velocity to zero, if  $n$  consecutive readings have been suppressed as vibrational noise. This is because a continuous sequence of noise vibration is a good indicator of a pause, or a statically held phone; hence, it is an opportunity to suppress inertial error. Figure 3(b) shows the effect of resetting the velocity. Even if small velocity drifts are still present, they have a tolerable impact on the displacement of the phone. As seen in Figure 3(c) the amount of displacement and the shape drawn are represented reasonably well. The direction of movement is inferred from the signs of the acceleration along the X, Y, and Z axes.

### (4) Differentiating an “A” from a Triangle

**Issue:** The imaginary blackboard in the air has no global reference frame for position. While writing the character “A”, assume the writer has already drawn the “/” and “\”, and now lifts the pen to draw the “-”. Observe that the phone has no idea about the global position of “/\”. Hence, upon drawing the “-”, the pen does not know whether it is meant to be added in the center (to indicate an “A”), or at the bottom (to indicate a triangle,  $\Delta$ ). This ambiguity underlies several other characters and shapes.

**Proposed Approach:** This is a difficult problem, and we plan to jointly exploit the accelerations along the X, Y, and Z axes. While writing an “A”, assume that the user has

just finished writing “/”. The pen is now at the bottom of the “\”. The user will now lift the pen and move it towards the up-left direction, so that she can write the “-”. The lifting of the pen happens in 3D space, and generates an identifiable impulse on the  $Z$  axis. When the acceleration in  $Z$  axis is above a certain threshold, we label that stroke as a “lifting of the pen”. This pen-lifting can be used as a trigger for the user going off the record. User movements in the  $XY$  plane are still monitored for pen repositioning, but are not included in the final output. When the phone is in position to write “-”, a small pause and change in direction is an indication for going back on the record.

## (5) Identifying Character Transitions

Even if pen-lifts are recognized, certain ambiguities remain. For instance, “B” and “13” may have the exact same hand-movement, including the pen-lift. The user’s intention is difficult to recognize, making character distinction hard.

**Proposed Approach:** We rely on a combination of multiple heuristics to mark character separations. The simplest approach is to require the user to include a special gesture between characters, like a “dot” or a relatively longer pause. Thus, “13” should be written as |.  $\supset\supset$ , while the “B” should be |  $\supset\supset$ . While this may be inconvenient, the phone pen also employs additional methods of delimiting characters. These methods rely on understanding what the user has written till now, and what the next “stroke” is likely to be. We will discuss this in detail in the next section, after we have discussed stroke-detection and a simple stroke-grammar to identify characters.

## 4 System Design and Algorithms

The above building blocks provide for a geometric representation of air-written characters. While the geometric version can be displayed or emailed as an image, conversion to text is likely to be more useful (for browsing and searching). This section develops the algorithmic components towards this goal.

### 4.1 Stroke Detection

Characters can be viewed as a sequence of strokes. The alphabet “A”, for instance, is composed of 3 strokes, namely “/”, “\”, and “—”. If the discrete strokes can be pulled out from the seemingly continuous movement of the hand, it is possible to infer the characters. To this end, we have analyzed the English alphabets and constructed a basic set of strokes, as in Figure 4.

To identify the strokes, P3 computes a running variance of the accelerometer readings. When this variance falls below a threshold, P3 marks those regions as a pause of the hand. The pauses demarcate the human-strokes, allowing P3 to operate on each of them individually. For determining the exact stroke, our basic idea is to correlate the human-strokes




Figure 4: Basic strokes for English characters.

against each of the ideal strokes. This form of correlation is not new, and has been used as standard primitives in classification and matching techniques [10, 7, 1]. We perform correlation over a varying window size of accelerometer readings. This is because the hand often rotates towards the end of the stroke, and the samples corresponding to the rotations should ideally be pruned out. Correlation is able to cope with such issues, showing a high correlation value when the ideal stroke *broadly aligns* with actual readings. Besides, even if some inter-stroke pauses are not identified, varying the correlation window-size yields the stroke boundaries. The intuition is that two consecutive strokes are typically different in the English alphabet, and thereby, correlating across the boundaries of the strokes (with a large window size) reduces the correlation value. Performance results indicate a reasonable reliability in stroke detection. The natural question, then, pertains to combining the strokes into a character.

## 4.2 Character Recognition

The PhonePoint Pen observes the logical juxtaposition of strokes to deduce the character that the human is trying to write. For this, we adopt a stroke grammar for English alphabets and digits. Figure 5 shows a pruned down version of this grammar for visual clarity. The grammar is essentially a tree, and expresses the valid sequence of strokes to form an alphabet. Moreover, the grammar also helps in stroke-recognition because it provides P3 with an ability to anticipate the next stroke. For instance, observing strokes “| \ /” in succession, P3 can anticipate an “M” and expect the next stroke to be a “|”. Thus, by correlating “|” to the stream of accelerometer readings (and ensuring a high correlation), the system can better identify the end-points of the next stroke. This helps in identifying the residual samples, which in turn helps in tracking the re-positioning of the hand in-between strokes. The benefits are cascading.

In certain cases, the user’s hand movement may be falsely classified as an incorrect stroke. A frequent example is “\” and “⌢”. Since the user’s hand has a natural rotational motion (pivoted at the elbow), moving diagonally for a “\” results in an arc, which then gets classified as “⌢”. Thus “N” may not be recognizable due to misclassification of the second stroke. To account for such possibilities, we have updated the grammar tree. For example, if | is followed by ⌢, we call it a “D” or “P”; however, if this is again followed by a “|”, we infer an “N” (since no alphabet is a sequence of “| ⌢ |”). We observe that such

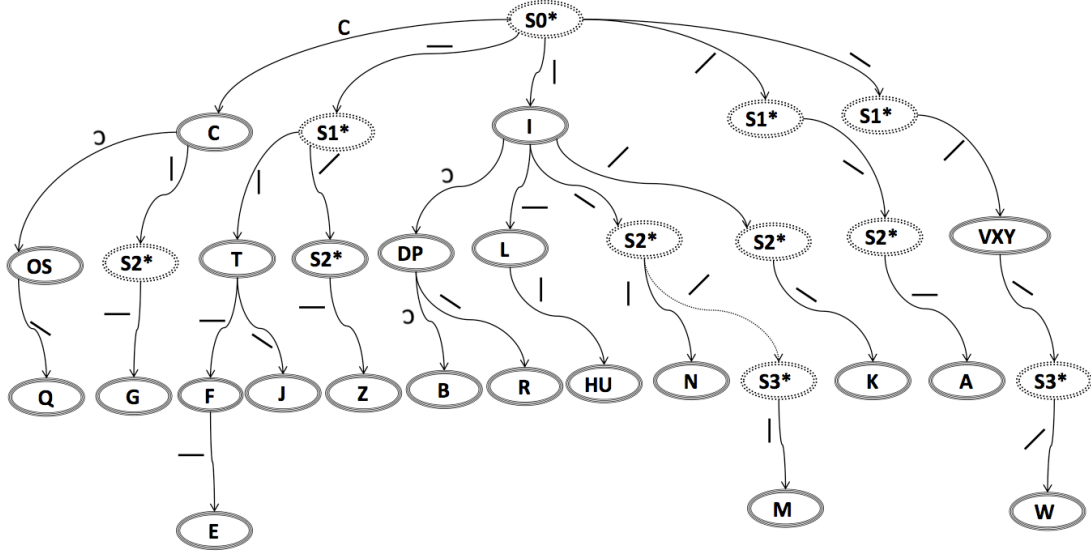


Figure 5: The basic grammar for character recognition. The edges are labeled with a stroke. Reaching certain states (solid circles) implies a valid alphabet, while others are intermediate states (dotted circles). Certain states have multiple alphabets in them, suggesting ambiguities in the basic stroke grammar.

opportunities are numerous in the stroke grammar, adding to the robustness of the system. We do not include this updated grammar in the paper and only show the example for N in Figure 6.

**Grammar Ambiguity.** Interestingly, the stroke grammar presents a number of ambiguities. For instance, “O” and “S” are composed of the same strokes, namely, “C” and “o”. P3 resolves this by simply observing the direction of movement in the second stroke. If the hand is moving upwards, computed from the sign of the Y-axis acceleration, the alphabet is declared as an “O”, and the vice versa. Another ambiguity is between “D” and “P”. In this case, P3 computes the relative sizes of Y-axis movements and compares them. If the sizes are comparable (second stroke greater than 0.75 of the first), the alphabet is deemed as a “D”, else a “P”. Finally, some kind of ambiguities are relatively harder. For instance, “X” and “Y” have the same strokes, and only differ in how the user repositions her pen. Since hand-repositioning has no preset movement, they are more prone to error. Thus, even though the “\” in “Y” is smaller than that of “X”, P3 makes a mistake in some cases. Finally, “O” and “0” cannot be discriminated.

### 4.3 Word Recognition

Recognizing the juxtaposition of characters, to recognize a word, adds to the ambiguity. For instance, “B” and “13” are identical in terms of the strokes used, and so are “H” and “IT”. Unless we find a signature to demarcate characters, the PhonePoint Pen will yield false positives. Towards this goal, we consider a combination of multiple heuristics. None of these heuristics are adequate in isolation, but may be reasonable when used in con-




Figure 6: Incorporating tolerance into the grammar tree translating it into a graph. The dotted edges are incorrect but permissible, e.g., alphabet “N” can be reached via multiple paths such as  $| \supset |$ .

junction. First, we make the observation that while transitioning from one alphabet to the next, users have a naturally longer pause (especially with upper case alphabets). Second, we observe that in some cases, if the hand moves in a *leftward* direction, it may be a hint about the start of a new character. This happens, for example, when a user has written across the imaginary plane in front of her, and moves back in space (towards left) to write more. Typically, since most of English alphabet strokes are gestured left to right, any opposite movement (from right to left), is a useful hint for character segregation. Third, we ask the users to gesture a “dot” between characters whenever they can remember. Thus, “13” should be written as “|.  $\supset \supset$ ”, while the “B” should be “|  $\supset \supset$ ”. Drawing a dot presents a unique signature to delimit characters, but slows down the user while writing. Thus, the user can use it only if she remembers or wishes to. If the delimiter is not used, the recognition accuracy is affected.

We note that not all cases are like “B” and “13”. Even without the delimiter, the stroke grammar will naturally separate some characters. In other words, given a sequence of strokes, P3 anticipates the next stroke to be from a specific subset of strokes. If the next stroke is not in this anticipated subset, then it implies the start of a new character. For example, given “|” and “—”, the phone can anticipate a “—” assuming that the user is trying to write an “F”, or “E”. However, if the next stroke is “C”, then the phone immediately infers that the prior alphabet was an intended “T”. Even if the delimiter is not present, such character transitions can be recognized to form words. Spell checkers can be employed on top of these methods to further improve word recognition accuracy.

## 4.4 P3-Aware Spelling Correction

Spelling correction tools accept a given word and compute a list of valid English words, sorted in the order of “edit distance”. The edit distance between two strings of characters is defined as the number of operations required to transform one of them into the other. The corrected spelling is typically the valid word with minimum edit distance. Importantly, multiple words may have the same (minimum) edit distance, and even the minimum edit distance may not be the best when the nature of the errors are guided by certain distributions. For example, the word MQM has an edit distance of 1 with valid words MOM, MAM, MUM. Since we can learn that P3 often confuses Q with O (the nature of the strokes are similar), but hardly confuses Q with A or U, a P3-aware spelling correction tool can suggest MOM with high confidence. A less trivial example occurs when P3 outputs, say, NIET. Words NET and MET have edit distances of 1 and 2, respectively. However, the spelling suggester could observe that P3 confuses “M” as “NI” with far greater probability than “E” as “IE”. Thus, one could predict that the user intended to write MET with reasonably high probability, even though its edit distance is higher. Formally, given a mis-spelt word  $w$ , the P3-aware spelling corrector computes the word  $\phi_w$  as follows.

$$\phi_w = \{i : \frac{P(w|i)}{P(w|j)} > 1\} \quad \forall \text{ valid words } i, j, i \neq j$$

The distribution of  $P(w|i)$  is learnt from our own data set, and can adapt to the user’s idiosyncrasies over time. We have implemented this tool and found improvement over dictionary based spelling correction.

## 4.5 Control Gestures

To write a short phrase, the words need to be separated by spaces. In certain cases, the characters may need to be deleted. Further, the user should be able to email the written/drawn content to her email address. These are a few control operations that are vital to improve the user’s experience. The PhonePoint Pen assigns a unique gesture to each of these, and recognizes them without difficulty. Specifically, the space is denoted by a long horizontal movement or two dots. The deletion is like using an eraser – the users shakes her hand at least four times briskly. To email, the user must draw a check mark in the air. With these functionalities in place, we present the implementation details of P3, followed by performance evaluation.

# 5 Implementation and Evaluation

We prototyped the PhonePoint Pen on a Nokia N95 phone. The 3D accelerometer was programmed to obtain 30-35 acceleration readings per second. We developed a server side implementation first on MATLAB. Basic MATLAB libraries allowed us to implement some signal processing techniques (low pass filtering) and simple statistical analysis. We ported this code on Python for on-phone processing, thus users write in air and the output is

shown on the screen. Some of the techniques were simplified (filtering operations modified to running averages and subtractions). The results from Python and MATLAB differed rarely.

The remainder of this section is organized in three parts: (1) evaluation metrics and methodology, (2) PhonePoint Pen evaluation with students, and (3) experiences from patients with cognitive/motor impairments conducted at the University Hospital.

## 5.1 Evaluation Metrics

The P3 evaluation is centered around character and word recognition accuracies. We define *Character Recognition Accuracy (CRA)* as the fraction of successful typed text recognitions, when a user writes individual alphabets/characters (used interchangeably). In addition to CRA, we also evaluate P3’s quality of geometric representation. For this, we display the geometric characters to a human, and ask her to recognize them. The correctly identified fraction is defined as the *Human Readability Accuracy (HRA)*.

To compute *Word Recognition Accuracy (WRA)*, we randomly generated English words from a dictionary and requested test users to write them in air. Longer words are naturally more prone to mistakes because every character and every transition will have to be precisely decoded. Thus WRA degrades with word-length. Nevertheless, since P3 outputs typed-text we can apply spelling correction to improve the final accuracy. Thus, we report WRA for basic P3, WRA with *English-Spelling-Correction*, and WRA with *P3-Aware-Correction*. We also report WRA with Human Readability (i.e., fraction of words correctly recognized by humans).

## 5.2 Evaluation Methodology

We conducted PhonePoint Pen tests mainly with students from computer science, and engineering. The test group comprised of 10 students in two categories: Trained and Novice. Novice students (6/10) were defined as users that practiced less than 10 characters before starting the evaluation. The rest were Trained students who practiced 26 characters each (each English alphabet approximately once, taking less than 5 minutes in total). Only one of the Trained users rehearsed for 75 characters before starting the tests. Besides university students, we also performed a study with a small population of 5 patients from the University Hospital – the primary purpose was to gain insights into P3’s applicability into assistive technology. According to our IRB approval, the patients were allowed to write up to 8 characters. The patients had no previous experience with our prototype, and performed the experiments under the supervision of care-givers. Although P3 broadly failed in the tests (due to occurrences that the system was not designed for), we will report the valuable experience and feedback we gained from neurosurgeons, physicians, and speech pathologists. We will discuss the modifications we have made to P3 based on these real-life feedback.




Figure 7: Alphabets D, U, K, E, U, N, I, V, E, R, S, I, T, Y, as outputs of the PhonePoint Pen. Although distorted, the characters are legible.




Figure 8: The raw acceleration data for the Alphabets E and Y from above.

### 5.3 Performance Evaluation

The main evaluation results are summarized as follows.

- Figs. 7 and 9 show sample words written with P3. Fig. 10 quantifies this by showing that average readability is 83% and 85.4% for characters written by Trained and Novice writers respectively. Employing stroke grammar, average character recognition is 91.9% and 78.2% for the same two groups (Fig. 11). Numeric digits experience similar accuracy.
- Zooming into results, most Novice users achieve similar CRA to Trained users (Fig. 13) – 2 weakest users achieve comparable accuracy by writing on a table-top (Table 1). Disambiguation approaches are reasonably effective (Fig. 14). However, the average character writing time is between 3.02 to 4.3 seconds (Fig. 15), the main limitation with the current system. Energy consumption with P3 is not a concern (Fig. 16).
- Word recognition degrades with increasing word-length, however, spelling correction helps – 80% for 5 letter words (Table 2). Human readability is lower: 70% for 5 letter words. Tests with hospital patients show low accuracy due to usability issues (Table 3). Advised by speech surgeons, we emulated a class of patients by writing with the left hand – P3 achieves 81.7% accuracy (Fig. 17).



ACM ZURICH PEN PAPER  
JANE LOL CU FYI AM PM

Figure 9: Samples of air-writing (few SMS lingo) recognized correctly by stroke grammar: ACM, ZURICH, PEN, PAPER, JANE (a name), LOL (laugh out loud), CU (see you), FYI, AM, and PM.




Figure 10: Human Readability Accuracy (HRA) for Trained and Novice users. Human's powerful cognitive abilities result in comparable recognition performance, even though the Novice's characters were qualitatively observed to be more distorted.

## Results: Readability and Recognition

Figure 7 shows the geometric version of alphabets D, U, K, E, U, N, I, V, E, R, S, I, T, and Y, written by a trained user (each alphabet written separately). The acceleration readings for E and Y are presented alongside. P3 correctly converts the acceleration to alphabets, while the geometric versions are human legible. Figure 9 shows some examples of air-written words. Evidently, the lack of a reference frame degrades the sense of proportion and relative placement of characters. Despite these distortions, the stroke grammar yielded correct results for all the words in Figure 9.

Towards a systematic evaluation, Figure 10 shows the *Human Readability Accuracy (HRA)* and the *Character Recognition Accuracy (CRA)* per-alphabet, per-user-category (526 characters were written in total). For HRA, each of the 526 alphabets were presented to arbitrarily selected students. Average HRA for Trained and Novice categories proved to be 83% and 85.4% respectively. This is likely because human cognition is powerful and is able to decipher even highly distorted characters (the key intuition with Captchas [19]). Thus, even though Novice users exhibited greater distortion in the geometric alphabets, human

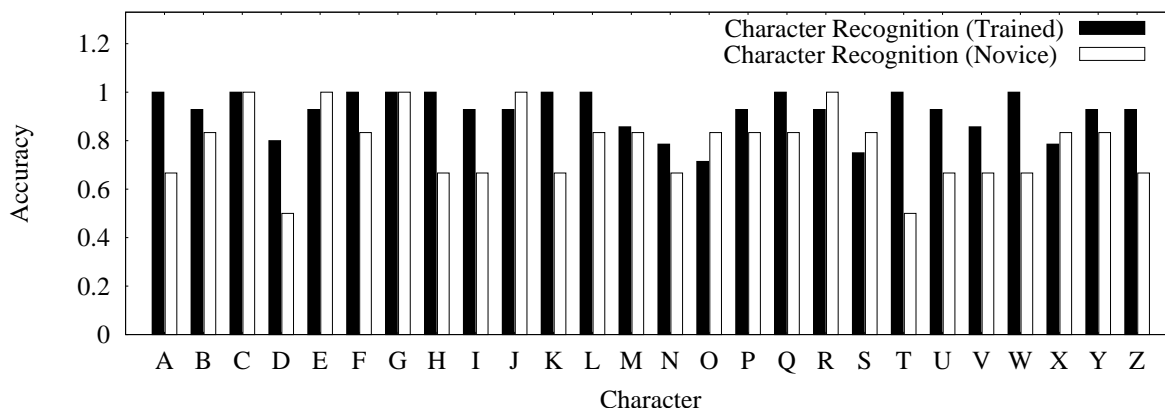


Figure 11: Character Recognition Accuracy (CRA) for Trained and Novice users per-character. The stroke grammar achieves a reasonably high accuracy, especially with Trained users.

readability for both the categories remained comparable. The expected difference between the two categories became evident in the CRA comparison. Figure 11 shows an average of 91.9% and 78.2% CRA for Trained and Novice categories, respectively. This suggests that 2-4 minutes of training has a positive impact – users learn how the system reacts to their hand-movements and adapt somewhat involuntarily. Numeric digits achieved comparable accuracies (shown in Figure 12).




Figure 12: Per-character CRA for numeric digits with 4 Trained users.

## Results: Per-User Accuracy

The accuracy of Novice category users was relatively lower – the following discussion zooms into the results. We re-plotted the results from the experiments on a per-user basis (Figure 13), and observed that the variance among the Novice users was quite high. Four novice users were able to achieve reasonably good CRA (in fact one of them was better than the Trained users), while two other novices weren't able to exceed a CRA of 70%. In response to this finding, we measured how these users performed when writing on a flat surface, like a table-top. Our hypothesis was that certain involuntary 3D hand motion, or intra-stroke wrist rotation, is likely to affect recognition – writing on a physical surface could improve performance.

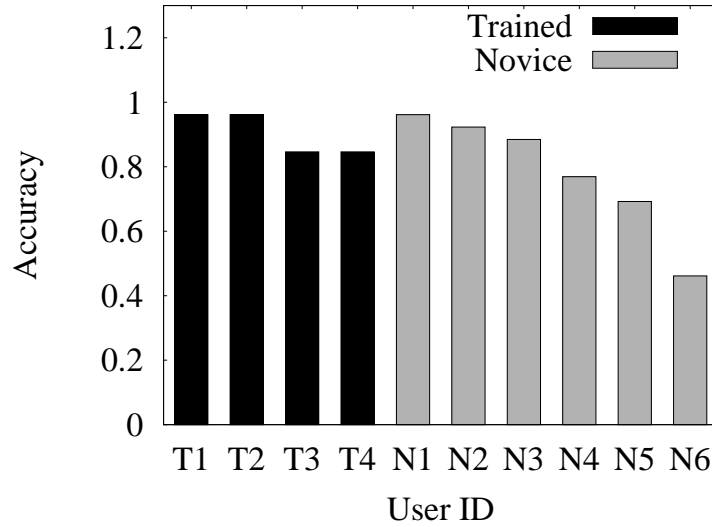


Figure 13: Average CRA per-user. Novice users have a wide variation in accuracy.

Table 1 shows the accuracy improvement when the two weak Novice users wrote the alphabets on a table-top. For instance, while writing “A”, the phone touched the table for the “/” strokes, was lifted and repositioned on the table again, and then the user wrote the “-”. The on-table accuracy improved substantially for the users; the weakest user experienced a jump from 46.2% to 76.9% CRA. Assuming that flat surfaces are often accessible, P3 may be acceptable even to the weak user.

Table 1: Writing In-Air Vs. On-Table

User	In-Air	On-Table
Weakest User	46%	77%
2nd Weakest User	69%	88%

## Results: Grammar Disambiguation

Recall that the P3 stroke-grammar exhibits inherent ambiguity. For example, D and P are written with the same strokes “|” and “⊃”. Similarly, characters {V, X, Y} and {O, 6, S} use common sets of strokes. We disambiguate between these characters by looking at the directions of the stroke (V has a downward followed by an upward movement, while X has a downward, pen reposition, and then again downward). We also track the movement of the phone during the pen-repositioning phase to extract more information about the user’s intention. Figure 14 presents the accuracy of disambiguation, along with the actual outcomes when the disambiguation fails. In majority of the cases, the character is decoded correctly. Also, among incorrectly decoded characters, some are not confused with their ambiguous counterparts (e.g., H and U). However, ambiguity still occurs, e.g., 0 with 6 in 28.6% and 6 with 0 in 35.7% cases; P with D in 7% cases.




Figure 14: Disambiguation performance with P3 (“Oth” in the striped portions implies “Others”).

## Results: Writing Speed

Figure 15 presents two CDF curves. One denotes the distribution of alphabet-writing time computed across all users. The median time was 4.3 seconds. We believe users displayed a tendency to write slower than necessary, partly because they were new to the system, and because they were keen on optimizing for accuracy. P3, however, can support quicker writing. To understand the speed limits with P3, we measured the minimum time incurred in writing each alphabet correctly. The second curve plots this distribution, and evidently, the median improves to 3.02 seconds. Nevertheless, even the best performance of P3 is quite slow, and is currently the key limitation with the prototype. We believe a number of opportunities exist that will increase the speed of air-writing. Greater number of accelerometer samples per second is an immediate one; we will discuss others as a part of our ongoing work.

## Results: Word Recognition

We requested users to also write English words in air (we did not include the two weakest users for these experiments). The words ranged from 2 to 5 characters and were chosen randomly from a dictionary. The users wrote 20 words for each word-length. Table 2 reports P3’s Word Recognition Accuracy (WRA). Spell checking with an English dictionary improves the accuracy; P3-aware corrections led to further improvements. This is because the P3-Aware spell checker better understands the relationship between user-intent and P3-output, and is able to make the necessary corrections. We also observed reasonable performance through human recognition.

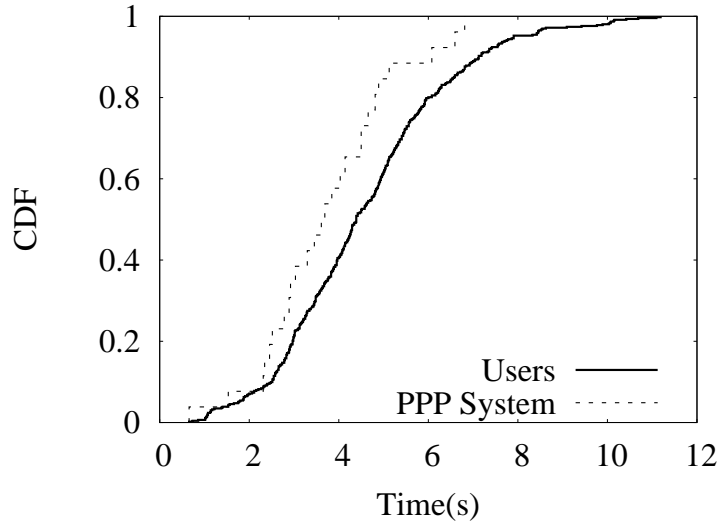


Figure 15: Distribution of time to correctly write English characters with P3.

Table 2: Word recognition.

Word Length	Phone Pen	Spell Correct	P3-Aware Spell Correct	Human Readable
2	17/20	19/20	20/20	11/20
3	18/20	19/20	19/20	12/20
4	13/20	18/20	19/20	10/20
5	13/20	16/20	17/20	14/20

## Results: Energy Measurements

We ran experiments to compute the energy footprint of the Nokia N95 accelerometer. We sampled the accelerometer at the same rate as PhonePoint Pen on a fully charged Nokia N95 8GB phone (Figure 16). The phone exhibited an average battery-lifetime of 40 hours, i.e., a user should be able to continuously write for 40 hours with P3.

## Results: Informal User Feedback

We performed informal survey and usability tests with around 20 students. In addition, the early prototypes of PhonePoint Pen featured in Slashdot and Scientific American received more than 100 comments. A video-demo on Youtube received more than 99,200 views. From these user surveys and online comments, we received valuable feedback. In general, people found P3 exciting; one user commented that *“even though I may not use it frequently, I would like to have this as an iPhone App”*. Some users were critical, and commented that they may look *“silly”* while writing in the air. A few users raised concerns about privacy, while others responded by saying that when observing a user air-writing, *“decoding the alphabets backwards is difficult”*. Many users liked the prospect of using the phone as the TV remote, and writing channel-numbers in the air. Finally, one comment was enthusiastic about being able to sign in air, perhaps when a FedEx package gets delivered to the door.




Figure 16: Accelerometer power consumption.

## 5.4 Experiences with Hospital Patients

In collaboration with Physicians from the Surgery/Speech Pathology department of University Hospital, we carried out PhonePoint Pen tests with patients suffering from various forms of cognitive disorders and motor impairments. Based on IRB approval, 5 patients were approved of writing 8 randomly chosen alphabets. The patients were selected with varying degree of motor-impairments (e.g., a hydrocephaleus lumbar drain trial patient exhibited good cognition but weak motor skills; a patient from a major car accident, 12 days prior, had a right side paralysis and spinal injury, but was able to write with his right hand; a 72-year old stroke patient had weakness on both limbs with sever tremors). The tests were carried under the supervision of medical practitioners and care-givers, who first learnt to use P3 from us. We were not allowed to observe the patients, however, we interacted closely with the care-givers to receive feedback. Importantly, P3 generated wide interest in the hospital, drawing neurosurgeons, speech therapists, and care-givers to witness the tests and comment on the potential applications and additional requirements. The overall experience proved to be invaluable. We report the main lessons from it, here.

(1) The P3 design requires users to press a button to start the application, and to re-press the button to stop. While this did not appear as an important design issue among university students, it proved to be a bad design choice for assistive technology applications. Patients found air-writing quite intuitive, but were unable to press the button appropriately. One patient pressed the buttons many times, another pressed the wrong one, and yet another found it hard to press. Table 3 shows the results. The unanimous advise from doctors was to eliminate button presses. Based on this feedback, we implemented the start and stop of the system through shaking. A user shakes the phone in air before beginning to write, and then shakes again to stop.

Table 3: Patient performance.

Patient ID	1	2	3	4	5
Accuracy	1/8	1/8	1/8	5/8	could not press button

(2) A neurosurgeon criticized that the P3 prototype required “shoulder, elbow, and wrist coordination”, a constraint that may be difficult to satisfy by hospitalized patients. His recommendation was to reduce the size of the letter so that it can be written with elbow movements alone. Moreover he suggested developing filters that would cancel the tremor in people’s hands, and thereby recognize the characters. Our ongoing work is focussed on learning the natural tremor of the person, and suitably canceling it from the entire accelerometer signal.

(3) One particular advantage of P3, even in light of specialized medical gadgets, is familiarity with cell phones. Physicians and care-givers emphasized the difficulties patients face in adopting new technological gadgets, particularly at the higher age group. Using the patient’s own phone to gesture “made a lot of sense”. They said, with a good degree of reliability, they envision a wide range of applications. Interestingly, several nurses showed enthusiasm at the prospect of the patient changing her TV channel by writing in the air (pressing the remote control button is again a difficult task for many).

(4) Exhaustive tests with patients are difficult due to IRB restrictions; the turn around time is also high. To partly overcome this, doctors suggested that it is valuable to test P3 by having normal right-handed users write with their left hand. Several speech-impaired patients only have minor problems with hand motions, and left-handed writing may be a credible emulation of these conditions. Figure 17 shows the CRA for consonants and vowels when 4 users wrote all the alphabets left-handed (we do not include per-alphabet accuracy in the interest of space). The average accuracy across all the alphabets and all users is 81.73%.

## 6 Limitations

Developing the PhonePoint Pen to the standards of a commercial product calls for further research. Nevertheless, we believe our current prototype has made credible advances towards the end goal. This section discusses the current limitations and opportunities for further improvement.

### Quicker Writing

The main limitation with P3 is its speed of writing – at best 3.02 seconds per alphabet on average<sup>1</sup>. Writing faster degrades accuracy. This is because the accelerometer exports around 30 samples per second, and is therefore inadequate to capture all the motions

<sup>1</sup>Of course, this should perhaps not be compared against typing on a keyboard or writing running hand.




Figure 17: Accuracy with left-handed operation emulating speech-impaired patients.

(especially in multi-stroke alphabets like E, H, W, etc.). While higher sampling rates will certainly be valuable [21, 22], we are investigating the benefits of using two accelerometers (available in OpenMoco phones). Further, we believe that there are opportunities in using the camera to determine hand-movements. By observing the changing camera-view over time (when the user gestures in air), the alphabets may be captured quicker.

### Writing long words and Drawing Pictures

Drawing capabilities require sophistication. The main problem stems from the difficulty in tracking the phone movement while the pen is being repositioned in 3D space. Thus, although the actual written words/shapes are identified, their relative placements are often incorrect. The problem is pronounced when the figure involves multiple pen-repositioning. Long words and sentences face similar problems. We leave these solutions to future work.

### Cursive Handwriting

Supporting cursive writing is certainly desirable with P3, however, significantly more difficult to accomplish. The problems of stroke-detection and character-recognition are exacerbated due to the continuous movement of the hand. One approach would be to apply pattern recognition algorithms on the entire stream of (noise-suppressed) accelerometer readings. However, such a scheme will not only require complex computation, but may also need a non-marginal degree of training. We have traded off this functionality for simplicity.

### Writing while Moving

If a person writes in the air while moving, the accelerometer readings will reflect that movement. Our current prototype assumes that the user is stationary while writing. Characterizing user motion, and subtracting them from the accelerometer signals, is a topic of future work.



### **Comment on Survey and Testing Population**

Students who have tested P3 are mostly students from the Computer Science and Engineering. These students are likely to have an understanding of accelerometers, and could have adapted to P3's behavior. In that sense, our accuracy results for Novice users could be partially optimistic. Nevertheless, with a little bit of training, even lay users should be able to adapt to P3.

## **7 Related Work**

Designing an alternative input technology is a rich area of research. Numerous sensors on mobile devices is fueling this area to rapid growth. Naturally, a large body of existing work relates to the PhonePoint Pen. We touch upon these works briefly (in the interest of space), while discriminating the contributions from this paper.

### **Air-gestures with 3D accelerometers.**

Gesture recognition has been widely studied through accelerometers, gyroscopes, vision based techniques, etc. [21, 22, 7, 8, 3]. Works that are closest to P3 include (1) a sensor mote-based 4-character recognizer [22], (2) a numeric digit recognizer with customized hardware [18, 21], and (3) *uWave*, a mobile phone based single-gesture recognizer. The first two works employ a highly capable accelerometer (around 100 samples/s). They use Principle Component Analysis (PCA), Hidden Markov Models (HMMs), and Dynamic Time Warping (DTW) algorithms, to achieve accuracies of 90 to 94%. However, the accuracy falls to 80% when the accelerometer is sampled at 40 samples/s. More importantly, the proposed systems are only able to write few numeric digits, *that do not require the user to reposition the pen within the same character*. Geometric figures are also not viable because gestures are identified through pattern matching, and hence, *the system does not compute the actual displacement and direction of motion*. PhonePoint Pen, on the other hand, tracks the user's hand movement, and develops methods for pen-repositioning, character transition, stroke-grammar, rotation avoidance, and character disambiguation.

*uWave* [10] is a mature work that allows a user to gesture with mobile phones, enabling simple operations like gesture-based user authentication, opening/closing applications, etc. The authors attain an impressive 99% accuracy with 8 gestures and negligible training. While this is valuable for a number of interfacing applications, we emphasize that character recognition entails an additional set of problems. Specifically, gestures are significantly tolerant to error; as long as the errors repeat across all gestures, the gesture can be identified. In contrast, the PhonePoint Pen requires a different approach to continuously track a more complicated motion of the hand.

### **Vision based gesture recognition.**

Cameras have also been used to track an object's 3D movements in the air [13]. Microsoft Research recently demonstrated a project titled "write in air" [2], that uses an apple in front of a camera to air-write alphabets. Computer vision based algorithms can precisely discern the movement of the apple (or any other object) to create both geometric and

textual representations of the alphabets. Noisy accelerometers and limited processing in mobile phones lack several advantages present in computer-connected cameras. Moreover, the system does not recognize words, side-stepping the problems of transition between characters. Signal processing based techniques are useful here, but not sufficient [1, 9, 20]

### **Wii, Logitech Air-Mouse, and Nokia NiiMe.**

A popular device capable of tracking hand movement is the Wii remote (or “Wiimote”) used by the Nintendo Wii console [13]. The Wiimote uses a 3-axes accelerometer to infer forward and backward rapid movements. In addition, optical sensors aid in positioning, accurate pointing, and rotation of the device relative to the ground. The optical sensor is embedded on the Wiimote and relies on a fixed reference (and a sensor bar) centered on top of the gameplay screen. The “Wiimote” can be augmented with the “Wii Motion Plus”, a pluggable device containing an integrated gyroscope to cope with hand rotation. These three sensors – the accelerometer, the gyroscope, and the optical sensor – can reproduce motions similar to real arm-motion. The Nokia N95 consists of only a (low-cost) accelerometer, and limited processing capabilities, in comparison to the Wii. Developing P3 on this Nokia platform presents unique research challenges.

The Logitech Air Mouse [12] targets people who use computers as multimedia devices. The Air Mouse provides mouse-like functionalities but the device can be held in air similar to a remote control. Accelerometers and gyroscopes together allow for linear and rotational motion of the pointer on the screen. Unlike the Air Mouse, P3 does not have a screen on which the human user may see and adjust the pen movement in real time. Thus P3 must estimate relative position of strokes and characters without any visual reference frame.

The NiiMe [4] project transformed the Nokia N95 phone into a bluetooth PC mouse. The PyAcceleREMOTER [16] project developed a remote control for the Linux media player MPlayer. By tilting the phone, the player’s play, stop, rewind, fast-forward, etc. are controlled. Other applications like *Inclinometer* provides car inclination while *Level Tool* allows measurement of the inclination of different surfaces by placing the phone on that surface. Lastly many video games for the N95 phone make use of the accelerometer, e.g., to guide a ball through a maze. Being able to write in the air, we believe, is a more challenging problem.

### **Smart Pen.**

Livescribe Smartpen [11] is a pen-like device capable of tracking a person’s writing. The device requires a special finely dotted paper to monitor movement of the pen. The pen recognizes the alphabets and numbers, and hence, can be downloaded to a PC. However, the dotted paper may not be always accessible, making ubiquitous note-taking difficult. Tablet PCs also suffer from this problem of ubiquitous accessibility.

## 8 Conclusion

This paper attempts to exploit the accelerometer in mobile phones to develop a new input technology. While today's users are mostly used to keyboards and touch-screens, we propose to mimic a pen. By holding the phone like a pen, the user should be able to write short messages in the air. The phone identifies the hand gestures as one of multiple strokes, compares the sequence of strokes against a grammar, and recognizes the air-written alphabets. The entire process requires negligible training, and can run entirely on the phone's processor. The written message is displayed on the phone-screen, and may also be emailed to the user if she desires. We believe that in the age of microblogging and tweeting, such input devices may be effective to note down information on the fly. Moreover, the pen may offer an intuitive user-experience, adding to the menu of current input methods. We call this system PhonePoint Pen, and demonstrate its feasibility through a Nokia N95 prototype and real user studies. The performance results are promising, while the user feedback (from the student community) is highly positive.

## References

- [1] AiLive LiveMove pro. AiLive Inc. <http://www.ailive.net/liveMovePro.html>.
- [2] Microsoft Research. Write in The Air, TechFest 2009. <http://www.youtube.com/watch?v=WmiGtt0v9CE>.
- [3] S. Agrawal, I. Constandache, S. Gaonkar, and R. Roy Choudhury. Phonepoint pen: using mobile phones to write in air. In *MobiHeld '09: Proceedings of the 1st ACM workshop on Networking, systems, and applications for mobile handhelds*, 2009.
- [4] A. Arranz. Niime. <http://www.niime.com/>.
- [5] V. Balakrishnan and P. H. Yeow. Sms usage satisfaction: Influences of hand anthropometry and gender. In *Human IT 9.2*, pages 52–75, 2007.
- [6] V. Balakrishnan and P. H. Yeow. A study of the effect of thumb sizes on mobile phone texting satisfaction. In *Journal of Usability Studies, Volume 3, Issue 3*, pages 118–128, 2008.
- [7] T. Baudel and M. Beaudouin-Lafon. Charade: remote control of objects using free-hand gestures. *Commun. ACM*, 1993.
- [8] X. Cao and R. Balakrishnan. Visionwand: interaction techniques for large displays using a passive wand tracked in 3d. *ACM Trans. Graph.*, 2004.
- [9] J. Kela, P. Korpipää, J. Mäntyjärvi, S. Kallio, G. Savino, L. Jozzo, and D. Marca. Accelerometer-based gesture control for a design environment. *Personal Ubiquitous Comput.*, 2006.

- [10] J. Liu, Z. Wang, and L. Zhong. uWave: Accelerometer-based Personalized Gesture Recognition and Its Applications. Mar. 2009.
- [11] LiveScribe. Smartpen. <http://www.livescribe.com/>.
- [12] Logitech. Air mouse. <http://www.logitech.com>.
- [13] Nintendo. Wii console. <http://www.nintendo.com/wii>.
- [14] Nintendo. Wii motion plus. <http://www.nintendo.com/whatsnew>.
- [15] Nokia. Virtual keyboard. <http://www.unwiredview.com/wp-content/uploads/2008/01/nokia-virtual-keyboard-patent.pdf>.
- [16] Pyacceleremoter project. <http://serk01.wordpress.com/pyacceleremoter-for-s60/>.
- [17] C. Soriano, G. K. Raikundalia, and J. Szajman. A usability study of short message service on middle-aged users. In *OZCHI '05: Proceedings of the 19th conference of the computer-human interaction special interest group (CHISIG) of Australia*, pages 1–4, Narrabundah, Australia, 2005.
- [18] C. Sung-Do, L. A.S., and L. Soo-Young. On-line handwritten character recognition with 3d accelerometer. *International Conference on Information Acquisition*, 2006.
- [19] L. von Ahn, M. Blum, and J. Langford. Telling humans and computers apart automatically. *Commun. ACM*, 2004.
- [20] J. O. Wobbrock, A. D. Wilson, and Y. Li. Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. In *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology*, 2007.
- [21] S. Zhang, C. Yuan, and Y. Zhang. Handwritten character recognition using orientation quantization based on 3d accelerometer. In *Mobiquitous '08: Proceedings of the 5th Annual International Conference on Mobile and Ubiquitous Systems*, 2008.
- [22] D. Zhuxin, U. C. Wejinya, Z. Shengli, S. Qing, and W. J. Li. Real-time written-character recognition using mems motion sensors: Calibration and experimental results. In *ROBIO '09: Proceedings of the 2008 IEEE International Conference on Robotics and Biomimetics*, 2009.