Cultural Heritage as a Vehicle for Basic Research in Computing Science: Pasteur's Quadrant and a Use-inspired Basic Research Agenda

David Arnold

Faculty of Management and Information Sciences, University of Brighton, UK

Abstract

Donald Stokes argued [Sto97] that for 50 years from the end of the second world war to the end of the 20th century there was an unhealthy taxonomy of research types which was formulated on a linear scale from pure to applied. The argument goes that the best research is only possible in environments which are free from contemplation of the potential uses to which results might be applied. In this paper current research challenges in the application of ICTs to cultural heritage information are reviewed in order to consider where these applications-linked needs require solutions that will advance the understanding of computational principles and help to develop new basic understanding in computer science, including shape manipulation and other aspects of importance in computer graphics and virtual environments. The paper draws extensively on the recent published EPOCH research agenda [AG07] for illustrations of the types of research which are required for the Cultural Heritage sector and the relationship between these and basic research challenges in Computing Science.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Three-Dimensional Graphics and Realism]: Virtual Reality

1. Introduction

Research in interdisciplinary fields such as ICTs for Cultural Heritage can be viewed from a variety of perspectives. In some areas basic, new computing science results must be developed specifically to address challenges which are unique to cultural heritage. In other areas the best of generic computing science research must be applied to cultural heritage situations, potentially creating novel working methods for Cultural Heritage professionals. This may constitute basic research from the perspective of the Cultural Heritage professional - enabling new ways of undertaking their own research. Similarly basic computing science first developed in the context of solving cultural heritage problems may later be discovered to have applications in other fields of endeavour. Such a range of situations can be understood neither as purely basic research in ICTs nor as purely applied research in Cultural Heritage Informatics.

EPOCH promotes inter-disciplinary use-inspired basic research both to increase the understanding of basic research issues in ICTs and to develop improved technology to meet specific needs of Cultural Heritage. Use-inspired basic research is often promoted nowadays as a means of ensuring that publicly funded research is firmly based on providing solutions that have exploitation potential, yet it remains tainted in the way that national and international organisations assess the worth of scientific advance. Donald Stokes's book "Pasteur's Quadrant: Basic Science and Technological Innovation" [Sto97] provides a framework to set the different types of research in context. Stokes analyses the relationships between differently motivated types of research and his evaluation merits more detailed discussion in the computer science community.

According to Stokes [Sto97], in the United States, the notion of "Basic Research" whose purity was guaranteed by a separation from conceivable applications, was promoted by the influential federal report "Sciences, the Endless Frontier" released in July 1945 by Vannevar Bush in his role as President Roosevelt's director of the Office of Scientific Research and Development. "Applied Research" was considered to be at the other end of the spectrum and somehow inevitably close to market and inevitably distorted by the need to address market forces. Of course the perceived proximity to market also meant that public funds were inappropriate to fund research which should be funded by the "direct" beneficiaries of the resulting products' inevitable profits. Bush argued that basic research should receive most public funding because the marketplace would not fund it adequately.

Research inspired by:		Considerations of Use?: No	Considerations of Use?: Yes
Quest for fundamental understanding	Yes	Pure Basic Research (Bohr)	Use-inspired basic research (Pasteur)
	No		Pure applied Research (Edison)

Figure 1: *Stokes' Quadrant Model (after [Sto97], p73)*

Stokes argued that scientific research should not be conceptualized as a linear progress, but rather that "considerations of use" and "quest for fundamental understanding" represent different measures against which any proposed research project could be rated. As a result Stokes introduced four quadrants of research (Figure-1):

- In a first quadrant he placed "Pure basic research" which is inspired by the quest for knowledge but not by potential use. A paradigmatic example of this is the work of the physicist Niels Bohr on the structure of the atom.
- In contrast, the quadrant labelled "Pure applied research" contains research developing practical solutions and marketable products. Stokes cites Thomas Edison for work on electric lighting, sound recording and other marketable, practical innovations.
- A third quadrant contains scientific work that is neither overtly theoretical nor directed at products. This work concentrates on the exploration of particular phenomena or the development of a taxonomy or other classificatory work. Rather than advance scientific knowledge or develop market-orientated solutions, the focus is more on already well understood research problems or formalising existing knowledge or academic practices (e.g. handbooks or guidelines). We will return to consideration of this quadrant later.
- The final fourth quadrant is reserved for use-inspired basic science. This has potential practical utility, but researchers who conduct such research do not lose sight of the goal of advancing scientific understanding. The paradigmatic example here is the work of Louis Pasteur.

Stokes suggested that "Pasteur's quadrant" should receive most of the interest in national research policies and public funding, providing a combination of advancing knowledge and potential exploitation and return on investment. The notion of use-inspired research has significant implications for how scholars conceive of research which may face some tensions with current academic research cultures. In fact, if researchers concentrate on basic research they will usually do so within the confines of specific "pure" scientific disciplines that have their list of research priorities and established review and reward mechanisms. These priorities and rewards may not be defined to address practical, societal and policy-related considerations explicitly.

Typically, research which can be located in Bohr's quadrant has received the highest prestige in peer assessment of the quality of research. This might be considered an odd result, given that such exercises often purport to be based on assessment of novelty, rigour and impact, and results with demonstrable application and equal novelty and rigour might be expected to have more obvious impact. There is little doubt however that attitudes which value "purely basic" research have had a material impact on communal research priorities and hence on the behaviour and careers of professional researchers. In this context, attempts to become more "use-inspired" may be considered to be misguided, despite the obvious link between research in Pasteur's Quadrant and some degree of intended usefulness of the results.

The UK's university research assessment exercise (RAE) sets out to measure the quality of research across the UK. One researcher has reported that *"it would be hard to place 'Intelligent heritage' in the RAE and hard to persuade practitioners that they ought to be associated with it. In simple terms, anyone caught doing 'intelligent heritage', heritage policy or applied computing is likely to be sidelined or dismissed in order to enhance an institutional response to RAE."* (William Kilbride reported in [NGV06]). Hence, there is a need to consider more deeply the values that could limit the perceived value of CH ICTs as an interdisciplinary field.

Ironically there is no doubt that both national and international policies on the public funding of research have moved strongly towards support for use-inspired basic research. Managed programs of research, which in the UK Computing Science community certainly stretch back to at least the Alvey program in the 1980's, have become the norm. Public funds are also allocated this way in the EU Framework programs. The frontispiece illustration in the UK's current science policy framework [HM 07] presents Pasteur's quadrant.

In some ways this thinking is beginning to permeate the Computing Science research community. In both the UK and North America there has been a slow but significant development of the notion of "Grand Challenges" - usually expressed as some vision that cannot be addressed with current technologies or systems and requires some significant results in basic computing science before the tools that can meet the challenge can be developed. At present nine grand challenges have been identified within the UK scheme [UK 07]. In North America the Grand Challenges program has been formulated around a series of workshops on specific research themes [Com07].

However although the notion of Grand Challenges appears to encourage a view that use-inspired basic research will find a more central place in the research community, in practice very few of the challenges so far formulated are grounded in application domains. In addition many of the managed programs of research struggle to encompass inter-disciplinary viewpoints and the motivations of eventual applications can lead to unrealistic expectation of the timescales for exploitation. The norm in European Union Framework programs has been for a report on "Dissemination and Exploitation Plans" to be a deliverable within the first six months of the funded period, when not only will the basic research envisaged in a project remain to be completed, but since the whole project is required to be pre-competitive collaboration, there must of necessity be additional development after the project, before products will generate revenues. This tension is in part grounded in the need to secure political support on the basis of predictions of direct economic benefit.

A further example would be the UK e-science programme which conceived of Grid technologies to facilitate large scale scientific experimentation with distributed resources, but produced a strong reaction in some Computing Science researchers who conceived the programme as a diversion of funds away from basic, real computing science to make the subject a mere servant of the natural sciences.

In fact in any Grand Challenge - and the longer agenda for e-science is undoubtedly one - there are serious computational aspects that current computing science cannot deliver. These cover engineering solutions, both hardware and software, and computational methods (data structures, algorithms, underpinning mathematical formulations, etc). There also are inevitably complex issues of deployment in terms of impact on business processes of other professionals, usability of tools, long term sustainability, etc whose impact on speed of technology adoption is normally underestimated but ultimately determine the acceptability and success of solutions.

In the following sections we review how different research topics could be mapped to Stokes' four quadrant model, by considering elements of:

1. The EPOCH research agenda for ICTs in support of Cultural Heritage applications [AG07]
2. The UK Grand Challenge - "Bringing the Past to Life for the Citizen" [Arn06]

2. Analysis of the EPOCH research agenda

The EPOCH research agenda [AG07] is the result of extensive consultation between technologists and cultural heritage professionals, based on user requirements in the cultural heritage sector. It also includes some natural extrapolation of the direction of travel foreseen by cultural heritage professionals. Adoption of results will be heavily dependent on additional business process adaptation, which can be compared to the 20-30 years it took from the early conceptualisation of CAAD systems until ICT based tools to solve architectural design challenges were embedded in the profession. This paper considers only the technical research challenges arising in the cultural heritage sector. In addition it is confined to the challenges facing professionals working with the tangible cultural heritage of monuments, sites, and museums. After considering five scenarios for future cultural heritage activity, the ICT research implied to make the scenarios possible, was grouped into sub-areas, described below.

2.1. Novel data capture

A recent event, convened by the European Commission and involving invited experts from digital libraries and the heritage sector, agreed that there were no types of artefact for which digitisation could be regarded as a solved problem. A very wide range of challenges come under this heading, magnified by concerns for accuracy, volumes of data, variety of capture conditions and material types.

Even in the digitisation of printed works to extract text sources there were significant gaps (e.g. OCR of pre-1840 Gothic typefaces - which covers a very large percentage of Germanic literature). Digitisation of materials such as glass, fur, jewels or historic fabrics are still challenging in any form (although progress continues [MBK05]) and largely unsolved for effective techniques to capture digital representations that will sustain a range of desirable analyses. Interactions between materials and with the complex and detailed geometry of many cultural artefacts present further challenges [MVSL05]. For example the intricate geometry of an ornate and decorative piece might interact with the material properties of antique filigree gold, natural crystal, carved ivory or jade.

Apart from materials issues, there are unsolved problems due to environmental challenges. For example, the digitization of masonry fragments on an archaeological site would require measurement techniques operating in situ (i.e. without lifting fragments).

2.2. Semantically-based representations

This area includes the need to access multi-lingual sources, which may no longer be spoken and the challenges of co-referencing - referencing items (e.g. events or people) from multiple sources, usually with incomplete information, synonyms or alternate spellings. Knowledge discovery in these circumstances might be thought of as experts "digging around" in archives manually. The challenge is to see if "digging around in the digital world" is possible and useful.

Another distinctive heritage need is to represent the development of knowledge over time. Knowledge of the past is often lost and then re-discovered. Rediscovery may be an extended process with many interim hypotheses (e.g. what a building was like or used for), refined as more evidence is uncovered. Although other disciplines share some of these characteristics (e.g. medical diagnosis or scientific discovery, itself often considered part of our heritage) cultural heritage imposes additional challenges. Most obvious amongst these are that:

- where information is factual, the facts were available at some time in the past even if they are now lost and
- facts are often overlaid with differing multi-cultural interpretations based on political convictions or belief systems.

Since the themes of surviving tangible cultural heritage often involve either commemoration of war or celebration of religious beliefs, it is unsurprising that the same artefact or environment may embody completely different meanings for different people. Again elements of similar issues may arise in other applications - for example virtual environments may need to represent and highlight different aspects for different users. For example, members of the emergency and security services would be trained to look for and act on different features at a reconstructed incident than the general public would notice, but the cultural heritage sphere is probably the most intense example of this application-driven requirement.

2.3. Multipurpose information visualisation and communication

Having captured the information base and organised our knowledge of it in ways which allow the extraction of information, there is the inevitable demand to be able to visualise information and create experiences or presentations for many purposes. There are many applications requiring effective communication with widely differing groups and few of the challenges can be realistically described as "unique to cultural heritage". An example of challenges which apply to many domains would be authoring tools which generate engaging, multi-modal experiences tailored for particular groups of users.

However there are specific needs which arise from both the sector itself and the nature of the data to be presented - most obviously how the uncertainty underpinning hypotheses and the multiple perspectives on the significance are presented to an audience. These challenges are exacerbated when the audience may have unknown background, interests, beliefs and motivations - for example as internet users - again in common with many applications over the internet, but probably of particular significance in cultural heritage.

From the perspective of the graphics community one area which it may be surprising to find listed as part of "visualisation and communication" in the EPOCH research agenda is "reconstruction" - which in graphics might often be considered part of data capture (i.e. modelling to create digital artefacts) but, for most cultural heritage professionals, visualisation is the creation of interpretations from the basic evidence and hence an analysis and "output process" from the primary evidence to inform the cultural heritage professional. Part of visualisation is in fact to test hypothesis to see whether the evidence that exists can be put together in ways which fit both the hypothesis and the constraints inherent in the evidence.

This difference of perspective is another concern from the application domain - that, since "anything can be modelled" and visualised, the technology actually enhances presentations that are pure fantasy or, worse still, based on scant evidence, and produces believable but unsupported conjecture. For these reasons an important part of creating reconstructions and other derived forms is to record the provenance, including cross-reference to the evidence base and reconstruction methods used, to document the background to a reconstruction. The London Charter [Kin06], which appears to be attracting interest and support, proves some principles in this area, using the term *paradata*. Part of this challenge includes a desire to be able to display reconstructions in ways which reflect the degree of uncertainty inherent in the resulting model. As yet there are no generally accepted paradigms for this but the use of non-photorealistic rendering techniques appears an attractive option.

Identification and provenance history also needs to be maintained during repeated manipulations in the creation of derived and collected works. This must work across the typical image processing, graphics and CAD manipulations with data structures involving hierarchy, instancing, etc. and data manipulations (mesh decimation, stitching, surface extraction, compression, feature extraction, etc).

At the interfacing level the domain shares all the challenges of other interactive applications, plus additional, and as yet unsolved, issues in the area of multi-cultural and multi-lingual interfaces. The issues shared with other application domains would include design of interactive systems which engage and entertain the public (c.f. games), and how to measure effective engagement so as to enable better interface design. Finally the design of interfaces to tools to support research methods is an area in common with other areas of the humanities.

The EPOCH research agenda concludes with two other areas where the research issues are not particularly reliant on specific application area:

- Mobile, distributed, and networked systems and
- Long-term availability (Preservation, data migration etc. including standards for data, business processes and legal frameworks, Digital Rights Management, Intellectual Property, Copyright, collected and derived works etc)

2.4. Mapping the Research to Stokes' Quadrants

In this research agenda there are aspects which can be considered as lying in each of the quadrants of Stokes' diagram.

In ICT, research undertaken with a view to generating fundamental understanding could be targeted at underlying theory or indeed, in Stokes's classifications, research where the understanding of representation and computational process can be applied to solutions which are independent of specific applications. Generic technologies, as mentioned in various places above, fit this definition of Bohr's Quadrant.

However the axis of "application consideration" is not a binary choice - any subset of computational methods will address classes of problems and these classes may tend to occur more in particular application domains. For the graphics area, any computation addressing geometric data targets applications where shape or spatial organisation is fundamental. The more specific the shape considerations are the more application-specific the research might be. Many basic ICT research topics might therefore fit in Pasteur's Quadrant, including all of the research targeted at intelligent tools. The argument in favour of the value of conducting research in this quadrant is that *different* basic research is undertaken where there are specific application domains in mind.

There are also serious motivations for research in Edison's Quadrant. Often this may be targeted at applying pure basic research results (Bohr's Quadrant) or indeed more generic research from Pasteur's Quadrant at applications in Cultural Heritage. Of particular significance is the need to prove that the generic results work effectively with the data arising in significant practical situations, which would typically exhibit special cases and data volumes which often may not have been tested in the original basic research. In the same way that the degree of application consideration is not a binary division neither is the "fundamental understanding" axis, which in some ways is used as an inverse of "closeness to market". However the amount of additional research that may need to be undertaken to convert "fundamental understanding," even if application specific, into deployable products is often considerable and frequently underestimated. This is further complicated because other disciplines, with a direct interest in the application, will be engaged in interdisciplinary research which relies on having the application in mind, but are also likely to add more value as the results become closer to deployment. If basic research is genuinely "use-inspired" and takes account of full interdisciplinary understanding then the gap between basic results and effective deployment should be minimised.

Stokes states that the fourth quadrant should not be thought of as "empty" just because it is not labelled, but in his view it includes research that "systematically explores particular phenomena without having in view either general explanatory objectives or any applied use to which the result may be put." Stokes particularly mentions research into taxonomies, which fits into the background work on the semantics of cultural heritage. There are other areas where systematic classification and analysis are needed to convert prototypes into fully functioning tools.

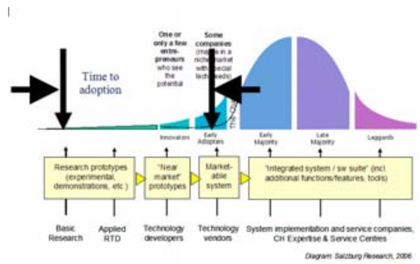

Figure 2: *Progression from basic research results to deployed solutions. The arrows indicate the extended time to adoption (after [AG07]*

For example historically there has been research into procedural and grammar-based procedural modelling of architectural forms and design styles [Yes75, Sti76, Sti77, Arn78, Sha07]. Recently research taking similar approaches has demonstrated the potential for highly efficient representation incorporating dynamic constraints and using grammar representations [MWH*06]. However the demonstrations to date have been limited in terms of the historic period, architectural styles and range of buildings modelled. Similar procedural models of architectural styles have been developed using the Graphics Modelling Language (GML) [GBHF05, HF04].

A systematic analysis of design styles is required as background to support the widespread adoption of grammar-based or other procedural modelling. These may not extend the basic computational principles (except when additional special cases are found) but they turn a prototype application (grounded on basic research) to a viable tool, extending our understanding of architectural form and possibly of the semantics of shape. Research to achieve coverage of an application domain from a computational perspective can be considered within the unnamed quadrant. Addressing the commonalities of different requirements in order to define best practice and standards also fits into this quadrant as standards are intended to be useful in many different applications contexts.

There are many data capture and representational formats which have been applied to cultural heritage artefacts and environments. Some of these allow capture of fine detail, for example in the use of scanning technologies. Others are designed to include encodings of different levels of detail, often built in explicitly via the modelling of each level and the selection and blending of individual levels within a visualisation tailored to a particular view of the environment. Some approaches assist the efficiency of displaying a visualisation by using segmentation of the environment into manageable pieces. Finally some representations allow for adaptive run-time generation of resolutions, which may use (for example) progressive meshes to generate additional detail according to knowledge of the underlying "perfect" geometry. This last approach (where it exploits knowledge of the underlying geometry) and occasionally the basis for segmentation of environments (e.g. the use of instances) may be seen to use knowledge of the elements of the environment to "work smarter" on the overall problem, but such techniques are barely scratching the surface of the potential, and indeed the need, to develop and use representations that are more fully grounded in our knowledge and understanding of the environments being represented. This paper discusses the problem, demonstrates the need and our extremely limited approaches to date in meeting these objectives, and sets directions for achieving integrated solutions to the research issues raised.

3. Inspiration for future "Use-inspired basic research"

The previous discussion has outlined many areas in which there are unsolved technical challenges. In this section a few cases grounded in examples drawn from heritage sites are used to illustrate the challenges and hopefully inspire others to take them up. For example, a brief consideration of the issues involved in digitally recording a major site in a useful way immediately exposes the limitations of current techniques. Consider for example the challenge of recording the current state and visualising the previous glory of a site such as the Temples at Karnak. Figure-3 shows the current entrance to the site.

Figure 3: *The public entrance at Karnak*

The site covers many hectares and, although much is standing, still more has been ruined. The current state includes individual features of interest on a very small scale (fractions of a centimetre for some purposes), but the needs of many visitors would probably be satisfied with much less individual detail but sophisticated reconstructions of the overall building mass, structure and decoration. At various times in the past elements of the temples have been remodelled and masonry from one state of the Temples has been re-used in later states. In addition, decorative detail, in particular representations of people from one regime, has been deliberately destroyed by elements from later regimes.

Such a site incorporates many of the challenges listed above:

1. Handling the requirements for accuracy and the data volume over the range of digitisation required.
2. Digitizing in the absence of access and portability of the material.
3. Representing the information content of the artefacts and environment that cultural heritage professionals would need to preserve and on which research about the site could be effectively based.
4. Representing multiple inter-related. For example the temples are covered in hieroglyphics which have information content of their own as well as well as surface shape in carving, material properties relevant to visualisation etc. In addition the inter-relationship between the carvings is significant and of course the carved masonry itself is structurally significant to the architecture. Finally the current state of the appearance is only one instance and the materials and their decoration have undergone numerous changes over the last 4,000 or so years.
5. Preserving in the long term and managing assets

Figure-3 shows the main tourist entrance to the site - an avenue lined with 36 statues. A sequence of pictures of the most accessible of these statues taken with a 5 Mpixel camera was processed using the EPOCH 3D webservice [EPO07] to produce a 3D model. The data for the sequence and resulting 3D model was about 100 Mbytes. Whilst the fact that such a service exists and can produce reasonable results with current technologies is remarkable, if an application required that all of the statues of this type on the site were recorded similarly then perhaps 10 GBytes of data would be needed.

There are however (inevitable) limitations to this type of model and hence the uses to which it can be put. The images involve a sample rate of the statue surface of around 3mm spacing. Clearly features smaller than 3mm are therefore hard to pick out, which limits potential analyses and might be critical for some applications. Increasing the accuracy of digitisation would add to the data volume and these statues are only a small part of the material at the site.

For example if the application envisaged is to be able to illustrate a multi-media CD-ROM then the model may well fulfil the requirements. If on the other hand the objective is to be able to use a time series of models to monitor environmental damage over time, then 3mm accuracy may well represent a significant limitation. Equally the use of such models as part of an extended scene in a real-time interactive virtual environment would not be possible currently. Much more efficient representations are required for real-time visualisation and these may or may not be derivable from the model's mesh.

Representing the single statue may in itself be interesting but is unlikely to address a complete application. Analysis of environmental damage would probably be intended to assist the site management as a whole - not just a single statue. The visualisation with an interactive walkthrough would normally also call for a model of an extended area of the site (if not a complete representation of one period of the site's development).

The challenge of representing multiple perspectives about an object to support analysis is really significant. Professional knowledge and understanding are used to inform the way artefacts and contexts are described and documented so that the information required to support research is available. However knowledge of the semantics of the domain is needed in order to design suitable representations of the scene.

Figure 4: *Carved Wall apparently showing a "spreadsheet"*

Consider the example in Figure-4, which shows information at a remarkable number of levels, carved into the wall. As a technologist the surface design has all the appearance of a spreadsheet, with cells containing hieroglyphics and apparently numeric data. However the information is incomplete at the level of the apparent original data and at the level of the individual elements of broken stone, where the breaks in some cases leave partial "cells". In the wider context of digitising the site appropriate representation of such situations should assist in making associations as other fragments are documented.

Figure-5 shows a different style of decorative carving in which the elements of story are told in a progression within the carving design. There is clearly additional information represented in the carving and the relationship between the various carvings and photo only shows a small part of the entire wall which extends much further. This single wall would represent a huge challenge of representation in its own right even though it is relatively accessible from the point of view of digitization.

Figure 5: *Part of an extended decorated wall*

4. In Conclusion

In this paper the author has sought to draw attention to the range of fundamental Computing Science challenges with origins very specifically embedded in the requirements of cultural heritage, in order to dispel some of the prejudices which have dogged the reception of the worth of use-inspired basic research.

Whilst "applied research" may be widely perceived in the UK's Research Assessment Exercise as having been held in low regard historically, there is no doubt that "Use-inspired basic research" embedded in Cultural Heritage applications present significant, fundamental research challenges in Computing Science. A uniformly poor reception of research in this area would certainly reflect shortcomings of the exercise rather than the rigour, novelty or impact of the research. The academic argument favouring Bohr's quadrant is that the basic research is not sullied by perceived or potential distortion from commercial pressures to reach particular conclusions. In addition political social or economic, rather than purely scientific, considerations could (and perhaps should) have an impact on resourcing particular lines of research. Since priorities would not have been drawn solely from scientific considerations it might be felt that this is also a potential distortion of quality, but the assessment of relative quality of results and decisions on priority topics to be resourced from limited public funds are essentially independent. They will come together when deciding which researchers and what methods are to be applied to publicly funded projects, given that there are limits on available resources for research.

Similar perceptions, or misperceptions, have been expressed about the way in which research related to applications is received in the selection of material for conference programs. Again - historically there has been some truth in that. It is hoped that the discussion in this paper will help those working in these areas and those who assess their work to differentiated between the quality of the underpinning work, the suitability for funding decisions and the determination of the relationship between underpinning concepts and application content. In some contexts the application content may be more valuable than the development of concepts, but this is more probable (e.g.) for an event where the focus is itself on application content.

In the mean time there are currently many sources of inspiration for those seeking interesting challenges on which to base their future research projects. Some of these have been drawn together in efforts to define the new EPOCH Research Agenda and will continue to be addressed as the UK Grand Challenge of Bringing the Past to Life for the Citizen is taken forward in the long term.

5. Acknowledgements

This work has been conducted as part of the EPOCH network of excellence (IST-2002-507382) within the IST (Information Society Technologies) section of the Sixth Framework Programme of the European Commission. Many fine researchers contributed to the EPOCH Research Agenda debate too many to name them all. The contributions of Guntram Geser of Saltzberg Research for drawing my attention to Stokes' work, discussions around the EPOCH Research Agenda and the production of Figure 2 are especially appreciated. My thanks too to Karina Rodriguez Echavarria for help with the production.

References

[AG07] ARNOLD D., GESER G.: *Research Agenda for the Applications of ICT to Cultural Heritage*. EPOCH Publications. ISBN 978-963-8046-80-2, 2007.

[Arn78] ARNOLD D.: Phd dissertation: A computer model of housing layout. University of Cambridge, 1978.

[Arn06] ARNOLD D.: Bringing the past to life for the citizen, grand challenge 9. http://www.cmis.brighton.ac.uk/gc9/, 2006.

[Com07] COMPUTING RESEARCH ASSOCIATION: Grand challenges. http://www.cra.org/grand.challenges/, 2007.

[EPO07] EPOCH: Epoch 3d webservice. http://homes.esat.kuleuven.be/ṽisit3d/webservice/html, 2007.

[GBHF05] GERTH B., BERNDT R., HAVEMANN S., FELLNER D.: 3d modelling for the non-expert users with the castle construction kit v0.5. In *Proceedings of VAST2005: 6th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (2005), pp. 49–58.

[HF04] HAVEMANN S., FELLNER D.: Generative parametric design of gothic window tracer. In *Proceedings of VAST2004: 5th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (2004), pp. 193–202.

[HM 07] HM TREASURY, HM DEPARTMENT FOR TRADE AND INDUSTRY, HM DEPARTMENT FOR EDUCATION AND SKILLS: Science and innovation investment framework 2004-2014. http://www.hm-treasury.gov.uk/spending_review/spend-_sr04/associated_documents/spending_sr04_science.cfm, 2007.

[Kin06] KINGS COLLEGE LONDON: The london charter. http://www.londoncharter.org, 2006.

[MBK05] MULLER G., BENDELS G. H., KLEIN R.: Rapid synchronous acquisition of geometry and appearance of cultural heritage artefacts. In *Proceedings of VAST2005: 6th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. 2005, pp. 13–20.

[MVSL05] MUDGE M., VOUTAZ J., SCHROER C., LUM M.: Reflection transformation imaging and virtual representations of coins from the hospice of the grand st. bernard. In *Proceedings of VAST2005: 6th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (2005), pp. 29–41.

[MWH*06] MUELLER P., WONKA P., HAEGLER S., ULMER A., VAN-GOOL L.: Procedural modeling of buildings. In *ACM Transactions on Graphics: Proceedings of ACM SIGGRAPH 2006* (2006).

[NGV06] NICCOLUCCI F., GESER G., VARRICCHIO T.: *Digital Applications for Tangible Cultural Heritage: Report on the State of the Union: Policies, Practices and Developments in Europe*. EPOCH Publication. ISBN 963-8046-68-6, 2006.

[Sha07] SHAPE GRAMMAR ORGANISATION: Bringing the past to life for the citizen, grand challenge 9. http://www.shapegrammar.org/biblio.html, 2007.

[Sti76] STINY G.: Two exercises in formal composition. *Environment and Planning B 3* (1976), 187–210.

[Sti77] STINY G.: Ice ray: A note on the generation of chinese lattice patterns. *Environment and Planning B 4* (1977).

[Sto97] STOKES D.: *Pasteur's Quadrant: Basic Science and Technological Innovation*. The Brookings Institution, Washington, 1997.

[UK 07] UK COMPUTING RESEARCH COMMITTEE: Grand challenges in computing research. http://www.ukcrc.org.uk/grand_challenges/index.cfm, 2007.

[Yes75] YESSIOS C.: Formal languages for site planning. In *Spatial Synthesis in Computer-Aided Building Design*, Eastman C., (Ed.). Applied Science Publishers, London, 1975.