

Using Poll Sheets and Computer Vision as an Inexpensive Alternative to Clickers

James Gain
Computer Science Department
University of Cape Town
jgain@cs.uct.ac.za

ABSTRACT

Classroom Response Systems, colloquially known as Clickers after the predominant hand-held input device, are widely used because they allow students to respond in class to questions posed by the lecturer. This improves active learning and interaction in large classes—students are more involved, and lecturers can assess understanding and even take remedial action. Unfortunately, Clicker systems are relatively expensive, particularly in a developing-world context. They typically cost \$200–700 for a base station and \$30–50 per Clicker.

In this paper we present an inexpensive alternative to Clickers. Poll sheets with coloured blocks printed on a white background are held up by the students and a camera-phone is used in panoramic mode to photograph the class. This image is then processed using computer vision to count and classify the students' responses.

While the 85% recognition rate we achieve is certainly not as accurate as Clickers, this approach nevertheless has many of the same benefits for active learning at a fraction of the cost: \$0.20 per poll sheet, assuming a laptop and camera-phone are already available.

Categories and Subject Descriptors

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces; K.3.1 [Computers and Education]: Computer Uses in Education

Keywords

Classroom Response System, Clickers, ICT4D, Active Learning

1. INTRODUCTION

Classroom Response Systems, or Clickers, are an educational technology that enables students to select answers to multiple choice questions presented in class, enter their choice on a keypad device (the eponymous Clicker) and have the results transmitted to the teacher's laptop for immediate display. This is similar to the 'ask the audience' option on the popular quiz show 'Who wants to be a millionaire?' More advanced systems allow numeric response in addition to four- or five-way choice.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SAICSIT '13, October 07-09, East London, South Africa
Copyright 2013 ACM 978-1-4503-2112-9/13/10 ...\$15.00
<http://dx.doi.org/10.1145/2513456.2513468>.

Clickers have been widely adopted, with one manufacturer citing their use in over 1,300 universities, with upwards of three million Clicker units sold¹. Studies report on their use for a range of class sizes (from 15 to 200 or more), subjects (from Computer Science to Psychology) and presentation formats (including lectures, tutorials and peer instruction) [1], and unreported educational practice is doubtless even more diverse.

This is more than just a race among educational institutions to adopt the latest technology. There is mounting evidence [4, 7, 5] that Clickers, combined with appropriate pedagogy, increase students' attention, responsiveness and lecture attendance; enhance awareness of students' problems among instructors; and even, in some cases, improve course grades. They are particularly useful in large-class settings, where other forms of active learning are difficult to incorporate.

However, Clickers are costly. The handheld Clicker units used by students typically range from \$30 to \$50, with sophisticated models costing up to \$100. The dominant technology is radio frequency-based, which requires a base station in each classroom, at around \$200-\$700. Indirect costs, such as IT support and training, and equipment updates and repairs, should also be factored in.

In cases where such costs are prohibitive, we propose the following poll-sheet alternative: students are issued with a double-sided sheet printed with red, green, blue and black squares on a white background (Figure 1[A]), which can be folded to display the student's choice in response to a four-choice poll (such as shown in Figure 1[B]); the lecturer captures class-response by taking a panoramic photograph using a camera-phone (Figure 1[D]); this is then uploaded to a laptop, where it is processed using standard Image Processing to count the number of squares of each colour (as described in Section 3 and shown in Figure 1[E]) and displayed as a bar chart (Figure 1[C]). Of course, response cards are not a new idea [3], but the difference here is that automated counting makes them usable in larger classes.

Admittedly, our approach is not without flaws. Unlike Clickers, responses cannot be tracked on a per-student basis, limiting individualized marking and diagnostics. Also, individual responses are visible to the class, which may inhibit some students. Finally, it can never be as accurate as Clickers, given students' propensity to accidentally angle, occlude, move, or otherwise complicate the image processing. Nevertheless, these shortcomings must be balanced against the cost. Printing a double-sided poll sheet costs approximately \$0.20, two orders of magnitude less than the cheapest Clicker.

2. PREVIOUS WORK

Clickers are not the only way to implement a Classroom

¹<http://www1.iclicker.com/higher-education-responseware>, retrieved 5 June 2013

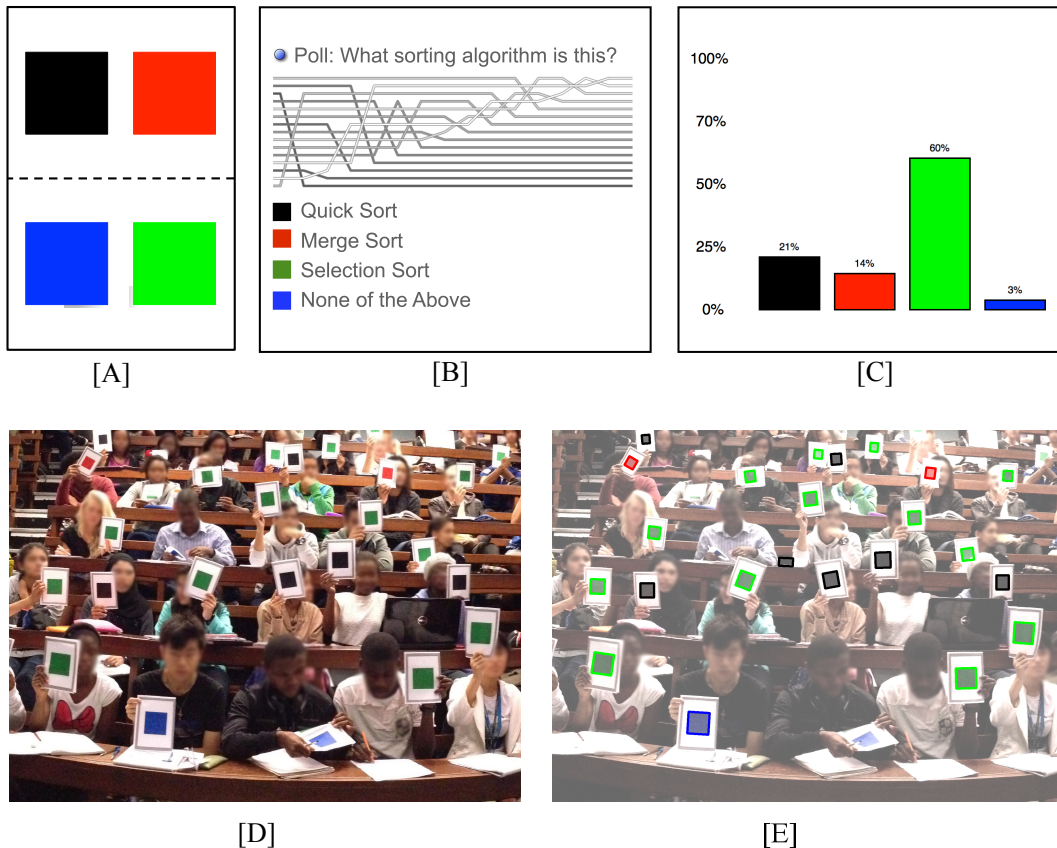


Figure 1: Poll Sheet Capture. [A] two-sided poll sheet, to be folded by student, [B] sample poll question, [C] class response graph, [D] subsection of panoramic class photograph (with faces blurred to protect anonymity), and [E] processed photo with identified sheets marked in blue.

Method	Max Class	Individ.	Anon.	Accuracy	Cost	Comments
Clickers	1000s	Y	Y	100%	\$200-700+\$30-50 per stud.	
Smartphones	1000s	Y	Y	100%	\$0	must be owned by every student
SMS	1000s	Y	Y	90-100%	\$0.05 per stud. per poll	potential delays
Fiducial Markers	25	Y	N	95-97%	\$0.10 per stud.	
Imaged Poll Sheets	250	N	N	85%	\$0.20 per stud.	not tested above 250

Table 1: Comparison of CRS technologies, tabulating the maximum class size, whether individual tracking and peer anonymity are supported, the capture accuracy and cost on a per student basis, for different methods.

Response System (CRS). Another option is to employ cell-phones. Students download a polling application and use this to vote wirelessly. While this has negligible cost, it is really only tenable if every student already owns a smartphone. More viable is a hybrid Clicker and cellphone system, as offered by some CRS manufacturers.

Some researchers [8] advocate a least common technological denominator strategy, using SMS cellphone messaging functionality. Certainly, basic cellphone adoption is closer to universal among students, even in a developing-world context. However, students now incur a per-poll cost, which can mount up if polling is employed actively across several courses. It also presupposes reliable and timely SMS transmission.

More directly comparable to our solution is recent work [2, 6] that explores the use of fiducial markers. These are paper sheets printed with a QR-like code that uniquely identifies a student. Votes are registered by holding the marker page in one of four orientations. These systems achieve better accuracy than we do by using a camera to progressively capture votes, but are limited to smaller class sizes because of the image area required for markers.

Table 1 provides a comparative summary of these different approaches.

3. IMAGE PROCESSING OF POLLS

To process the panoramic photographs of class polls we employ a variety of standard image processing techniques (using Python bindings for the OpenCV Computer Vision library²). Our strategy is to prevent misidentification of extraneous elements as poll responses even if this means we miss some valid but difficult-to-detect cases. That is, we accept an increase in false negatives in order to reduce false positives.

We begin by running a Canny edge detector, to extract a binary silhouette image. A high minimum threshold is appropriate here because it reduces soft background edges. In contrast, poll sheets with their saturated coloured squares on a white background produce a strong edge response. Next, blob detection is used to segment out regions surrounded by edges. Because Canny edge detection does not always produce a closed loop around the central

²opencv.org

marker (for instance, if a student places a finger across the poll sheet) we run a single iteration of image erosion to separate blobs with a narrow necking connection. An enclosing non-axial rectangle is then fitted to each blob.

At this stage there are a significant number of false positives. We cull these based on a variety of shape and colour tests. Specifically, the size and aspect ratio of the enclosing box must fall in certain ranges in order to represent a valid square marker. The size constraint is fairly loose to compensate for varying distance from the camera. Fortunately, in raked lecture theaters we can use vertical image position as a proxy for distance and a scale factor on the size check, because markers that are higher will be towards the back rank of the lecture theatre. Of course, this refinement must be disabled in a flat-space venue. We also use fill ratio—the number of pixels in a blob relative to the rectangle area. This ratio will be close to 1 for blobs that fill their enclosing rectangle, as opposed to those having a more amorphous shape. For shape tests, the relevant parameters are set as follows:

$$\begin{aligned} 0.008I_h < r \min(B_h, B_w) < 0.1I_h \\ 0.5 < a < 2.0 \\ 0.6 < f < 1.1 \end{aligned}$$

where a is aspect ratio, I and B are the image and enclosing box, respectively, with w and h being their width and height parameters, f is the fill ratio and $r = B_h/I_h + 0.1$ is the raking scale factor.

On the colour front, we expand the enclosing rectangle slightly and use the resulting offset corners to check for the white background around markers. The blob’s mean colour and variation are also tested against their exemplars (red, green, blue and black). In all cases we allow considerable variation to compensate for differences in lighting, poll sheet angle and various forms of occlusion. In the end it is the variety of tests rather than their individual strictness that cuts down on misclassification.

If a blob passes all of these tests, then its mean colour is used to determine the marker classification.

Red	Poll Colours			Black	Acc.	Time
	Green	Blue				
71 (61)	28 (31)	12 (11)	4 (7)	85%	12.3s	
10 (10)	98 (86)	4 (4)	13 (16)	88%	11.0s	
33 (24)	28 (25)	34 (29)	20 (20)	85%	10.0s	
15 (16)	54 (37)	22 (21)	12 (13)	81%	14.9s	
6 (7)	4 (8)	0 (2)	66 (66)	91%	13.0s	
56 (48)	7 (8)	1 (1)	2 (9)	76%	13.1s	
12 (21)	81 (67)	6 (3)	18 (20)	76%	14.0s	
19 (24)	21 (17)	26 (25)	9 (7)	84%	12.9s	
7 (7)	11 (10)	3 (3)	14 (14)	97%	10.9s	
76 (61)	22 (19)	20 (19)	5 (5)	85%	14.9s	
56 (54)	21 (13)	1 (1)	0 (5)	81%	11.4s	
1 (4)	1 (1)	0 (1)	47 (46)	90%	9.7s	

Table 2: Results of Image Processing based on 12 test images, taken live in a real class setting. The first four columns represent a human count of answers in each colour (with the image processed count in parentheses). We calculate accuracy as the sum of absolute differences between recognized and actual markers as a proportion of the total number of markers present. Time is the duration of the image processing and does not include time required to capture and upload the panoramic image.

The efficacy of our approach was tested on 12 poll images (as shown in Table 2), taken during a first-year Computer Science course in 2013. Two venues were tested: a wide 300-seater venue and a steeply-raked 160-seater

venue (first and last six entries in Table 2, respectively). The input image resolutions were relatively consistent (from 7600×2348 to 8832×2386). The number of votes ranged from 35 to 125, with a good distribution between images - in some one poll colour predominates and in others there is a mixture of responses. We gave very little instruction to the class, apart from the fact that their poll sheets needed to be visible to the camera. The software was run on a single core 1.7 GHz Intel Core i5 processor with 4 GB RAM, as representative of the type of laptop commonly used in a lecture setting.




Figure 2: Problem Cases. [A] Wash out, [B] Orthogonal Rotation, [C] Hand Occlusion, and [D] Body Occlusion.

On average, we achieve 85% recognition accuracy in under 15 seconds. The problem of recognizing poll sheets is complicated by wash-out effects, related to the angle between poll sheets and the classroom lighting, perspective distortions, when poll sheets are inclined relative to the camera, and occlusion by hands or other students. Some representative problem cases are shown in Figure 2. On a positive note, there is little time-dependence on the number of poll sheets in our results, so there is no reason that this approach should not scale to larger classes as long as the proportion of problem cases does not increase.

4. RESULTS AND DISCUSSION

Image-based polling as outlined here was trialled for 5 weeks in a first-year university-level Computer Science course in Introductory Programming, and we discovered a number of considerations around using this approach in practice.

First, the processing time alone (10–15s in Table 2) is not a true reflection of in-class delay. The time required to capture a panoramic photograph (9–12s) and upload it to a laptop (5–10s) must also be factored in. Realistically, it can take up to 40 seconds to capture and display an imaged poll.

Consideration must also be given to how to issue poll sheets. We experimented with the two obvious strategies: handing out and retrieving poll sheets on a per-lecture basis and issuing them permanently to students. While the former approach requires some management, it avoids the all too prevalent issue of misplaced poll sheets. (On a side note, we recommend printing poll sheets for large classes outside of normal office hours).

In terms of participation, on average, 72% of the class voted during polls, (with a wide range of 51% to 90%).

(A) Polls: "The in-class polls (using printed poll sheets) improved my understanding of the course material."

1. Strongly Disagree	2. Disagree	3. Neutral	4. Agree	5. Strongly Agree	Mean	N
3%	5%	22%	46%	24%	3.83	177

(B) In-class Exercises: "The in-class exercises (begun by you as a student, completed by the lecturer in class) improved my understanding of the course material."

1. Strongly Disagree	2. Disagree	3. Neutral	4. Agree	5. Strongly Agree	Mean	N
2%	2%	23%	50%	24%	3.95	177

(C) Lecture Recording: "The lecture recordings improved my understanding of the course material."

1. Strongly Disagree	2. Disagree	3. Neutral	4. Agree	5. Strongly Agree	Mean	N
3%	5%	29%	34%	28%	3.76	177

(D) Recap Workshop: "The recap workshop improved my understanding of the course material."

1. Strongly Disagree	2. Disagree	3. Neutral	4. Agree	5. Strongly Agree	Mean	N
1%	4%	38%	40%	18%	3.69	85

Table 3: Survey Results. Imaged polls are rated comparably by students to other learning supports, such as in-class exercises, video recording of lectures, and revision workshops.

These results are in line with reported proportions of Clicker usage [9]. We noticed that participation tended to be higher when poll sheets were handed out during lectures, which is another reason to adopt a distribute-and-retrieve approach.

We also surveyed the class to see how they rated image-based polls, in relation to a number of other active learning initiatives (see Table 3). The surveyed alternatives included: in-class exercises (undertaken in pairs, with a solution presented by the lecturer), a half-day revision workshop (which involved working through a problem sheet under the guidance of the lecturer and tutors), and video recording of the lectures. Students were asked, on a 5-point Likert scale ranging from Strongly Disagree (1) to Strongly Agree (5), whether they felt that the initiative in question improved their understanding of the course material. The mean for Polls (3.83) was below In-class Exercises (3.95) but above Lecture Recording and the Recap Workshop. However, these differences are not statistically significant. Interestingly, one student commented in the survey: "they [polls] made me want to take part in the class quizzes, which made me listen in class. Overall, it was effective." This lends support to the notion that polls are a spur to active learning.

While we found the recognition accuracy adequate for teaching purposes, particularly since there is no tracking of individual responses, there is certainly room for improvement. Larger markers could be used, perhaps with a single square per A4 sheet and this would have only minor cost implications. However, we believe that the large A3 sheets used in previous work on fiducial tracking [2, 6] would prove unwieldy in practice. Also, the students could be given more instruction on visible voting to prevent angling and occlusion, although it would be impossible to get rid of this entirely. Finally, a video camera could be used to track markers more accurately over several video frames, but allowance would need to be made for wide venues.

5. CONCLUSION

We have presented a low-cost Classroom Response System, based on poll sheets and image processing, averag-

ing 85% accuracy in polling medium-sized classes (tested with up to 125 students). While our approach does not support peer anonymity and per-student response tracking, it is significantly less expensive than Clickers, making it appropriate where cost is the deciding factor or lecturers wish to investigate polling before committing to a full Clicker solution.

6. REFERENCES

- [1] J. E. Caldwell. Clickers in the large classroom: Current research and best-practice tips. *CBE-Life Sciences Education*, 6(1):9–20, 2007.
- [2] A. Cross, E. Cutrell, and W. Thies. Low-cost audience polling using computer vision. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*, pages 45–54. ACM, 2012.
- [3] W. L. Heward et al. Everyone participates in this class: Using response cards to increase active student response. *Teaching Exceptional Children*, 28(2):4–10, 1996.
- [4] D. Johnson and S. McLeod. Get answers: Using student response systems to see students' thinking. *Learning & Leading with Technology*, 32(4):18–23, 2005.
- [5] J. K. Knight and W. B. Wood. Teaching more by lecturing less. *Cell Biology Education*, 4(4):298–310, 2005.
- [6] M. Miura and T. Nakada. Device-free personal response system based on fiducial markers. In *Wireless, Mobile and Ubiquitous Technology in Education (WMUTE), 2012 IEEE Seventh International Conference on*, pages 87–91. IEEE, 2012.
- [7] J. Roschelle, W. R. Penuel, and L. Abrahamson. The networked classroom. *Educational Leadership*, 61(5):50–53, 2004.
- [8] E. Scornavacca and S. Marshall. Txt-2-learn: improving students' learning experience in the classroom through interactive SMS. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pages 1–8. IEEE, 2007.
- [9] P. J. White, D. G. Delaney, D. Syncox, O. A. Akerberg, and B. Alters. Clicker implementation models. *EDUCAUSE Quarterly*, 34(4):1–4, 2011.