Computer Support for Vicarious Learning

Rachada Monthienvichienchai
M. Angela Sasse
Department of Computer Science, University College London, Gower Street, London, WC1E, UK
+44 (0) 20 7679 3039
P.Monthien@cs.ucl.ac.uk and A.Sasse@cs.ucl.ac.uk

Abstract: This paper investigates how best to implement computer support for vicarious learning by taking a principled approach to selecting and combining different media to capture dialogues that take place during collaborative learning activities. The main goal is to create vicarious learning materials that contain the appropriate pedagogic and production quality, whilst minimising cost, both in terms of money and the time required to produce such materials. In particular, it proposes that multimedia and television production principles can be used to aid production of high quality vicarious learning materials for collaborative learning of diagram construction in the field of computer science. Additionally, real world problems experienced by the project are reported in this paper. This paper concludes that the specifics of how to create appropriate vicarious learning material is currently a neglected field and invites further efforts in tackling this non-trivial issue.
Keywords: vicarious learning, collaborative learning, digital video, courseware authoring, higher education.

Introduction

What is Vicarious Learning?

Vicarious learning (VL) is learning that takes place while observing learning dialogues between a student and a tutor/lecturer. In previous research on vicarious learning, the evaluation criteria were entirely pedagogical, such as "how effective are vicarious learning materials?" [5][7][11] and "do vicarious learning resources promote learning?" [4][6]. This research builds on these researches, and extends them to identifying how best to implement computer support for vicarious learning (CSVL). The main goal of our research is to create tool support for vicarious learning materials that contain the appropriate pedagogic and production quality, and can be delivered with the skills, resources and equipment available to most education institutions.

Vicarious Learning and Computer Science Education

Diagram construction is a major part in the syllabus of computer science and information technology students. Each year, students find it difficult to learn how to construct system models such as dataflow and entity-relationship diagrams. To help the students, many lecturers show common mistakes from previous years in the form of completely constructed diagrams. But since the process of constructing such a diagram involves more than the finished product can communicate, most students acquire the ability to construct such diagrams only after several attempts of their own, with feedback from the lecturer. This type of feedback – delivered 1-to-1 or in small groups – is very demanding on lecturers' time. To get more learning revenue from this time-consuming process, the tutorial sessions were captured as digital videos, and integrated into course notes and online computer-aided assessment (CAA) system. In effect, students who are not present in the tutorial can vicariously learn the same material and more through this setup. Hence, this particular case study looks at how best to produce Digital Video-based (DV-based) vicarious learning materials for teaching diagram construction in the field of computer science.

Observing Collaborative Learning Activity

Students can also acquire necessary skills by observing one of their colleagues constructing a diagram or, more often, by collaboratively constructing the diagram as a group. Additionally, if dialogues between the student(s) and the tutor/lecturer can be captured during an attempt at constructing a diagram, students who were not present will not only see the resulting mistakes but also how and why the mistake came about in the first place and of course, the reasoning behind constructing the final correct diagram as well! In such cases, the collaborative learning activities will also benefit those students who were not present during such interaction.

What makes CSVL different from just giving out a video?

The main advantages of supporting the production and access of vicarious learning material via computers as digital video (DV) over conventional video are the ability to integrate such resources with other online learning materials (such as online courseworks, lecture notes) and to create a certain degree of personalisation of the learning material. DV can contain URLs to online lecture notes or relevant online formative/summative assessment modules. Additionally, DV can be divided into segments that can be much more easily navigated than conventional analogue video.

As real life interactions will be recorded, students will be able to relate directly to problems that they and their colleagues are having during that academic year. In the simplest form, CSVL will allow students to revisit their collaborative attempts at tackling a problem. Such revisits strengthen their understanding of the topic by promoting the use of "episodic memory" [12]. Episodic memory is a person's personal memory encoded with respect to specific time and place, i.e., remembering the lecturer explaining what "episodic memory" is on a hot summer afternoon in his/her office. Such memory is more resistant to forgetting. Additionally, students will also be able to view other groups' attempts at tackling the same or similar problems by either accessing materials from the current or previous academic years.

Producing VL Materials

As mentioned earlier, the main goal of this research is to investigate how best to implement computer support for vicarious learning. In order to focus the investigation, the following assumptions have been made:

Perception of quality of material impacts upon teaching effectiveness.

Interaction with vicarious learning materials influences learning outcome.

The main design problem that is being tackled is the derivation of a design specification for a cost-effective CSVL system that can be easily implemented by individual or small groups of lecturers as well as being adaptable to their individual needs. The system needs to be adaptable enough to meet requirements for installations in small rooms (individual tutorial) and medium-size rooms (group tutorials).

Why focus on individual implementations of CSVL? Vicarious learning can be seen to work in professional productions. Examples on television include Open University education programs containing brainstorming and debates between lecturers and students and public debates often discussing issues concerning current affairs. However, a lecturer who wants to use such programs can face several obstacles. Firstly, the content of the program may not suit the need of the students/courses that the lecturer is teaching. Secondly, the materials are not grounded or linked to the students' own experience and thus does not make use of the benefits of episodic memory. Such professional productions are also costly and require large amount of time and effort from the lecturer if he or she wants to produce similar material him/herself. For this reason, this research is asking whether computer support can bring the benefits of vicarious learning to individual lecturers. If so, how can it be done? If not, what are the main obstacles?

Pedagogic Design

Within the context of pedagogic design, vicarious learning can form parts of various pedagogic designs. However, as the main aim of this project is to investigate the various aspects of technology used for implementing vicarious learning, the Laurillard's model of learning [2] was adopted and maintained through out for consistency. In this case, vicarious learning forms part of the on going negotiation of meaning between the lecturer and the students. A simplified diagram of this model is given below.

Figure 1: Simplified diagram of Laurillard's Model of Learning [9]

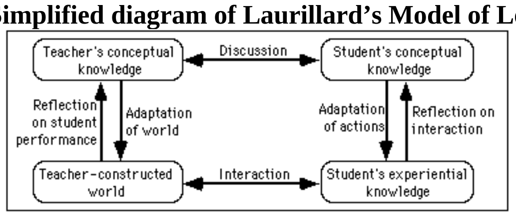

Another model of learning which may also be applicable is Wenger's Learning architecture model [16], where students are working towards being part of a community of practice, e.g., community of systems analysts.

Why perceived quality of production matters – a pilot study

Most video-based vicarious learning systems consist of video recordings of the interactions between lecturers and students, accompanied by notes on the topic under discussions, and perhaps the page of the lecture notes on the topic. Two such systems are The Dissemination System used at Heriot-Watt University [19] and a similar one at University of Edinburgh [19]. Both systems are used for teaching Human-Computer Interaction (HCI) courses. On the other end of the spectrum, the Classroom2000 project at Georgia Institute of Technology [17] demonstrates what high-quality recording and dissemination technology can deliver. By combining video/audio recording with time-stamped slides and annotation, the Classroom2000 project has managed to put the essence of a lecture online. However, this emphasis of putting lectures online means that Classroom2000 does not focus on student actions and feedback, which is essential for vicarious learning. Additionally, the system is not cost-effective to implement in individual lecturer's room or small tutorial rooms. Classroom2000 requires expensive equipment such as electronic whiteboards and backup equipment all of which contribute to the total price tag of around $100,000 (according to Classroom2000 online promotional video).

So, how would a simple production of a single-camera recording a tutorial session be perceived? As a pilot study for this project, computer science students were shown some samples of videos containing vicarious learning resources and were asked to make general comments about them. In each video, a lecturer and a student have a staged dialogue on general HCI topics while seated around a table. At various points, either the lecturer or the student would refer to materials on pieces of paper or the whiteboard. The video was produced using a single static video camera and, after encoding into RealMedia streaming format, shown on a 1-inch in diagonal window supported by mono sound. Some of the most prominent comments from the students concerning such a production were:

> Poor sound quality made it very difficult to understand what each person was saying.
>
> The video was so small and blurry that there was "little point in it being there".
>
> It was impossible to see what was written on the whiteboard or the pieces of paper on the tabletop.
>
> The content was perceived as "deadly dull."
>
> It was hard work going through each video and picking out what the main message was.
>
> Students said they would be reluctant to refer to such material, they would prefer to use lecture notes and books.

What the pilot study revealed was that the perceived quality of the material not only affects comprehension of the content being offered by the resource, but also its perceived practicality, i.e., whether such resource will be used or not. A previous study of real-time multimedia in a distance education context [13] found that perceived quality of materials can be affected by unsatisfactory audio quality resulting from incorrect setup of end-system audio hardware and equipment setup, rather than by network effects such as packet loss and jitter. Comments from students in this pilot study showed that perceived quality of recorded material can be similarly affected by insufficient end-system video hardware and incorrect equipment setup.

How to make a difference?

So what affects the perceived quality of video-based vicarious learning materials? This case study looks at the student's perception of the quality of materials in comparison with other materials such as lecture handouts, books and self-made notes. Audio/visual production variables play a crucial role in the perceived quality of the final material, We wanted to determine which variables affect students' perception of the quality of recorded instructional material, and to what extent.

Multimedia Production Principles

Reeves [10] identified a wealth of factors that can influence the perception of multimedia material. The following factors are, in our view, particularly relevant to authoring video-based multimedia:

> **Image size:** larger images are evaluated more positively than smaller ones and better capture the viewer's attention. Additionally, content on larger images are better remembered than those on smaller ones.
>
> **Fidelity:** sound quality was found to have more impact on the viewer's perception of the quality of the material being shown. More importantly, audio fidelity was found to affect attention to media and the viewer's memory for audio information.

Synchrony: viewers of videos containing audio-video asynchrony commented that people in the video were "less interesting, less pleasant, less influential, more agitated, and less successful in their delivery" than those in synchronized video. This is very significant, as educational materials should not have any of the attributes that have just been mentioned.

Scene Changes: scene changes can be used to provide structural cues to direct the viewer's attention as well as serving as markers in the viewer's memory. So it is possible through the use of cuts to create structure within the production as well as keeping the viewer's attention at the appropriate level. In order to create appropriate cuts and scene changes without using cameramen, it will be necessary to use more than one camera to capture the discussion.

Television Production Principles

The field of communication arts (or media studies) suggests that the number of camera angles also affect the viewer's perception of the quality of the production. The two main factors that have not already been mentioned are:

Shot Selection

Multimedia principles already point out that, if the production will be viewed in small windows, the camera should close in on the speakers to create the sensation of larger size. Television production principles add to this several points concerning the placement of cameras and how speakers should appear on screen (shot selection). These ranges from wide-angle shots used for orientation of the viewer to over-shoulder shots for suggesting relationship between speakers. Overall there are around seven main shots that can be used during an interview each with specific purpose [1].

Specifically in the context of this case study, it is also important to make sure that the camera shots that are used allow a clear view of any artifacts – such as models or drawings.

Number of Cameras

As well as the shot selection, the number of cameras also influences the quality and style of production [15]. Single camera productions require a camera operator that can skillfully move around the set in order to film each speaker as he/she speaks without any obvious interruptions. In a two-camera production, camera movements are greatly reduced by "cross-shooting" (the field of vision of each camera appear to cross). Additionally, the camera operator can zoom in or out to get extra "over-the-shoulder" shots.

However, if a third camera is introduced, the job of the camera operator becomes minimal or nonexistent. That is, it is possible to create all the necessary shots with three cameras for an interview situation, for instance, without needing to zoom in or out. This is important in relation to the practicality of creating vicarious learning materials. If no camera operators are required, materials can be created without the camera crew associated with traditional video productions

Post-production Issues

In addition to multimedia and television production variables, there are some post-production design issues that may also affect the student's perception of quality. Compared with text, video and audio are slow media for transmitting information – reading a book yourself is much quicker than having it read to you. It is therefore necessary to provide navigational support to allow students to quickly access certain part of the material, without having to go through the entire content of the material.

Too Many Cameras?

Experiment 1 – Finding the right balance

According to television production principles, increasing the number of cameras will increase quality of production. *But is there a point of diminishing return?* In terms of return on investment, is the amount of information captured that much better when using three cameras than one?

An experiment was conducted to investigate how increasing the number of camera angles in video productions affects the students' perception of the quality of the production. Three cameras were used to capture 3rd year undergraduate communication arts students producing their first television program. Lecturers in this discipline usually introduce students to the theories of producing a "show" in the classroom, but it is difficult to teach students how things actually work or relate to each other in the studio. The aim of the video was to show fellow students the dynamics and the dialogues that go on within the control room – problems that other students encounter, how they solved them and how they make decisions within the group to accept certain takes. The point of the video is to minimize the amount of time students spend inside the studio learning how to produce a show rather than actually producing one. As studio time is rare and expensive, this will allow students to make better use of their studio time.

The production session takes place in two adjacent rooms – the control room and the studio. The producer, director, soundman, switcher and computer graphics supervisor work in the control room while actors, cameramen and stage manager work in the studio. It is the action in the control room that are of direct pedagogic relevance for this production, but the actions in the studio are also relevant in the wider pedagogic context.

Experimental Setup

As mentioned earlier, three cameras were setup to capture the student's production session. The aim is to create for the experiment *three productions that differ only in the number of camera angles that are used*. Consequently, it was necessary to produce the three productions from just a single take to ensure that the content of the productions are identical. Each camera was positioned as follows:

Camera 1: Wide angle shot inside the control room with each person and control equipments visible. Each person has his/her back towards the camera as they are working on the equipments within the room. Additionally, a small portion of the studio room was also visible.

Camera 2: Wide angle shot inside the studio itself. This camera captures the actors, cameramen and stage manager.

Camera 3: Close up shot of each person inside the control room. Each person face, in contrast to camera 1's angle, is now visible.

Figure 2: Screenshots showing the three camera angles that were used in the experiment (including the PIP insert)

In addition to these three camera angles, it was possible to add another angle, which was drawn from the actual production that the students were creating – i.e., the show that they were making – this was presented as a picture-in-picture (PIP) insert. This is similar to being able to see what people are writing while they have a discussion around a whiteboard. From these three cameras and the PIP insert from the student's production, the following productions were created:

Production	Camera 1	Camera 2	Camera 3	PIP
1	X			
2	X			X
3	X	X		X
4	X	X	X	X

The productions are ordered according to the amount of effort it would take a lecturer to edit each production, where production 1 requires the least amount of effort.

How increasing the number of cameras affect perceived quality

Twenty-one students were grouped into three focus groups. Each group viewed the following productions:

Focus Group	Prod. 1	Prod. 2	Prod. 3	Prod. 4
1	X	X		
2		X	X	
3			X	X

Students were not directly prompted to comment about the quality of the video but instead were first asked whether there were any differences between the two video productions.

Results Summary
The transcripts of each focus group were analysed in terms of comments that the students made specifically about the quality of each production. There were significant differences in terms of perceived quality between production 1 and production 2 and likewise between production 2 and 3. In both cases, increasing the number of camera angles resulted in the students thinking that the quality of the production has been improved. However, there were no perceived improvements in quality when moving from production 3 to production 4. As a matter of fact, focus group 3 actually said that production 4 was 'more boring than production 3'. This indicates that *while increasing the number of cameras does not always increase the perceived quality of the production, the number of cameras used does affect the viewer's perception of production quality*. The result suggests that the optimum setup for this particular case would be to use two cameras to capture the dialogue and one PIP insert to keep track of what the students are doing with props.

Results Analysis
Why did the addition of another camera in production 4 not increase the perceived quality of the production? When breaking down the content of the video in terms of the Laurillard's model of learning, a pattern has emerged. If the content of each production is categorised in terms of the stages of interactions in the model, each additional camera angle in successive productions adds another stage of interaction of the model to the production – all except in the case of production 4. Production 2 adds the result of the student action, production 3 adds the actual action of the students in the studio, but production 4, while adding another camera angle, does not add a stage of interaction that has not already been covered by previous productions.

Consequently, in order to determine how many cameras will be needed to effectively capture a situation for use as vicarious learning material, it is necessary to *first identify what actions correspond to which part of the Laurillard model and then implement the necessary technology to capture those actions*. For example, a 1-camera setup will be as pedagogically effective as a 2-camera setup when capturing a student and a lecturer talking about a concept, but a 2-camera setup will be more pedagogically effective should any written materials be introduced into the discussion.

CAPTURING DIALOGUES IN THE REAL WORLD

This section reports on a case study that investigates real-world usability problems that lecturers face when producing video-based educational multimedia.

Experiment 2 – Producing video-based educational multimedia

What kind of problems do lecturers really encounter when creating their own video-based materials? In order to investigate this issue, a case study was conducted. Rather than start from scratch, the system used was based on current knowledge of what is required to produce good quality video-based materials. A lecturer's room has been setup with the following equipments so that tutorials can be captured:

- 3 digital video cameras
- 1 MIMIO [20] electronic whiteboard
- 1 Intel-based and 1 PowerPC-based PCs configured for DV-editing and whiteboard content capture
- 1 Intel-based PC configured as a streaming video server
- 1 high quality omni-directional microphone
- 1 document scanner

So far, the total cost of the system is under $10,000. While it has less functionality than Classroom2000, it is only one tenth of the cost.

Video configuration
The three video cameras were set up inside a lecturer's office to create two overlapping two-camera productions. This is required because it would allow the students and the lecturer to sit more freely as well as allowing the option of having one camera capturing solely the interactions around the

whiteboard. This enables productions with the following camera angles (where camera 2 and 3 are interchangeable) without the need of cameramen:

Camera 1: Wide Angle shot

Camera 2: Over shoulder shot of lecturer (student's view)

Camera 3: Medium close-up shot of student (lecturer's view)

Figure 3: Camera Setup

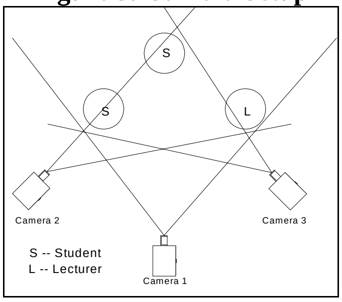

From this setup, it is possible to produce not only a two-camera production, but also one and three camera productions by omitting or adding a camera, respectively.

Problems with Capturing Real Educational Dialogues

Capturing real (as opposed to scripted) educational dialogues, presents significant problems. The problems can be broadly categorised into those relating to the pedagogic content of the captured material and the task of video-editing.

Pedagogic content

In a small-group discussion (1 lecturer and 2 students), the written materials referred to are often on paper, not on whiteboard as in the materials in the Dissemination System [18]. Students bring their work on paper, and discussions and amendments centre around the pieces of paper on the desktop. Desktop and surrounding cameras did not get a good angle on the content of paper – even at close range both physically and through zooming. It was very difficult or impossible to make out what is being written on the pieces of paper in sufficient detail. Why is the whiteboard not used in this instance? Another UCL project, ReLaTe [21], found that the use of the whiteboard is in the "tutor's domain" [3][14]. Students are reluctant to "expose" their work on the whiteboard and prefer the familiarity of their own paper. For students, writing on the whiteboard means publishing their ideas, which they must be prepared to defend.

Additionally, unscripted dialogues take time to get going. There are often periods of silence when discussion materials, such as problem sheets or student's work, are being read or inspected by the tutor or other students. Trivial or off-topic chats also take place. These do not have any obvious pedagogic value and will require time to edit them out of the final production.

Video editing

In order to produce good quality material, captured tutorial sessions need to be edited. Lecturers face significant problems here. Firstly, it is not possible to capture from multiple cameras at full resolution at the same time. This is a significant problem if more than one camera is being used, since any additional camera will require more time in post-production to transfer to computer. For example, for a 3-camera production, one hour of video would require one hour of live capturing and two hours post-production capture.

However, the most time consuming task in the process of creating CSVL material is the editing of each captured session. Off-line editing of captured dialogue is extremely time-consuming. TV studios do real time (live) editing, but an educational institution cannot afford dedicated staff to support each captured session. Additionally, each camera's microphone was picking up the audio from all the speakers making it impossible to visually detect from the visual representatives of the audio tracks who was speaking and when, so that the choice of which camera should be used could be made.

Consequently, it was necessary to first listen to a five or six seconds segment of a clip in order to determine which camera angle to use, then repeat this process until the end of the clip is reached. This "stop-go" approach is very time-consuming and is not practical for a lecturer to do at all.

Video editing for a single camera production is less time consuming. It only involves identifying which part of the captured video may be pedagogically useful, and editing out parts that are irrelevant. However, as mentioned in section 3.2.2, it is desirable to use more than one camera in a production. This creates a serious practical problem, as adding extra cameras means more post-session capture and more complex editing tasks. With two cameras, the lecturer not only has to choose which part of the captured session to use but also which version or camera angle to use. While choosing when to use a certain camera angle may be as simple as switching to *camera X* while *person A* speaks, this is very time-consuming when conducted in post-production. In television studios, this task is often performed in real time with the producer or director making the decision to "switch" cameras while the production continues to roll. This "studio" style of creating the final production would be much easier and less time-consuming for the lecturer.

Implication for Design

This section presents the conclusions derived from the pilot study and the two case studies in terms of what is now required from the multimedia community to provide sufficient support for lecturers producing video-based educational multimedia.

The Ideal Production

So, what constitute the "ideal" video-based vicarious learning material? Together with the previous findings from the literature and the case studies conducted, the following attributes have be derived:

1. The size of the video should be at least 360x288 pixels (half-screen DV) in a 1024x768 screen resolution.
2. The audio should be captured at the highest possible sampling rate and bit depth – 16-bits stereo audio sampled at 48kHz.
3. Maintain audio-video synchrony.

However, these attributes are often dictated by the bandwidth that is available while viewing the material online and therefore will be more of a design constraint rather than a design feature. In contrast, the following features can be incorporated irrespective of bandwidth size:

4. Cuts and scene changes can be controlled and incorporated by design into productions.
5. Cuts or scene changes should be used to create structure within the production as well as keeping the viewer's attention at the appropriate level. In order to create such cuts or scene changes, more than one camera will be required and each camera will need to be placed in the appropriate position to obtain the best possible angle
6. The video must incorporate camera angles that give the viewer sufficient access to props (e.g., what is written on the whiteboard, notes made on paper, etc…)
7. Allow the viewer to navigate through the video with respect to a clear pedagogic framework.

Derived system requirements

A system that efficiently supports production of video-based educational multimedia must enable the author to create materials with the greatest degree of quality while minimizing the amount of time and effort required to create those materials, i.e., it must be practical to use.
A survey of the current range of video-editing suites reveals that there are no packages that efficiently support, in terms of time and effort required, the task of authoring in-house video-based educational multimedia. Apple's iMovie is very easy to use but does not support the ability to have more than one synchronized camera angle to choose from at any one moment.

While Adobe Premiere and Apple's Final Cut Pro support multiple camera angles, they do not allow the author to choose camera angles "studio" style in post-production. Both packages force the user to use the "stop-go" approach to video editing. A system is required that can:

1. simultaneously capture DV footage from more than one camera
2. play each captured camera angle (video file) at the same time (once they've all been downloaded and synchronized) so that the author can see exactly what each camera has captured
3. allow the author to choose which camera to use in real-time without having to stop the videos being played (as if the production is being made live in a studio). This task should be as simple as clicking on the camera angle that is wanted at that time to switch to that camera in the final production.

This will critically improve the usability of video-editing systems in three major ways:

1. The pre-editing time will be cut substantially if footages from different cameras can be downloaded in parallel.
2. The editing time will be cut substantially compared with the "stop-go" approach.
3. The editing task and user-interface will be simplified to just picking which camera is "live" or active at certain time.

Figure 4: While iMovie is easy to use, it does not support the use of multiple camera angles (video tracks)

Another feature that is desirable but not essential is an automated video switching system. This system will automatically detect when someone is speaking and choose the appropriate camera angle. Such automation would allow the lecturer to edit a multiple camera production in the same way as he/she would edit a single camera production, i.e., he/she only needs to choose which part of a captured section to use. There are already systems that can perform this task, however, all are very expensive and therefore outside the resource of most educational establishments. Additionally, such a system will take away from the lecturer the precise control of pedagogic content that is to be in the production. This side-effect is highly undesirable in the educational context [8].

Conclusions

Producing video-based educational multimedia is more than just pointing and shooting a video camera. Empirical works reported in this paper support the notion that perceived quality of production is significantly affected by the number of camera angles that are used in the production. While the current range of video-editing software packages have the features to produce video-based educational multimedia material of sufficient quality, the tasks in which they demand the user to go through in order to produce such materials are not compatible with the way the lecturer actually needs to work. To make video-editing packages usable in the educational context, the editing task needs to be both quicker and easier. As the number of camera angles in production affects the perceived quality of that production, packages need to support the use of multiple camera angles in productions without forcing the user into a “stop-go” style of interaction to edit videos – this is currently a significant usability bottleneck for users of these packages.

Future Work

There is currently no large-scale empirical data on how students interact with vicarious learning materials. More importantly, there are no data on how the interaction pattern is affected when there are pedagogic shortfalls resulting from production problems, e.g., missing out vital contents through using insufficient camera angles. These variables are independent from the usual variables investigated in video technology, such as bandwidth, resolution and synchronization. This needs to be thoroughly investigated as pedagogic effectiveness is at stake.

In conclusion, in order to effectively use video-based educational multimedia to support vicarious learning, it is not necessary to have a fully equipped television studio. Such materials can be created in-house and still be pedagogically effective. However, the education community is currently being short-changed by the video-editing suites that are available on the market today. Therefore, in conclusion, it is possible to bring the benefits of vicarious learning to individual lecturers but there are some significant problems and issues that still need to be addressed in order to implement the most effective support for such method of learning and teaching.

REFERENCES

[1] Armer A. (1990), Directing Television and Film. 2nd Edition. Wadsworth Publishing Company, California

[2] Laurillard, D. (1993). Rethinking University Teaching: A Framework for the Effective Use of Educational Technology. London, Routledge.

[3] Hughes J. and Sasse M.A. Internet Multimedia Conferencing - Results from the ReLaTe Project. Proceedings of ICDE World Conference, Pennsylvania, 1997

[4] Lee, J., F. Dineen, et al. (1998). "Supporting Student Discussions: It Isn't Just Talk." *Education and Information Technology* **3**: 217-229.

[5] Lee, J., F. Dineen, et al. (1999). Vicarious Learning: cognitive and linguistic effects of observing peer dialogues. *American Educational Research Association (AERA '99) Annual Meeting.*

[6] Mayes, J. T. and C. J. H. Fowler (1999). Learning technology and usability: a framework for understanding courseware. *Interacting with Computers* **11**: 485-97.

[7] McKendree, J., K. Stenning, et al. (1998). Why Observing A Dialogue May Benefit Learning: The Vicarious Learner. *Journal of Computer Assisted Learning* **14**(2): 110-119.

[8] R. Monthienvichienchai, M. A. Sasse & R. Wheeldon. Educational metadata - friend or foe? In P. Dillenbourg, A. Eurelings & K. Hakkaranen (Eds.): Proceedings of the European Conference on Computer-Supported Collaborative Learning (Euro-CSCL2001) Maastricht, March 2001. pp. 508-519.

[9] Pennel, R. (1996). Essential Aspects of the Teaching-Learning Process. http://edtech.nepean.uws.edu.au/mmn8/laurillard.html. , Educational Technology Centre, UWS Nepean. Accessed on 24/06/2000

[10] Reeves, B. and Nass, C. (1996). The Media Equation: how people treat computers, television, and new media like real people and places. New York: CSLI Publications.

[11] Stenning, K., J. McKendree, et al. (1999) Vicarious Learning from Educational Dialogue. *Proceedings of CSCL1999* (Palo Alto, California, December 1999), Lawrence Erlbaum Associates, Mahwah, NJ.

[12] Tulving, E. (1983). The Elements of Episodic Memory. NewYork: Oxford University Press.

[13] Watson, A. and Sasse M.A. The Good, the Bad, and the Muffled: the Impact of Different Degradations on Internet Speech in Proceedings of ACM Multimedia 2000, Los Angeles, California, October 30 – November 4 2000, 269-276

[14] Watson A. and Sasse M.A. Assessing the Usability and Effectiveness of a Remote Language Teaching System in Proceedings of ED-MEDIA 96 - World Conference on Educational Multimedia and Hypermedia, June 17-22 1996, Boston, Mass. pp 685-690. AACE.

[15] Watts, H., 1997, On Camera: Essential Know-how for programme-makers. Lontec, Borehamwood, Herts.

[16] Wenger E and Snyder. WM. (2000). Communities of Practice: The Organizational Frontier. Harvard Business Review, Jan-Feb, pp 139-145

Websites

[17] **Classroom2000 Project** at http://www.cc.gatech.edu/fce/eclass/index.html (Georgia Institute of Technology)

[18] **The Dissemination System** at http://www.hcrc.ed.ac.uk/Vicar/TT/

The Computers in Teaching and Learning course was a module of an MSc in Human Computer Interaction at Heriot-Watt University.

[19] **Human Communications** 1h: at http://www.cogsci.ed.ac.uk/school/study/ug/hc1h/

A first year undergraduate course taught at HCRC (University of Edinburgh).

[20] **MIMIO Whiteboard** at http://www.mimio.com/

[21] **ReLaTe**
http://www.ex.ac.uk/pallas/relate/
ReLaTe is developing and testing video conferencing software for use in language teaching. It is a joint project between the University of Exeter and University College London (UCL). Originally funded by BT as part of the BT/JISC SuperJANET initiative, it has now received further funding from the JISC/JTAP programme