

10 Conclusions

In this final chapter of the thesis, the work of the previous chapters is drawn together, discussed and reflected upon, in order to examine what the research as a whole has discovered, to highlight insights into the field of game-based learning in Higher Education and to consider what has been learned and how it might be applied.

The first section of this chapter revisits the research questions that were introduced in Chapter 1, and considers to what extent each has been answered by the work undertaken. Section 10.2 provides a discussion of the contribution to knowledge that is made by this work, a critique of the research methods used and consideration of ways in which the design of the research could have been improved, and provides a discussion of the results from the various studies detailed in this thesis. The chapter finishes by considering possible future directions for research in the field of collaborative game-based learning in Higher Education.

10.1 Overview of the research questions

Four research questions were posed at the start of this thesis, relating to whether there is a compelling rationale for using games for learning in Higher Education, the best way to design such educational games, how to evaluate the educational effects of computer game-based learning, and a consideration of whether different types of game-based learning lead to different levels of learning or engagement. In the sub-sections that follow, each of these questions is discussed in relation to the findings of this thesis.

10.1.1 What is the rationale for using computer games?

The first research question considered if there was a rationale for using computer game-based learning in Higher Education. From the literature review undertaken in Chapter 2 there emerged a number of sound educational reasons for using certain types of games, based on constructivist theories of learning. If games are experiential, active, problem-based and collaborative then they have the potential to be valid environments for learning, not specifically because they are games but because they exhibit the characteristics of constructivist learning

environments. However, this thesis has provided no evidence to support the common assumption, found in the literature on educational gaming, that games are a valid way to teach because they are intrinsically motivating, or that using games as a way to motivate students to learn, without even knowing that they have done so, is appropriate for adult learners.

There is also evidence, from the background studies described in Chapter 3, that there is no link between an individual's propensity to play games recreationally and that same individual's motivation to play games for learning. In fact, what has emerged is evidence that games may be engaging and support learning, but only if they are well designed and pedagogically sound, with the learning outcomes closely aligned with the goals of the game, and not simply because they are games. There is also evidence from these background studies that, although we cannot assume that games are, in themselves, motivating for everyone, this is not necessarily a reason not to use them. Adult learners, in particular, want to learn something in the most efficient manner, and if a game is perceived as being the most effective way to learn something then that in itself will be motivational.

To summarise, there is a strong rationale for games being a viable and appropriate educational medium, not for the motivational reasons commonly cited in the literature but because well-designed computer games share many of the characteristics and learning goals of well-designed educational activities.

10.1.2 Designing games to be usable and enhance learning

The second research question addressed the best ways in which computer games could be designed to support learning and also to increase usability of the gaming interface and functionality. Chapter 5 presented and discussed a selection of the literature on the design of educational multimedia interfaces and the design of computer-based learning activities. The later part of the chapter describes an analysis that was undertaken to synthesise this research with first-hand evaluations of 16 online games. These were used to produce two sets of six criteria that can be used to support the design and development of computer game-based learning, one relating to the interface design and one relating to the educational design.

The educational design criteria for game-based learning applications, which were highlighted in Chapter 5, focus on the degree to which an application supports active learning and engenders engagement, and whether it is appropriate for the intended teaching situation, encourages reflection, and provides an equitable experience for all users, as well as providing ongoing support and help. The interface design criteria, also described in that chapter, examine the level of flexible interaction provided by the interface, the support for player community and transparent navigation features, and whether there is an appropriate level of user control, system robustness, and quality of visual design. These criteria have the potential to be applied to inform the design of new game-based learning activities and also as a tool with which to evaluate the possible worth of existing games in specific learning situations.

Chapters 6 and 7 describe the design and development process of two examples of computer game-based learning applications with identical learning outcomes: the Time Capsule and the Pharaoh's Tomb; both were designed with the previously identified criteria in mind and evaluated against them. The Time Capsule is a collaborative activity, based on an existing face-to-face classroom activity, that exhibited many features of game-based learning but did not provide an environment that could be explored, use scored outcomes or support competition. The Pharaoh's Tomb is a multi-player adventure game set in a virtual environment, within which the players can navigate, explore and interact with objects. The evaluation of both of these activities against the educational design and interface design criteria showed that they implemented good practice in design, as far as was feasible within the pragmatic constraints of the project.

10.1.3 Measuring the educational effectiveness of games

The third research question focused on ways in which the educational effectiveness of games can be measured and, in particular, ways that were appropriate to measure the effectiveness of the Time Capsule and the Pharaoh's Tomb. Chapter 8 examined different ways in which learning can be measured and provided reasons why established methods (e.g. the pre-test/post-test model) would not be appropriate for the games developed as part

of this work. The development of a questionnaire to test students self-evaluation of their learning with the examples of game-based learning was described.

Chapter 8 then proposes engagement as an alternative indicator of educational effectiveness and describes the rigorous development and testing of a questionnaire to measure post-experiential engagement. From the literature on engagement, coupled with factors emerging from the interviews described in Chapter 4, five factors were drawn out that were hypothesised to contribute to overall engagement with an activity; these are an appropriate challenge (consisting of motivation to attempt the task, clarity of what is required and a perception that the task is achievable), control over the environment, immersion and absorption in the activity, intrinsic interest in the activity and subject matter, and a perception of the activity having a purpose and not being pointless or a waste of time. An 18-question Likert evaluation questionnaire, based on eliciting engagement as a combination of these five factors, was developed and tested.

10.1.4 Comparing different examples of game-based learning

The fourth, and final, research question considered how differences in game design would affect the educational experience of the students using the game. To address this question, the final area of research involved a series of three comparative experiments, comprised of two small pilot studies and one larger final study. The implementation of these three experimental studies is summarised in Table 10-1.

	Pilot Study 1	Pilot Study 2	Final Study
Location	Napier University	City University of Hong Kong	Napier University
Date	March 2006	September 2006	November 2006
Time Capsule (<i>n</i>)	7	7	36
Pharaoh's Tomb (<i>n</i>)	8	12	42
Total participants (<i>n</i>)	15	19	78

Table 10-1: Summary of the implementation of the comparative experiments

In each study, each student was allocated to one of two conditions, using either the Time Capsule or Pharaoh's Tomb activity. After completing the activity, levels of self-reported learning and engagement were measured using two

questionnaires. The design, implementation and results of these experiments were described in detail in Chapter 9.

There was little evidence for a significant difference in either learning or engagement between the two experimental conditions (i.e. the use of the two different games), although there was some indication that the Pharaoh's Tomb increased self-reported learning in some areas. However, since these findings were relatively inconclusive, it would be difficult to draw any definite inferences from these data. There was, however, a significant difference in the levels of control participants perceived themselves to have using each activity, with students having a greater sense of control when using the Time Capsule. This provides evidence that computer game-based learning materials need not exhibit all the characteristics of games to be just as engaging and appropriate for learning as games that do, and, in fact, a more straightforward application may lead to greater feelings of control during the game.

10.2 Contribution to knowledge

This section highlights and discusses the significant original contributions to knowledge that have been presented in this thesis. These have been in two areas. First, the research has provided insights into and a greater understanding of the potential use of game-based learning in Higher Education. It has provided a pedagogic rationale for employing games in certain teaching and learning situations, increased understanding of the characteristics that affect motivation and engagement with adult learners, and provided evidence of differences in educational effectiveness from different types of collaborative computer game-based learning applications. Second, it has produced practical tools to understand and evaluate game-based learning: a framework for defining different examples of game-based learning through their characteristics; two sets of criteria to support the design, development and evaluation of engaging and usable educational games; and a robust and rigorously tested questionnaire to measure post-experiential engagement. The next two sub-sections of this section provide a critique of the research methods employed and a discussion of the implications of the results.

10.2.1 Increased understanding of the nature of game-based learning

This thesis has provided an increased understanding of game-based learning in Higher Education by presenting a sound and explicit pedagogic rationale for the use of game-based learning; by increasing the understanding of motivation for game playing both for leisure and learning; and by showing that game-based learning applications do not need to be set in virtual worlds that can be explored to be engaging and appropriate for learning. Indeed, the virtual world may even be a learning environment that decreases the learners' sense of control and disadvantages students who have difficulty navigating the environment. This work argues that while games do have educational potential in Higher Education, they are not the solution to educating future generations that they are sometimes claimed to be (e.g. Prensky, 2001; Oblinger, 2004).

Academics in the field of educational games research often rely on the fact that games are perceived as intrinsically motivational for most people as a rationale for using them. This thesis asserts that this assumption is simply not true, particularly in the case of adult learners in Higher Education, and provides evidence that even in a population of computing students fewer than two-thirds were motivated to learn with games and, more crucially, there is no link between being motivated to play games recreationally and a motivation to play games for learning. This thesis has highlighted that adults will use the way to learn that they perceive to be most effective (Knowles, 1998) and this does not preclude games – but they must be seen as an appropriate way to learn and not simply as a gimmick or motivational tool.

Despite showing that the accepted rationale for using game-based learning is unsubstantiated, this thesis does not argue that game-based learning is inappropriate for learning in Higher Education, but instead presents a rationale based on the characteristics of certain types of game – adventure games, simulations, strategy games and role-playing games in particular – and constructivist learning environments (Wilson, 1996) that support experiential (Kolb, 1984), problem-based (Boud & Feletti, 1991), and collaborative (McConnell, 2000; Wenger, 2000) learning. This rationale provides a context for game-based learning within learning theory and highlights the importance of using games that are designed for learning and not simply because they are

games – close alignment of the gaming objectives and the learning outcomes is essential to ensure that engagement with the game leads to engagement with learning. The work described here has also added to knowledge regarding game playing and motivation, presenting a number of motivations for recreational game playing in adults. An appreciation that a range of motivations exist and an understanding of these different motivations will better enable the creation of games for both education and entertainment that meet the needs of a range of different types of game players.

As well as providing insights into the nature of game-playing for learning and student motivation, this thesis also contributes evidence as to the relative merits of different types of game-based learning for educational use. In the comparative study described in Chapter 9, the Pharaoh's Tomb, a competitive and immersive adventure game, was compared with the Time Capsule, a computer-based implementation of a more traditional activity for teaching group skills. The experiment showed that there was very little difference in learning or engagement between the groups of students that had taken part in each of the activities, but, those who had used the simpler interface of the Time Capsule had a significantly greater perception of control than those who used the more complex virtual environment. This shows that when game-based learning is considered it is worth remembering that it is not necessarily the activities that exhibit more of the characteristics of games that are most engaging or suitable for learning. When considering the use of game-based learning it is worth remembering that a simpler environment or less game-like activity may be as effective pedagogically as an activity that is completely game-like. It is more important that any game-based learning application is thoughtfully designed to meet the needs of the learners in a way that integrates appropriately with the curriculum, teaching situation and intended learning outcomes.

10.2.2 Tools to evaluate game-based learning

As well as contributing to the theoretical knowledge on game-based learning in Higher Education, this thesis has provided three practical tools for evaluating game-based learning. These are: a framework for classifying game-based learning activities; two sets of objective criteria to evaluate game-based learning, as well as to inform the design and development of activities; and a

tested questionnaire for evaluating post-experiential engagement with an activity.

Chapter 2 of this thesis provides a number of different definitions of games used by academics, developers and researchers in the field of gaming; these are amalgamated to present an inclusive definition of game-based learning and a single set of characteristics onto which all activities of this type can be mapped (see Table 2-1). This framework removes the somewhat artificial boundaries between different types of active and game-like learning activities and allows them all to exist on a continuum depending on their characteristics, thus presenting an open yet flexible definition of game-based learning.

This thesis also builds on and synthesises existing guidelines in multimedia design, interface design and the design of constructivist learning environments to present two sets of objective criteria specifically designed to support the design, development and evaluation of computer game-based learning activities in Higher Education. The first set provides criteria for the effective educational design of game-based learning on sound pedagogic principles, and the second provides a means to evaluate the design of game-based learning interfaces to ensure their applicability for supporting learning. As well as being used as an evaluative tool, these criteria can be used to inform and influence the design and development of future educational games.

The final contribution of this thesis is provided in the form of a practical tested tool for measuring engagement with a learning experience. This is based on a five-factor model of engagement, which in turn is based upon the work of Malone (1980a; 1980b), Csikszentmihalyi (1992) and the first-hand research described in Chapter 4 of this thesis. This questionnaire was rigorously tested, refined and used to evaluate engagement in the comparative experiments described in Chapter 9. It was also shown that the questionnaire as a whole and each of the factors used to measure engagement were internally consistent. This questionnaire is one of the important outputs of this work and has the potential to be used in a variety of teaching and learning situations to compare engagement under two or more conditions; it is easy to administer and straightforward to analyse.

In all, the research presented here and three tools described make a significant contribution to furthering the understanding of the field of computer game-based learning and to providing robust conclusions and outputs that can be confidently taken forward and applied in future research.

10.2.3 Critique of research methods

In this section the research methods used in this work are critiqued and issues that arose, associated with the research design, and lessons learned are discussed.

One major issue with the work presented here, particularly coming from the educational perspective that it does, is the focus on quantitative data collection methods rather than qualitative ones. In the field of computer game-based learning much of the research is anecdotal or based on qualitative studies that lack rigour, so it was important that this research was based on a strong evidential foundation, using a range of methods and a relatively large population. Qualitative research is best suited to exploration of an issue, and generation of topics or ideas, and was used at two points in this research: initially to explore the range of issues surrounding game playing and learning, and during the evaluation phase of the activity development. Triangulating the research using a mixture of methods, each to their best advantage, aimed to provide a more robust set of outcomes.

A pragmatic reason for the heavy leaning towards quantitative data collection and analysis techniques was the difficulty in recruiting volunteers to take part in the research, even with a monetary reward. While students could be asked to take part in an activity as part of their learning and could realistically be expected to spend a few minutes completing a questionnaire, it was not possible or practical to spend the large amounts of time with participants that qualitative research requires. If participation were no object, then further qualitative work to explore the perceptions of the same individuals using each of the activities would be an interesting way to extend the work and gain a deeper understanding of the preferences and attitudes of the population.

Using self-reporting of learning and engagement as indicators of educational effectiveness is not as robust a method as actually assessing learning itself in some manner. However, since this type of testing was not feasible for the trials, because the additional time was not available for testing, and there was difficulty in testing the learning outcomes without influencing subsequent learning, it is felt to be an acceptable way of measurement in the circumstances. Again, if resources and access to participants were no object, a larger intervention using pre-tests and post-tests at varying intervals would provide more robust data.

There might also be an issue with engagement in that students may experience too much engagement or be engaged in the game play rather than in the learning. For that reason it was particularly important that there is a clear and strong mapping between any educational game and the learning outcomes, so that engagement in the game would necessarily equate to engagement with the learning outcomes. While using engagement is not ideal as a single measure, it was triangulated with self-perceived learning and was felt to be a practical way to assess learning given the constraints of time and availability of participants.

A further research criticism was the lack of time given to briefing and debriefing in the comparative studies. This was because of practical and administrative constraints imposed by inflexible timetabling systems. The possibility of insufficient detailed briefing and debriefing sessions has implications for an understanding of the purpose of any experiential learning activity, which depends on reflection about key learning outcomes. Such understanding and reflection is a central feature of adult learning theory and constructivist perspectives of learning. The limited time for briefing and debriefing during these studies meant that these aspects may not have been adequately implemented, thereby limiting potential learning from the game activities. In an ideal situation, at least as much time would be given to the reflective processes around the game as to the game itself.

Another difficulty in the research design is related to the population used in the sample – computing students. There are characteristics associated with this sample, for example computer literacy, experience with computer games, typing

speeds and attitudes towards computers, which are unlikely to match those attributes in the general population. For this reason care must be taken when generalising from these results. However, one of the reasons for choosing computing students was that they were considered to be the group most likely to accept computer game-based learning and that if it would work with this group then there would at least be evidence for the viability of further research; if there was no evidence for its acceptability with this group of students then it was hypothesised that it would be less likely to be acceptable elsewhere. Also, the demographics of this group are such that they might be expected to play games and therefore be motivated to learn using games. There was, however, no evidence that this is the case in this group, and motivation is unlikely to be higher in the general population.

10.2.4 Discussion of results

In terms of providing a rationale for the use of computer game-based learning in Higher Education, this thesis has shown that there seems to be no evidence for the assumption that students will like to learn with games because they are intrinsically motivating, or even that people who are motivated to play games in their leisure time will be motivated to use them to learn. It is not appropriate to use games for their motivational power alone, without considering the appropriateness of the game for the topic under study and the particular group of students. However, there are strong arguments in favour of the use of educational games, but only if they are designed for learning, supported in the learning context and perceived by the students as the best way to learn. The right type of game, used in the right type of environment can foster collaboration, provide active and experiential learning, and be highly engaging.

This thesis argues that it is not the fact that a learning activity is a game, per se, that gives it the potential to be engaging, but the fact that it is designed as an effective educational experience, embodying constructivist learning principles. The characteristics of game-based learning applications also share many of the characteristics of effective learning environments, and a range of activities that embody the characteristics of game-based learning have the potential to be equally engaging. Game-based learning does not have to exhibit every characteristic of games to be effective and appropriate for learning.

In Chapter 2, game-based learning applications were defined as having certain characteristics, and an inclusive and open view was taken of what is defined as game-based learning. Originally it was intended to use these characteristics to distinguish between and test two activities that were seen as a 'game' and a 'non-game', but this distinction did not turn out to be a useful one, as individuals' perceptions of game-based learning differ. As well as comparing two examples of game-based learning it may have been more meaningful to test game-based activities against a more traditional teaching method, such as a lecture, and there may be a far greater differential between these two teaching mediums. However, it was felt that there was clearly more educational benefit in an active learning method, so the students taking part in the lecture would be at a disadvantage, which would be unethical.

From observations, the students taking part in the trial seemed to be generally enthused about using the applications, the Pharaoh's Tomb game or the Time Capsule activity, and this appeared to be partly because of the fact that it was something different, unlike the teaching they were used to. There is undoubtedly a novelty effect associated with introduction of any new type of learning method, and it is possible that any motivational or engagement effects would diminish if games became commonplace in education.

There was also observed to be a greater polarity of reaction among students who used the Pharaoh's Tomb game, with some finding it instantly immersive and others finding it more difficult to get started. This diversity was not apparent with the collaborative Time Capsule activity. When students found it difficult to engage it was usually associated with an inability to navigate through the gaming environment. A small number of people seemed completely unable to visualise the game as a three-dimensional space and themselves within it and therefore found it very difficult to navigate around the environment. It is important to be aware that any type of game may be inaccessible to some students, and although making the applications fully accessible was not an aim of this research, the potential of disadvantaging certain students would be an important consideration if this type of application were to be used on a wider scale.

Although games have a clear educational potential, it is still not clear whether the disadvantages of games for learning, such as the expense, the time taken to develop and learn, and the difficulties of integration into the curriculum and assessment, outweigh the advantages. For educational games to be effective it is important that they are designed as constructivist learning environments, that the gaming outcomes match the learning outcomes, and that appropriate reflective activities are designed around the game. For this reason, using commercial games, while having greater production value and leading to greater engagement, may be inappropriate for learning unless specifically designed for this purpose.

This thesis argues that certain computer game-based learning has the potential to be effective in certain situations, but that it is not the answer to reforming education. It has a place alongside a whole host of other teaching and learning methods, electronic and traditional, and should be the most effective way to learn something rather than simply being a gimmick to motivate and engage students into learning without being aware of it.

Overall, this work has provided a range of contributions to knowledge, both theoretical and practical. It is hoped that these contributions will be used to inform and underpin future work in the field, both by the researcher and by others. The final section of this chapter considers a range of potential areas for future research in the field.

10.3 Areas for future research

A number of directions for future research are suggested by the work presented here. Five areas, which are of particular interest, are discussed in the sub-sections that follow. These are: the potential for evaluating larger-scale development and implementation of game-based learning; the use of in-depth qualitative investigative techniques; additional quantitative studies; investigation into the perception of three-dimensional environments; research into novel and experimental interface designs; and the potential for the design of asynchronous games for use in virtual learning environments to support distance learners.

10.3.1 Large-scale implementation

One of the limitations of the two game-based applications developed as part of this study was that they only represented one hour's worth of learning. An obvious future direction for this research is to implement collaborative game-based learning on a larger scale, for example its use across a unit or programme built around a game-based environment, or delivered in an immersive game-like environment, which would enable a more rigorous evaluation of the effects on learning over time. In particular, such work would enable effects arising because of the novelty factors associated with game-based learning to be evaluated.

There is great potential for collaborative adventure-gaming environments to support social interaction and group work in a range of contexts and situations, with a range of students, and the development and evaluation of such environments is a natural extension of the research described in this thesis. There is also scope for evaluating a range of associated issues such as the organisational and managerial issues associated with implementing a change in teaching such as this in the Higher Education context; technical and software design issues for supporting larger-scale multi-user interaction across local and distributed networks; and teaching and learning issues such as appropriate ways to assess game-based learning courses.

10.3.2 Qualitative evaluation

A second limitation of the research described here is its focus on quantitative data collection and analysis techniques, which meant that the collection of rich and in-depth qualitative data was limited. This suggests that a second potential area for future research would be to undertake more qualitative explorative work in the field of game-based learning, for example ethnographic investigations into virtual gaming communities or a larger phenomenographic study into the ways in which people perceive computer games in Higher Education, looking at different populations of individuals to discover the factors that make games appropriate for learning in different student populations, in terms, for example, of subject, age, gender and previous experience with computer games. This type of rich qualitative research would make an attempt to better understand the nature of game-based learning from the perspectives of the range of individuals

that are involved in the phenomenon. Qualitative analysis of this sort would also enable the investigation of more sensitive topics of study relating to games and game playing such as game addiction, the link between aggression and certain types of game, and gender roles and stereotyping in gaming environments.

10.3.3 Quantitative studies

As well as a need for further qualitative work, more robust and rigorous quantitative studies are also required. Two areas in particular are highlighted here for further study: examination of evidence of the link between engagement and learning and comparative experimental designs using a control group.

Although there is an intuitive link between engagement in an activity and learning from it, there is little research evidence detailing the nature or extent of the relationship. Correlation of measures of learning and measures of engagement using the same group of learners in future studies would allow this relationship to be explored in greater detail and provide empirical evidence of the extent to which increased engagement is related to increased learning.

Use of a control group, for example students using a traditional format such as a lecture, was not used in the study described here, primarily for ethical reasons in that it was felt by the researcher that the students in a lecture condition would be at a disadvantage compared to groups learning in a more active manner. However, there is a clear need for empirical studies of this type so that traditional teaching models can be compared with new learning methods, in order to get a true picture of the educational value of innovations. The ethical issues could be countered by ensuring that studies using a control of this nature did not affect students' assessments in any respect, for example, by recruiting and paying students as an activity separate to their studies.

10.3.4 Three-dimensional environments

A small percentage of users of the Pharaoh's Tomb appeared to find it particularly difficult to navigate in a three-dimensional virtual environment. It would be possible to undertake further research in this area to examine the links between visual awareness and spatial perception and to consider ways in which three-dimensional interfaces could be adapted to support navigation and, in

particular, users with less experience of them. It would also be interesting to examine the cause and effect relationship: do some individuals not play three-dimensional games because of a lack of spatial ability or do they have less well-developed spatial abilities because they play fewer computer games of this type? If the latter, then there would be strong evidence for the acquisition of spatial skills through gaming.

10.3.5 Interaction devices

A fourth area of interest for future research would be to consider alternative methods of interaction between game interfaces and players, with a particular focus on investigating the emerging range of interface devices used on console games machines, for example dance mats, microphones, and controllers that are sensitive to movement. This would enable the consideration of the development of games for learning physical skills and provide games that would appeal to people who are motivated to play games for the physical challenge.

Future research in this area could also examine the range of ways that individuals in collaborative virtual environments could be represented, how they could be controlled by a range of interaction devices, and the methods of interaction with other people who are present in the environment. Issues of accessibility are particularly relevant here, as the use of game-based learning environments may create an inequitable experience for those learners with less experience or with physical or cognitive disabilities.

10.3.6 Asynchronous communication

The game-based learning activities created for this research were created to be played synchronously because this was thought to be more appropriate for the skills being taught but also because it was more practical to use and evaluate these games in a synchronous manner. However, this mode would not always be practical, particularly in online or distance education where students are not necessarily online at the same time. A final area for possible future research is in the development of asynchronous games, which could provide more flexibility for the students. Issues for potential study include the design of game-play for an asynchronous environment, and ways of ensuring and testing engagement over time.

There are a number of potential areas of future research in the field of game-based learning and the five here only provide some examples of work that leads directly from the work covered in this thesis. The research described here has aimed to provide insights into the nature of game-based learning and its applicability to Higher Education as well as practical tools for investigating the area. This field is one in which the amount of interest and research is fast increasing and it is hoped that the work presented here will be of use to other researchers by providing a sound theoretical and pedagogic underpinning for the use of game-based learning in Higher Education as well as practical tools for supporting and evaluating future research in the discipline.