

Photonote evaluation: aiding students with disabilities in a lecture environment

Gregory Hughes
gregory.hughes@cl.cam.ac.uk

Peter Robinson
pr@cl.cam.ac.uk

Computer Laboratory, University of Cambridge
15 JJ Thomson Avenue, Cambridge UK CB3 0FD

ABSTRACT

Visual material presented in lectures can be enhanced for students with disabilities by using high-resolution digital-still cameras. The Photonote system uses a digital-still camera to capture visual information, a digital-video camera to capture a lecturer and a second digital-video camera to capture a sign-language interpreter, if necessary. The visual information is enhanced using computer-vision algorithms and presented alongside the recorded video and audio to provide an accurate representation of a lecture which can be used by students with disabilities for **review** purposes. This paper presents the Photonote system and a user study evaluating its effectiveness at aiding students with disabilities in a lecture environment.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture; I.4.3 [Image Processing and Computer Vision]: Enhancement; K.3.1 [Computer and Education]: Computer Uses in Education

General Terms

Algorithms, Experimentation, Human Factors

1. INTRODUCTION

Taking accurate notes during a lecture can be a difficult task for many students with disabilities due to the inherent visual nature of the material, oral nature of the lecture, and/or the physical act of taking notes. By utilising high-resolution digital cameras it is possible to capture the contents of a whiteboard or overhead projector (OHP) in real time to allow future presentation to a student with a disability without requiring the lecturer to adapt to new technologies[5]. Our system, called Photonote, combines enhanced high-resolution imagery of displayed information with video of a lecturer and sign-language interpreter to provide a lecture review system for students with disabilities. Currently these students employ many methods to compensate for their disability, including university-appointed note-takers, copies of friends' or lecturers' notes, or simply aiming to teach themselves the material from text books.

2. RELATED WORK

Various systems have been designed to capture and record lectures, however none has been designed specifically to aid students with disabilities. The eClass project was started at the Georgia Institute of Technology with the initial goal of creating a "classroom environment in which electronic notes taken by students and teachers could be preserved and augmented with audio and video recordings"[1]. Using a 72" touch screen display, called a LiveBoard, a professor can write notes into proprietary software called Zen-Pad. This system required training, and even the researchers of eClass[1] admit that lecturers generally use whiteboards three times the size of the 72" LiveBoard which their system implements. The Lecture Browser[6] system was designed to process a lecture for later viewing, however much like eClass, this system requires the lecturer to be aware of the system and to upload PowerPoint files to a server after the lecture has finished. AutoAuditorium[4] is a commercial system designed to automate the lecture-recording process. At minimum it is composed of two static cameras and one automatic pan/tilt camera. The pan/tilt camera is used to follow the presenter during the lecture, one camera is used to capture an overall view of the theatre, while a third camera is used to capture the output of a digital projector. This system does not require the lecturer to be conscious of its presence. However, this system is designed to create an aesthetically-pleasing video rather than provide assistance for students with disabilities. **Apreso[2] is a popular lecture-recording system used by many universities including Temple University, the University of Pittsburgh, Princeton University, and the University of Tennessee. This system can capture presented material using many different methods. The primary capture method forces the lecturer to use a computer to present visual information. Apreso can also be configured to capture visual information with a standard-definition video camera. As described in Section 3.1, standard definition video is not capable of capturing enough detail to read presented visual information.**

Numerous other lecture-recording systems are both being researched and available commercially; such as Lectopia[3] and Tergrity[7]. There are two primary differences between Photonote and any system, either in development or in production. First, Photonote is designed to allow a lecturer to present information in their normal manner without requiring the lecturer to be aware of the system while still being able to capture intricate visual information. Second, Photonote is designed specifically to address the needs of students with disabilities; a large group of students neglected by the



Figure 1: A screenshot of the Photonote application. (a) Visual information windows (b) Video of the lecturer (c) Video of the sign-language interpreter (d) Thumbnail of the visual information (e) Play/Pause button (f) Time slider

majority of lecture recording systems.

3. PHOTONOTE

Photonote is a system we developed to aid students with disabilities which impair their ability to take accurate notes to review material presented in a lecture. The Photonote software (See Figure 1) displays three keys pieces of information simultaneously. The largest area within the application displays a high-resolution, enhanced image of the visual information which was presented during the lecture. Two videos are also displayed, one of the lecturer and one of a sign-language interpreter. Audio from the lecture is replayed for the student through a pair of standard headphones. All four of these sources are synchronised in time. A user can play/pause the videos as well as jump to any point in time within the lecture. The visual information window can be scrolled and zoomed by the user while the software is replaying the lecture.

3.1 Capturing information

Two digital-video cameras are used to capture the lecturer and the sign-language interpreter while a digital-still camera is used to capture the visual information which is presented during a lecture. Audio from the lecture is captured alongside the video of the lecturer. Although video would be the obvious choice to capture the visual information, time-lapse digital photography was chosen based on laboratory experimentation which determined that the resolution of a digital-video camera was insufficient to capture the details of a lecturer's presentation. Using a standard Snellen Eye Chart, it was determined that a standard-digital-video camera had a visual acuity of 20/100. Therefore, a digital-video camera would need to be five times closer to an object than the average human in order to distinguish the same detail without using optical zoom. Using the same method, an 8 Megapixel (MP) digital-still camera was found to have better (>20/20) visual acuity as an average human. This allows the camera to be placed anywhere in a classroom that an average human would be able to see the presented information. The camera used, a Canon PowerShot S80, is able to

capture one frame every three seconds which provides ample frequency since presented information on a whiteboard or OHP will not change more than once within a three second time-span. The camera is controlled by a Windows PC using a USB 2.0 interface to allow for the quick transfer of the captured photos.

3.2 Information enhancement

Visual information can be enhanced for students with disabilities using a number of different techniques. In our system, we employ computer-vision techniques to compensate for the angle at which the visual information is captured, enhance the contrast by running an adaptive threshold algorithm, and also remove obstructions and highlight new additions (See Figure 2). The final result of this process (Figure 1(a)) provides a very legible copy of the visual information which can be zoomed and changed to suit the needs of any individual using the system.

The visual enhancement begins with a perspective transformation which compensates for the angular difference between the camera and the source of visual information. This technique is generally used to create three-dimensional renderings in a two-dimensional image space. A perspective transformation can also be used to convert from one image plane to another image plane as in our scenario. A possible source of visual information is treated as a plane and can be specified using the four corners of the visual information, such as the four corners of a whiteboard. The other plane can be specified by using the four corners of the image. The perspective transformation matrix can be derived by solving the following set of linear equations where (x_n, y_n) are the four corners of the piece of visual information and (u_n, v_n) are the four corners of the image.

$$\begin{pmatrix} x_0 & y_0 & 1 & 0 & 0 & 0 & -x_0u_0 & -y_0v_0 \\ 0 & 0 & 0 & x_0 & y_0 & 1 & -x_0v_0 & -y_0v_0 \\ x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1u_1 & -y_1v_1 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -x_1v_1 & -y_1v_1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 & -x_2u_2 & -y_2v_2 \\ 0 & 0 & 0 & x_2 & y_2 & 1 & -x_2v_2 & -y_2v_2 \\ x_3 & y_3 & 1 & 0 & 0 & 0 & -x_3u_3 & -y_3v_3 \\ 0 & 0 & 0 & x_3 & y_3 & 1 & -x_3v_3 & -y_3v_3 \end{pmatrix}^{-1} \begin{pmatrix} u_0 \\ v_0 \\ u_1 \\ v_1 \\ u_2 \\ v_2 \\ u_3 \\ v_3 \end{pmatrix} = \begin{pmatrix} a \\ b \\ c \\ d \\ e \\ f \\ g \\ h \end{pmatrix} = \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{pmatrix}$$

The four corners of the visual information can be defined by the user during a simple calibration phase where various information sources can be selected. Alternatively, we have experimented with using a square detection algorithm to locate whiteboards, however this was found to return undesirable results in a large classroom or when multiple potential sources of visual information are present.

After compensating for the angular difference between the information and the camera, whiteboards and OHP output can be further enhanced by using a modified version of the adaptive threshold algorithm developed by Wellner for use on the DigitalDesk[8]. This algorithm traverses the image left

to right and top to bottom while keeping track of the average colour of the last x number of pixels. When a pixel is encountered that is significantly darker than this running average, it is coloured black, while all other pixels are coloured white. The algorithm was modified slightly to exclude the dark pixel values from the running average. This algorithm helps to threshold an image that is not evenly lit, however it requires that the image has more white space than dark space, such as a piece of paper, an OHP projection, or a whiteboard.

Temporary obstructions can be removed and new information highlighted by analysing the image after the adaptive threshold algorithm is executed. The image is traversed pixel-by-pixel, feeding each pixel into a blob-detection algorithm, along with a threshold. The blob-detection algorithm then searches surrounding pixels to find pixels with colour values within the given threshold. Each pixel within the threshold is marked as part of the blob, and then the pixel is fed as input back into the blob-detection algorithm. This method is recursive, ending when no more touching pixels are found with colour values within the threshold. The algorithm finally returns the total number of pixels in the blob, along with the blob's height and width. A threshold value is determined depending on the image being analysed. Blobs containing fewer pixels than this threshold value are considered to be text, and anything larger is considered to be a person, or some unimportant obstruction.

Obstructions can now be removed by analysing the large blobs and resorting to information from the same area in previous frames where no obstruction was present, which may require the system to check a number of different frames to find one without an obstruction. Although this enhancement can be useful to many students, it can present false information. If the lecturer was obstructing the whiteboard while erasing and writing new information the system would continue to display the old information.

Highlighting and detecting the addition of new text is important because the user can lose spacial orientation while zoomed in on visual information and may be unaware of new information presented in other locations. By detecting new text the system can either automatically centre on the new information, or indicate to the user that new information has been added elsewhere.

4. EVALUATION

To evaluate the effectiveness and usefulness of the Photonote system for students with disabilities, we conducted a user study. The goal was to determine if our electronic review system could replace the job of a human note-taker utilised by students with disabilities which impair their ability to take accurate notes during a lecture. Our hypothesis is that students with disabilities will perform equally well, or better, on an examination when using their own notes and Photonote to review lecture material, compared to an examination where the students use their own notes along with the notes of a note-taker to review material.

4.1 Participants

The study participants included people both with and without disabilities. The study was conducted using a total of 33 participants aged between 18 and 61 with a mean age of 27. The participants were divided into 5 groups based upon their disability type. The disabilities being investigated in-

Table 1: Groups of participants for the user study.

Group (subgroup)	Disability	Number of members
1	Hearing	6
(1a)		2
(1b)		3
(1c)		1
2	Vision	9
(2a)		5
(2b)	4	
3	Mobility	2
(3a)		1
(3b)	1	
4	Learning	6
(4a)		4
(4b)		2
5	None	10
(5a)		3
(5b)		4
(5c)	3	

Table 2: Participants groups and the method utilised during the review period.

	Lecture 1	Lecture 2
"A" Subgroups	Method 1	Method 2
"B" Subgroups	Method 2	Method 1
"C" Subgroups	Control	Control
Test Notation	Exam 1 (E_1)	Exam 2 (E_2)

cluded hearing impairments, visual impairments, learning disabilities, and mobility impairments. **The disabilities of all participants with disabilities inhabited their note taking ability in some manor.**

Each group was then split into subgroups (see Table 1) based on their time of arrival at the review session.

Participants were paid a total of \$100 each for their time and transportation costs. All participants were provided with an informed consent letter. Participants were chosen for the user study based on their response to a preliminary questionnaire. This questionnaire asked questions such as age, highest education level and disability. Preference was given to those applicants that currently used a note taker in a university setting, followed by applicants who were currently enrolled in a university. Once the study started, all participants were identified only by a unique identification number.

4.2 Methods

The study had a repeated measure design where participants used two different methods to prepare for an examination. All participants attended two consecutive lectures on the 6 January 2007. Oceanography was chosen as the subject for the lectures as it is a topic that few people know much about, yet it is still interesting to a general audience. **It is unknown if any participant had previous knowledge, however the repeated-measure design of the study should render this variable irrelevant.** The first lecture covered micro-organisms, and the second lecture covered sea-floor spreading. The lecturer was Professor John Merrill, Ph.D. of the Oceanography Department within the University of Rhode Island. An American Sign Language (ASL) inter-

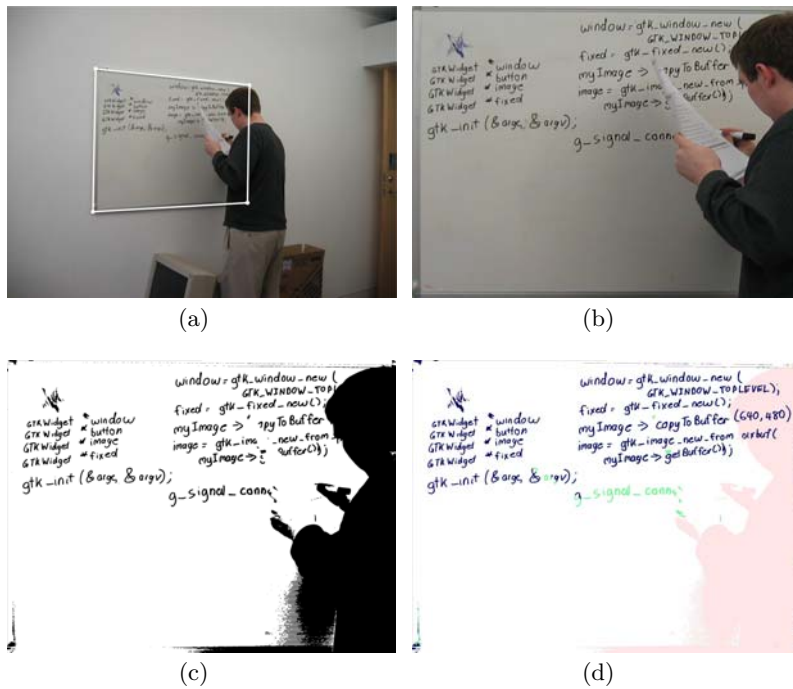


Figure 2: Demonstrates the enhancement process for a whiteboard within a classroom.

preter was provided to translate the lectures for participants with hearing impairments. During the lectures, participants took notes as they normally would in a lecture environment.

On the 13 January 2007, a one-and-a-half hour review session was held where students reviewed the two lectures using one of the two methods. Method One entailed using the Photonote system to review the material presented in a lecture along with the participant's own lecture notes. For Method Two participants used their own lecture notes to review material, along with the notes of a note taker if the participant had a disability. Each individual using Method Two received notes from the same two note-takers which eliminated variance in the provided notes. Participants had 45 minutes to review each of the two lectures. During the first 45 minute session, everyone in an "a" subgroup used Method One to review Lecture One, while everyone in a "b" subgroup used Method Two (See Tables 1 and 2). During the last 45 minutes, Lecture Two was reviewed using Method One by everyone in "b" subgroups, while everyone in "a" subgroups used Method Two. Participants within "c" subgroups were used as a control to determine how difficult the exams were and did not use the Photonote system at all. However, one of the participants in subgroup "c" was provided with the note-taker's notes for both review sessions because the participant had a disability. Participants in group 5 were not given a copy of a note-takers' notes at any point as it is only a provision provided to students with disabilities. This group was used to see if the system aids students without disabilities.

A 45 minute written examination followed the review period to test the students' retention and understanding of the information presented in the lectures. The examination consisted of 20 true-or-false questions and 30 multiple-choice questions. All odd numbered questions pertained to Lecture One, while even numbered questions pertained to

Lecture Two. After the examination, it was discovered that question number 26 had no correct response, thus there was a total of 25 questions on material from Lecture One and 24 questions on material from Lecture Two. Although they were combined as one examination, for the purpose of this study they will be considered as two separate examinations. The odd questions will be considered Exam One (E_1), while the even questions will be considered Exam Two (E_2). Furthermore, we can define M_1 and M_2 to be the exam that a participant took after utilising Method One and Method Two respectively. M_1 is equal to E_1 and M_2 is equal to E_2 for all "a" subgroups, and vice-versa for all "b" subgroups.

Following the examination a separate questionnaire was filled in by all participants. The final questionnaire asked questions regarding the participants' experiences with Photonote, the difficulties they normally experience in a lecture environment, and their familiarity with a computer.

4.3 Materials

The study was conducted at the University of Rhode Island's Providence campus in the United States. The lectures were held in a 60 seat classroom where information was presented by the lecturer using an OHP. Two digital-video cameras were used to capture the lecturer and the sign-language interpreter separately. A computer-controlled webcam was used as a backup in case either video camera failed. The digital-still camera used to capture the visual material being presented was a Canon PowerShot S80. A Canon PowerShot A80 was used as a backup camera in case the S80 failed.

When using Method One to review for an examination, participants used computers in a nearby computer lab. All 24 computers were Dell Optiplex Gx620 desktops with 2.8 GHz processors, 2 Gb of RAM and a 256 Mb ATI Radeon video card. The terminals used a piece of software called




Figure 3: Density histogram of E_1 and E_2 with $N(0, 1)$ superimposed.

“Deep Freeze” which would restore the terminal to a stored state upon reboot thus removing any installed programs or added files. This required our software to be designed such that it can be installed and used without requiring a reboot of the system.

In the beginning of a Method One review session, participants were given a five-minute introduction to the Photonote software. During this introduction, participants were taught how to use the basic controls of the software and were told that they could ask for assistance during the review session if necessary.

5. QUANTITATIVE RESULTS

As previously stated our study was a repeated measure design, however the measures are obtained from different examinations, namely E_1 and E_2 . To account for this difference, z-values were used as a standard score for each participant’s exam results.

The histogram of exam z-scores seen in Figure 3 shows that the results of the two exams approximate to a normal distribution. We also need to check that the two exam scores are correlated; meaning that when a participant does well on one exam they also do well on the other exam. The correlation coefficient, $r=0.79$, between E_1 and E_2 shows that there is a reasonably strong positive correlation between the two exam results. A Pearson product-moment correlation test was used to determine that this correlation was highly statistically significant ($p < 0.01$). Given this correlation of the two exams, we can say that the exams were relatively equal and it is reasonable for us to compare M_1 and M_2 when examining the validity of our hypothesis and examining other correlations within our dataset.

In order to examine the effectiveness of our system in helping participants **review** for an examination, we must investigate the relationship between M_1 and M_2 . The correlation coefficient is 0.81 and a Pearson product-moment correlation test showed that the correlation was highly statistically significant ($p < 0.01$). This demonstrates that participants are performing consistently on both exams, M_1 and M_2 .




Figure 4: Density histogram of ΔM with $N(0, 1)$ superimposed.

This result tells us that the exam scores are correlated, but it does not show that students using our system perform equally well, or better, on an examination as those using their own notes, and a note-taker’s notes if they have a disability.

To examine the differences in exam scores when using Photonote compared to using a note-taker’s notes, we define the value ΔM to be the difference between M_1 and M_2 . The resulting value will be positive if our system helped the participant, and negative if our system impeded the participant. A histogram of these values can be seen in Figure 4. If Photonote helped or did not impede the participant, then we would expect the true mean of ΔM to be greater than or equal to zero. ΔM has a range of -1.47 to 1.33 with an average of -0.21 and a standard deviation of 0.64 for all participants, excluding the control group. ΔM has a range of -1.47 to 1.33 with an average of -0.18 and a standard deviation of 0.63 for all participants with disabilities, excluding the control group.

A t-test was used to test if the true mean of ΔM for the participants with disabilities would be greater than or equal to 0 in the general population which resulted in a p-value greater than 0.05 and less than 0.95. This result demonstrates that no statistically significant impact can be inferred, showing that the system is not universally helpful to students with disabilities but, more importantly, it did not universally impede students with disabilities.

An ANOVA test can also be used to test our hypothesis that the true mean of ΔM is greater than zero. Executing an ANOVA test on this data can be accomplished by examining variances in exam scores, while grouping the participants according to whether they used our system to **review** for the exam in question, and by which exam a particular score corresponds to. The results, as seen in Table 3, tells us about interactions between our variables. This ANOVA result reinforces our previous claim that the two exams were equally difficult because it shows that there was no statistically significant interaction between whether Exam One or Exam Two was taken and the exam score. We can also see that the utilisation of Photonote had no statistically sig-

nificant interaction with the exam score, which helps support our hypothesis that Photonote can be used to replace a note-taker. Table 4 shows the ANOVA results only for those participants with disabilities. Again, none of the interactions are significant preventing us from rejecting the null hypothesis.

These results show that it is likely that our system would be a suitable alternative for some students with disabilities as it can effectively replace a note-taker for some students, although it is not a universal solution. Further analysis was used to seek an explanation as to why the system helped some students but not others.

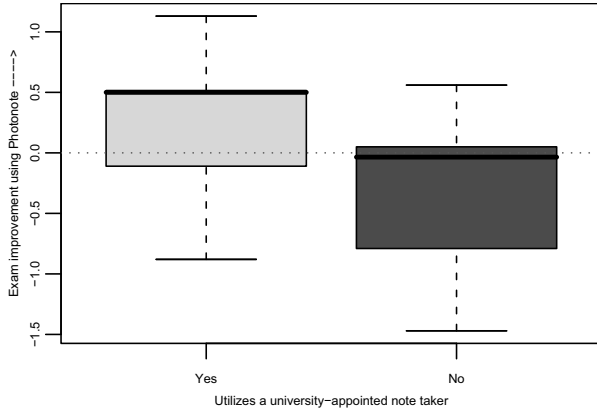


Figure 5: Box plot of change in exam results versus utilisation of a university-appointed note-taker. Boxes encompass the range of the standard deviation either side of the mean. The solid bar shows the median.

6. VARIABLES AFFECTING PERFORMANCE

Information obtained from the participants' questionnaires was used to determine why some participants were aided by Photonote while others were not. Further statistical tests were used to check if any of the following criteria had an effect on ΔM .

age; computer proficiency; considered by themselves as disabled; difficulty in school; difficulty reading a blackboard; difficulty reading a whiteboard takes helpful notes; existing methods are better than Photonote; finds note-takers helpful; finds taking notes easy; gender; has utilised a friend's notes; has utilised a lecturer's notes; has utilised a sign-language interpreter; has utilised a university-appointed note-taker; highest education level; normal study duration; registered with a university as having a disability; reported difficulty in school; reported grades in school; type of disability

And finally participants' opinion on whether

Photonote could enhance existing methods; Photonote could replace existing methods; Photonote was helpful or not; studying was easier with Photonote

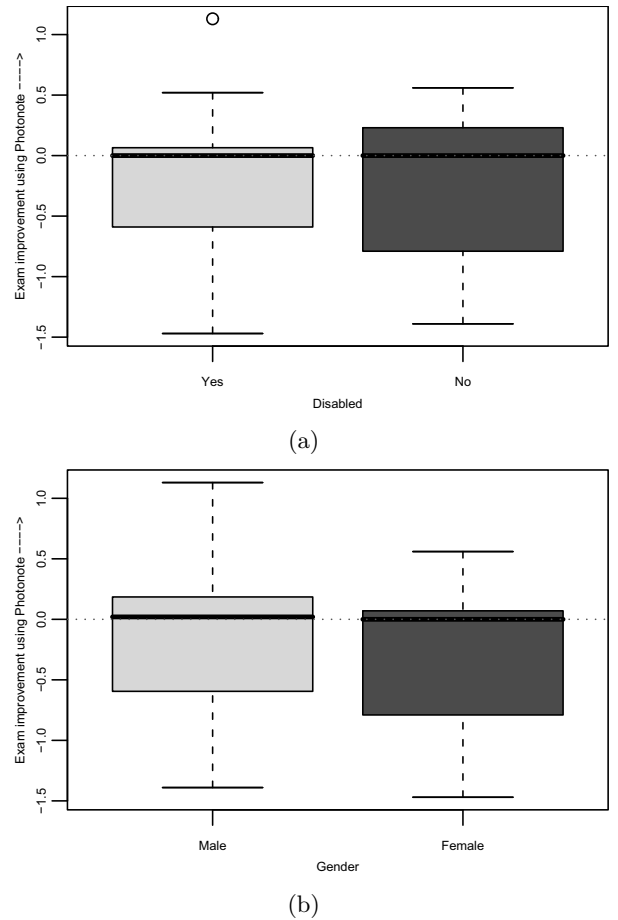


Figure 6: (a) Box plot of change in exam results versus whether the participant considers themselves as having a disability. (b) Box plot of change in exam results versus gender.

The examination of these variables, although difficult to prove statistically, does help us to try to understand why some students did better using the Photonote software, while others did worse.

Most tests were inconclusive, however an interesting interaction was found between ΔM and whether the participant normally uses a university-appointed note-taker. This interaction is visible in Figure 5, which demonstrates that if a person habitually uses a university-appointed note-taker, then they are more likely to benefit from the use of the Photonote system. This is a crucial finding as the original purpose of our system was to replace a university-appointed note-taker for students with disabilities.

Combinations of variables were also examined, for example males with disabilities or females who have difficulty in school. One variable interaction was particularly interesting: the combination of gender and whether a participant considers themselves disabled. In Figure 6 it can be seen that there were no significant differences across gender or whether participants consider themselves to be disabled or not. However, Figure 8 shows that our system helped females without a disability and males with a disability. Figure 7 shows the grouping of gender and disability type, highlighting that Photonote mainly helped males with hearing impairments and learning disabilities.

Table 3: The results of an ANOVA test comparing a participants exam result to which exam was taken to the usage of our system for review for that exam.

Error stratum	Variable	Df	SSq	MSq	F val	Pr(>F)
participant	using Photonote × exam	1	0.141	0.141	0.077	0.783
	Residuals	27	49.280	1.825		
within	using Photonote	1	0.650	0.650	3.060	0.092
	exam	1	0.0124	0.013	0.059	0.811
	Residuals	27	5.736	0.212		

Table 4: The results of an ANOVA test comparing a participants exam result to which exam was taken to the usage of our system for review for that exam. These are the results of only the participants with disabilities.

Error stratum	Variable	Df	SSq	MSq	F val	Pr(>F)
participant	using Photonote × exam	1	0.006	0.006	0.003	0.959
	Residuals	20	41.595	2.080		
within	using Photonote	1	0.359	0.359	1.723	0.204
	whichExam	1	0.061	0.061	0.292	0.595
	Residuals	20	4.159	0.208		




Figure 7: Box plot of change in exam results versus gender and the disability type of the participant.




Figure 8: Box plot of change in exam results versus gender and whether the participant considers themselves as having a disability.

7. CONCLUSION

While looking at the numbers and statistical tests, it is easy to forget that the participants are human and that it is very difficult to find universal solutions, especially when developing technologies for individuals with disabilities. Although the analysis suggests that our system could replace a note-taker for students with disabilities, it is not a universal solution, just as the utilisation of a note-taker is not a universal solution for all people with the same disability. As seen in Figure 4 some students were dramatically impaired when using our system instead of using their normal method of review.

Students with disabilities learn in many different ways and find individual ways of overcoming their disabilities. It is quite ambitious to think that a single method of adaptation would be suitable for all students with disabilities. Students with similar disabilities can have different needs, and therefore find their own way of adapting. When looking on the individual level we see that our system proved to be very beneficial to some students, improving their performance on an examination.

One of the goals of this study was to run a real-world test to see how the system would perform in a real lecture environment. Because of this goal, our sample size is relatively small and thus more prone to individual deviation. Having a relatively short examination, as in this study, leads to large difference in participant performance. However it is an examination reflective of a typical course exam in a university environment in the United States.

Through this user study, we were not able to prove our hypothesis absolutely. However, we were able to find convincing evidence to suggest that our system is a viable alternative to a note-taker for a student with a disability. We also showed that our system helped about 50% of the participants with disabilities, as seen in Table 5, despite the mean of ΔM being less than 0. Further data analysis has shown that Photonote tends to help males who consider themselves as having a disability more than females, and it is definitely helpful to students with disabilities who currently utilise a note-taker within a lecture environment. Since this is the first study conducted testing this software and lecture-capturing technique, it shows great promise for future improvements based on user feedback.

Table 5: Photonote's impact on participants.

Group (subgroup)	Desc.	Photonote aided	Photonote impaired
1	Hearing	2	3
(1a)		0	2
(1b)		2	1
(1c)		-	-
2	Vision	3	6
(2a)		2	3
(2b)		1	3
3	Mobility	1	1
(3a)		0	1
(3b)		1	0
4	Learning	5	1
(4a)		3	1
(4b)		2	0
5	None	4	3
(5a)		2	1
(5b)		2	2
(5c)		-	-
Totals		15	14

8. ACKNOWLEDGMENTS

Our thanks to all of the people who participated in our study, Peter DeCesaro, Mark Oliver, Professor John Merrill, Pamela Rohland, the Gates Cambridge Trust, and the University of Rhode Island.

9. REFERENCES

- [1] S. D. Abowd. Classroom 2000: An experiment with the instrumentation of a living educational environment. *IBM Systems Journal, Special issue on Pervasive Computing*, 38(4):508–530, 1999.
- [2] Anystream Incorporated. Apreso. Online. http://www.apreso.com/ac_product_overview.asp; Accessed 12 Jul 2007.
- [3] Arts Multimedia Centre, The University of Western Australia. Lectopia. Online. <http://ilectures.uwa.edu.au/>; Accessed 12 Jul 2007.
- [4] AutoAuditorium. Autoauditorium system home page. Online. <http://www.autoauditorium.com/>; Accessed 16 Aug 2005.
- [5] G. Hughes and P. Robinson. Time-lapse photography as an assistive tool. In *Proceedings of the 3rd Cambridge Workshop on Universal Access and Assistive Technology (CWUAAT)*, pages 87–89, Cambridge, UK, 2006.
- [6] S. Mukhopadhyay and B. Smith. Passive capture and structuring of lectures. In *MULTIMEDIA '99: Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, pages 477–487, New York, NY, USA, 1999. ACM Press.
- [7] Tegrity Incorporated. Tegrity campus 2.0. Online. <http://www.tegrity.com/>; Accessed 12 Jul 2007.
- [8] P. D. Wellner. Interacting with paper on the DigitalDesk. Technical Report UCAM-CL-TR-330, University of Cambridge, Computer Laboratory, Mar. 1994.