Applications of Mixed-Initiative Interfaces and Intelligence to Electronic Portfolios *

Patrick Lougheed and **David Brokenshire** and **Brittney Bogyo** and **Mayo Jordanov** and **Vive Kumar**

School of Interactive Arts and Technology
Simon Fraser University
14th Floor, 13450 102 Ave.
Surrey, British Columbia, Canada V3T 5X3
{patl,debroken,bbogyo,mayoj,vive}@sfu.ca

Abstract

Electronic portfolios are a rapidly growing area of research, primarily in the educational and medical arenas. While existing systems are fairly simplistic, without knowledge of their domain or their users, there are possible applications for enhancing this technology to allow for better user interaction. We outline our current research interests with regards to enhance portfolios via mixed-initiative interfaces and intelligence. We also present our methods of implementing our ideas and their requirements.

Introduction and Background

Portfolios are an old concept used by many for everything from job searching to scrapbooking. A portfolio can be *anything*, from the old "memories" shoebox that many keep under their bed to store photos from childhood to far more complex undertakings. They are largely based on the idea of "reflection" - the ability to look back and reminisce, look critically at past actions and output and improve ourselves based on this critical look.

While the concept of portfolios is not a new one, a recent expansion of their scope from merely a paper-based job searching tool to a digital educational tool has greatly enhanced both the breadth and depth of the research on the subject. A number of new tools have been created, both as research platforms and for commerical use, helping users create their own portfolios. These systems are typically implemented as editors around a set of artifacts that the user has created, and their responses to them. These artifacts and responses can then be placed into a coherent structure and ultimately placed somewhere where others can view the portfolio and give feedback upon it.

Portfolios have been implemented in numerous jurisdictions - from primary and secondary school (British Columbia Ministry of Education 2004) to higher education, to corporations and federal and state governments to health care organizations (McKimm 2001) - as learning aids. They have advantages in this regard by allowing the learner to reflect on themselves and their achievements and see improvements or declines over time, as well as areas where they can improve themselves. Major corporations are beginning to require portfolios for some positions - Bell Canada, for example, will not consider an applicant for a senior-level management position without a portfolio of some sort. There are of course the traditional venues for portfolios - art houses, freelance artists and architects have long built their portfolios to add them in the search for the next contract, and these portfolios now exist more and more in either both mediums or are solely digital. There are nearly as many different structures of portfolios as there are authors of portfolios, and many different uses (Smith & Tillema 2003).

Unfortunately, one current problem with the state of the art in electronic portfolio is that of user buy-in. While many institutions and corporations have adopted the idea of electronic portfolios whole-heartedly from a management perspective, there are still a plethora of issues keeping users from using portfolios to their advantage. These reasons include a lack of communication from those providing the service with regards to the benefits of having an electronic portfolio, and a lack of the critical mass of users and readers of portfolios.

We also believe that part of the reason for this lack of uptake has to do with the current tools; while today's platforms are orders of magnitude better than those of five or ten years ago (when portfolios were created in Adobe Acrobat and Microsoft Word, applications not targetted to portfolios at all) they are still fairly bare-bones in terms of the services they can offer users. At this point nearly everything is done manually by the user, with no possibility of help from the system itself. There is also the problem of the learning curve for the user, as many of the people we're asking to adopt electronic portfolios are not comfortable with computers, let alone more advanced applications such as portfolio tools.

Our aim is to enhance our tool and attempt to overcome some of these issues with the aid of the mixed-initiative research and technologies, and address at least the technical issues via this method.

Current Platform

Our tool, SPARC (the Student Portfolio Architecture and Research Community) (Bogyo *et al.* 2004), currently in it's second major revision, and is being both in production use at Simon Fraser University for more than 400 students, and as a platform for research in various different areas. One such area is the incorporation of system intelligence to aid the user in the creation of their own portfolio(s). SPARC is implemented under the principle of 'everything is an artifact' - files, URLs, reflections on these, other people's feedback, etc. It can manage multiple portfolios for a single user, and has an artifact repository. The backend is based on SOAP and web services, allowing an infinite number of front-ends, of which we have implemented two, one web-based and one a local Java program.

Portfolio creation in SPARC (and in most other systems) generally happens cyclically, in about three different phases. The first phase is artifact cataloguing, where the user (and possibly the system) adds artifacts to the portfolio tool. The second is laying the portfolio out - placing the artifacts in a tree structure, reflecting on the artifacts, etc. - while the final phase is publishing the portfolio. The first two phases happen simultaneously in many cases, and the entire process repeats itself based on external feedback, new artifact creation as time passes, or different requirements of the portfolio. Also note that a user can have multiple portfolios simultaneously at different points in the process.

The System's Initiative

SPARC currently has a 'stealth mode' of artifact acquistion - it can tie into existing assignment submission systems to automatically add a user's submission as an artifact in their repository. These submission systems can be home-grown, such as our in-house Course Management System, or commerical solutions, such as WebCT, both of which we are in the process of connecting.

This mode is completely system driven - it requires no interaction from the user outside of the initial submission. In many cases even this can be avoided; one example is group based assignments, which one user can submit and it will be automatically added to each team member's artifact repository.

The User's Initiative

Almost all other SPARC interaction is at the user's initiative. This includes bringing in other artifacts, creation of the portfolio structure, and ultimately publishing the portfolio.

To this point little of this can be automated, as there's many different possible goals that the user may have for creating their portfolios, and while it is likely possible to catalogue these goals and build them into the system, one of the current problems faced by the portfolio community is buy-in - many users do not see the end benefit to themselves, and often use any excuse to avoid using the system. Making the user jump through hoops, asking them questions about their goals before they even begin creating their portfolio will needlessly drive users away. In our case, we've chosen to pursue mixed-initiative interactions as a possible method of counteracting this problem.

Areas of Improvement Using Mixed-Initiative

At a simplistic level, implementing a wizard-style interface to allow for interactively creating the initial portfolio structure would be fairly reasonable. However, there is currently insufficient research on different types of portfolios and the best way to do the layout/creation of the different types. Without this research, any attempt to simply create a wizard would be wasted effort. A better approach would be to implement the intelligence to track users over time and begin making inferences on the knowledge we can glean from this, then applying this interactively to users. This, of course, becomes an infinitely recursively process, as our initial attempts to help the user may meet with resistance, or we may be off the mark in our attempts to help, in which case there would need to be a way for the user to communicate this to us, and we can incorporate this feedback for our next attempt, if we choose to make one.

Overcoming these issues requires several different approaches in a coherent structure. We need to be able to recognize certain patterns and methods of portfolio creation, and be able to point the user to possible improvements in their methods, such as new categories they may want, adding newly entered artifacts into the structure, populating new structures with currently existing artifacts, etc.

We should also be able to aid the user in creating the initial structure of the portfolio by recognizing certain initial patterns. For example, if the user begins creating their structure with categories such as "Skills," "Work Experience," and "Education," we should be able to recognize the fact they are creating a job searching portfolio and aid them in filling out portions of the structure. While this is similar to a typical wizard-style interface, we gain this knowledge from looking at other users and how they setup the structure. The system doesn't explicitly recognize the portfolio as a job searching one, it just sees analogues between portfolio structures and attempts to apply this information to the creation.

With our 'stealth mode,' we have a large area of possibilities. We can track, over time, where artifacts from different assignment sets (for example, from a class) are generally placed in the portfolio. With this knowledge, we can suggest placements for the artifact in the portfolio. As well, by grouping the artifacts together, if a user creates a new portfolio and begins adding artifacts in the set to certain positions, we can automatically add the rest (with the user's permission, of course).

We can also generate these artifact sets over time, by looking at the user's trends rather than than the information we get from the assignment server, and trends from other users. By looking for artifacts that the user generally places similarly, or that their peers generally place similarly, (assuming we can get information along with the pushed artifacts on class, section and assignment number, for example) we can dynamically generate these sets.

Methods

As mentioned, SPARC is currently a functioning architecture, and has been in use for a year in first-year studies at

Simon Fraser University's Surrey campus. We are currently working on the second major revision. The courses that utilize SPARC also make use of several different types of assignment submission software - three different systems over the last year - giving us a good platform to experiment with.

Our initial attempts at integrating mixed-initiative pieces to SPARC will involve the necessary intelligence to recognize user's actions and be able to reason on them. In this we will be using software developed within our group (Shakya & Kumar 2005) to log, parse, and break down user's movements and actions within the system. The information, as it is broken down, is instantiated into an ontology upon which we can later rely for the reasoning and inference aspects of the cycle.

Knowledge Representation

Use of this software has several prerequisites. The need to represent the information that the parser has pulled out of the information is first and foremost among these. While any of a multitude of different techniques may work, we've chosen to instantiate an ontology based upon this information. Utilizing an ontology allows us to set rules about how things may happen in a portfolio (without arbitrarily constraining the user, of course), as well as allowing us to get information back out over time. The use of an ontology also allows for inferring information from the data collected.

Any knowledge representation method we select needs to be able to fulfill several criteria. It needs to be able to link users into sets of users that can be dynamic over time; it needs to be able to store information on the portfolio structure; it needs to have information on how the current structure was arrived at - the process, what external agents were involved, such as submission tools, the artifacts created, etc.; it requires information on other user actions, to be able to learn the methods the user utilizes to create and update their portfolio; and it needs to be able to store our current knowledge of the user's goals, so that we can use this information and also update it as time passes.

We have been working for some time on an ontology for electronic portfolios, and have made large strides in this area (Brokenshire, Bogyo, & Kumar 2004; Lougheed *et al.* 2005). The ontology at present allows for storing the information on structure and how it was arrived at. Currently we are in the process of expanding the scope of the ontology for the purposes of the mixed-initiative connection and use.

Once we have the ontology itself, we need to make the parser aware of the ontology.

Transferring Information

At present the parsing software takes input the form of eXtensible Markup Language (XML) log files detailing the user's actions. These actions can be anything from highlighting and re-reading certain portions of the reflections they've written to creating new artifacts to publishing the portfolio. The software has such a wide range of inputs due to it's connection to multiple projects and their different requirements, and this allows us to collect large amounts of data to inform ourselves of the user's goals.

Currently the parser only takes log files after the fact, once actions have been complete. We plan to make modification such that it can take in input on the fly. If possible, we'd also like to be able to pass and information structure directly to the parser instead of relying on passing XML output, which is inherently slower and more time- and processor-intensive, taking away from our real goals.

Once the parser has the information, it instantiates the ontology with the information is gets from the log files. This ontology can then be queried directly by the application for inferences about the learning upon which the application can then act.

Conclusions

While electronic portfolios have been a source of increasing research over the last few years, there has been little attempt to integrate a more helpful and useful interface with intelligence on top of existing systems. We believe this may be one of the reasons for the relatively poor uptake by users of electronic portfolios.

We have presented our goals and the methods to achieve them in integrating and expanding the current research in mixed-initiative systems, specifically interfaces and intelligence, with our existing technology for electronic portfolio creation.

References

Bogyo, B.; Lougheed, P.; Kumar, V.; and Bogyo, B. 2004. ePortfolio systems serving multiple clientele. In *Proceedings of ePortfolio 2004* (2004).

British Columbia Ministry of Education. 2004. Graduation portfolio assessment and focus areas: A program guide. `http://www.bced.gov.bc.ca/graduation/ portfolio/moe_grad_portfolio_p1_p2% .pdf`.

Brokenshire, D.; Bogyo, B.; and Kumar, V. 2004. Towards an upper level ontology for eportfolios. In *Proceedings of ePortfolio 2004* (2004).

2004. *Proceedings of ePortfolio 2004*, European Institute for E-Learning.

Lougheed, P.; Bogyo, B.; Brokenshire, D.; and Kumar, V. 2005. Formalizing electronic portfolios in the SPARC ePortfolio tool. Accepted for publication at the Applications of Semantic Web Technologies for E-Learning workshop at the Third International Conference on Knowledge Capture.

McKimm, J. 2001. Using portfolios in medical education. Imperial College of Medicine.

Shakya, J., and Kumar, V. 2005. Recognizing opportunities for mixed-initiative interactions based on the principles of self-regulated learning. Accepted for publication at the American Association of Artificial Intelligence Fall Symposium Series workshop on Mixed-Initiative Problem Solving Assistants.

Smith, K., and Tillema, H. 2003. Clarifying different types of portfolio use. *Assessment & Evaluation in Higher Education* 28(6):625–648.