

RiTa: Creativity Support for Computational Literature

Daniel C. Howe
Brown University
Computer Science Dept.
dhowe@cs.brown.edu

ABSTRACT

The RiTa toolkit for computational literature is a suite of open-source components, tutorials, and examples, providing support for a range of tasks related to the practice of creative writing in programmable media. Designed both as a toolkit for practicing writers and as an end-to-end solution for digital writing courses (the focus of this paper), RiTa covers a range of computational tasks related to literary practice, including text analysis, generation, display and animation, text-to-speech, text-mining, and access to external resources (e.g., WordNet). In courses taught at Brown University, students from a wide range of backgrounds (creative writers, digital artists, media theorists, linguists and programmers, etc.) have been able to quickly achieve facility with the RiTa components, to gain an understanding of core language processing tasks, and to quickly progress on to their own creative language projects.

Author Keywords

Computer education, Computational literature, Digital writing, Software libraries, Creativity support tools

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

While significant research in the Creativity Support (or CST) community has focused on tools to aid creative practice, there has been less work addressing artists and designers interested in directly employing procedural methods and/or thinking [4,7]. Some areas where such practice has received increasing interest in recent years include procedural product design, generative architecture, algorithmic musical composition, and computational literature (the focus of this paper), to name just a few. Teachers of introductory art and design courses that leverage procedural techniques are often faced with the challenge of setting up practical programming tools for student assignments and projects. In these new and rapidly evolving fields, this task can be a difficult one. Students often enter courses with vastly different backgrounds and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5–10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00

skill sets, and their creative projects tend to integrate a variety of programming tasks. A typical approach to this problem has been to employ multiple programming environments, with each providing support for some specific task of interest. For example, an introductory course in computational (or 'digital') literature might use Perl for text parsing and web-crawling, Apple's built-in 'talk' facility for text-to-speech, Adobe's Flash for text display and animation, Max/MSP for audio support, and one of several research-oriented natural language packages for statistical analysis. By relying on the built-in features of these languages and platforms, the instructor can avoid developing a software infrastructure on their own.

One of several unfortunate consequences of this strategy is that significant time must be devoted to teaching the specifics of each new environment. This increases the delay before students are able to move on to more substantive topics, whether related to software engineering, natural language processing, or the creative practice of computational literature itself. Further, students cannot build on previously learned material in subsequent assignments. This lack of scaffolding is especially problematic when student projects tend to span a variety of 'core' tasks and thus require multiple environments to be bridged in a final project; an often formidable task. For example, a somewhat typical student project that involves extracting text from the web, altering it in some way, and visually displaying the results, accompanied by text-to-speech, might use most or even all of the environments mentioned above. It seems clear that with these shortcomings, a fresh approach is warranted.

In this paper we present a novel, creativity-oriented approach that directly addresses the above challenges and provides a flexible means for organizing the practical component of an introductory computational literature course. We describe RiTa, a creativity support library developed by the author in conjunction with a series of courses taught at Brown University from 2007-2009¹. RiTa is implemented in Java, optionally integrates with the Processing language environment, and runs on common platforms including Windows, OS X, Linux, and Unix. It is freely available under an open-source Creative Commons license at <http://www.rednoise.org/rita/>.

¹ Specifically 'Electronic Writing' (LR0021) and 'Programming for Digital Arts & Literature' (CSCI1950.)

THE PROGRAMMING ENVIRONMENT

An important first step in designing creativity support software for procedural practice is the choice of an appropriate programming environment² [1]. A number of considerations, discussed elsewhere [3], influenced our choice of Processing [5] (and Java) for the context of digital literature. First, it was important that the environment provide a relatively shallow learning curve, so that novice programmers could receive immediate rewards for their efforts. Second, it should support rapid prototyping and short develop/test cycles. Third, it should be widely used so that questions, examples, and projects can be easily located on the web. Fourth, it should facilitate both structured (object-oriented) and ad-hoc programming styles. Fifth, it should provide end-to-end support for tasks which, though perhaps not central to the practice of computational literature, are often necessary for fully realizing a work (e.g., simple yet robust access to sound, network, and graphics.) Sixth, all library functions should run in real-time (there should be no need for offline processing). Seventh, student programs should be easily publishable and include source code to facilitate knowledge-sharing. Finally, all programs should be executable in a web-browser environment to eliminate any dependencies on hardware, operating system, and configuration, which can waste valuable time in workshop-style settings.

DESIGN CRITERIA

Once a base environment in which to implement the toolkit was selected, we identified several criteria which we felt would be important in the design and implementation of RiTa, somewhat generally following the methodologies laid out in earlier work on design methods[6].

Ease-of-Use. The primary purpose of the toolkit was to allow students to effectively implement their own creative language works. Thus the more time students spent learning the toolkit, the less useful it would be.

Consistency. The toolkit should use consistent naming, syntax, functions, and design patterns. Therefore if one object was created via the traditional call to 'new', another should not be created via a static or factory creation pattern

Extensibility. The toolkit should easily accommodate new component implementations, whether for replication of existing functionality (for exercises), or for the addition of new (for the needs of specific projects).

Documentation. The toolkit, its data structures, and its implementation should be carefully and thoroughly documented and updated. All naming conventions should be carefully chosen and consistently used.

² Environment in this context refers to the combination of programming language, libraries, development support (e.g., IDEs), and associated tools for writing, compiling, debugging, running, and publishing programs.

Small/Light. To enable download and execution in web browsers, the library should be as efficient as possible with resources, most importantly browser memory (the most recent core version of RiTa, not including WordNet and Text-to-Speech voices, was 2.6 MB.)

Modularity. The interaction between different components of the toolkit should be minimal, using simple, well-defined interfaces. In particular, it should be possible to complete individual projects using small parts of the toolkit without concern for how they interact with the rest. This allows students to learn the library incrementally over a semester.

Share-ability. Students programs should be easily exportable as web applets with optional HTML generation of source code. To facilitate this, the toolkit should be efficient enough to ensure that student programs can run in common web browsers.

Performance: Library functions should be fast enough that students can use the toolkit for interactive projects with all functions returning in 'perceptual' real-time.

Transparency: The library and its underlying support APIs (Java, Processing, Eclipse, etc.) should all be freely available with easily browsable and well-documented source code.

Platform-Agnostic: The library should contain no operating system-specific behavior, allowing it be used in all major browsers on all common platforms.

CORE OBJECTS

The RiTa toolkit is implemented as a Java library comprised of seven independent packages. The core object collection is comprised of approximately 15 classes within the rita.* package, all of which follow similar naming and usage conventions. The rest of the packages provide support for these core objects, but are not directly accessed under typical usage. Each core object (described briefly below) defines the basic properties, methods and support structures for a specific task.

RiText: The basic utility object for strings of text and associated features. Contains a variety of utility methods for typography and display, animation, text-to-speech, and audio playback.

RiAnalyze: Analyzes phrases, annotating each contained word, and the phrase itself, with a range of (customizable) feature data. Default features include word-boundaries, part-of-speech, stresses, syllables, and phonemes.

RiMarkov: Performs analysis and text generation via Markov chains (aka *n-grams*) with options to process single characters, words, sentences, and arbitrary regular expressions.

RiGrammar: Implementation of a probabilistic context-free grammar (with specific literary extensions) to perform

generation from user-specified grammars.

RiLexicon: The built-in, user-customizable lexicon equipped with implementations of a variety of matching algorithms (min-edit-distance, soundex, anagrams, alliteration, rhymes, looks-like, etc.) based on combinations of letters, syllables, stress, pos, and phonemes.

RiTokenizer: A simple tokenizer for word and sentence boundaries with regular expression support for customization.

RiStemmer: A simple stemmer (based on the Porter algorithm) for extracting roots from words by removing prefixes and suffixes.

RiGoogler: A utility object for obtaining unigram, bigram, and weighted-bigram counts for words and phrases via the Google search engine.

RiWordNet: An intuitive interface to the WordNet ontology providing definitions, glosses, and a range of *-onyms* (hypernyms, hyponyms, synonyms, antonyms, meronyms, etc.) Can be transparently bundled into a web-based, browser-executable program.

RiChunker: A simple and lightweight implementation of a phrase-chunker for non-recursive syntactic elements (e.g., noun-phrases, verb-phrases, etc).

RiKWICKer: An implementation of a simple KeyWord-In-Context (KWIC) model for efficient indexing and lookup of words-in-phrases within documents.

RiSpeech: Provides basic cross-platform text-to-speech facilities with control over a range of parameters including voice-selection, pitch, speed, rate, etc.

RiSample: Simple library-agnostic audio support for RiTa that handles playback of .wav and .aiff samples and server-based streaming of .mp3s.

RiHtmlParser: Provides various utility functions for fetching and parsing text data from web pages using either the Document-Object-Model (DOM) or regular expressions.

RiTextBehavior: An extensible set of text-behaviors including a variety of interpolation algorithms for moving, fading, and adjusting text properties.

RiTaServer: Provides remote support (local and distributed) for core objects with potentially expensive initialization routines, e.g., building a large n-gram model with the RiMarkov object.

DOCUMENTATION

RiTa is accompanied by extensive documentation that explains the functionality provided by the toolkit and describes how to use and extend it. In addition to descriptions, examples, tutorial and a comprehensive

reference, the *project gallery* provides students with access to a wide range of existing projects (implemented in RiTa) by other students and artists, all with linked source code. Students can access this archive either for inspiration on projects or for assistance in addressing particular issues. These projects demonstrate proper documentation strategies, a particularly important element for those working with rapidly evolving technologies. Finally, students may, with instructor approval, add their own projects to the gallery, a goal which both inspired students and encouraged participation in the larger community of practicing digital artists.



Figure 1. An interactive student work created with RiTa.

IN THE CLASSROOM

RiTa can be used to create student assignments of varying difficulty and scope. In the simplest exercises, students experiment with a core RiTa object, attempting to generate interesting literary outputs by adjusting its properties and/or supplying new inputs. They are asked to reflect on the potential of the technique represented by the object, (e.g., context-free grammars and the RiGrammar object) and to identify any limits that they encounter. The variety of existing objects provides a range of opportunities for creating such simple assignments. As students become more familiar with relevant concepts, they can be asked to make minor changes or extensions to an existing module, perhaps addressing the limiting elements they had previously experienced. A more challenging task is to re-implement some or all of an object's interface themselves, demonstrating not only that they understand the core concepts, but that they can also implement them efficiently in code. Finally, students are asked to employ the object in question (their own implementation if they have done one) in a creative project of their own design. Here, RiTa provides some useful starting points, including the examples and project gallery described above. More importantly, student projects are conceived and developed with extensive feedback (largely in workshop-style critiques) from both class members and the professor. This iterative process serves to ensure that projects are not only of sufficient technical and intellectual merit, but also can be

feasibly implemented within the specified time frame (generally 2-6 weeks).

EVALUATION

We used RiTa as a basis for the assignments, discussions and student projects in multiple iterations of 'Electronic Writing' (LR0021) and 'Programming for Digital Arts & Literature' (CSCI1950) at Brown University. These courses were open to both undergrads and graduates and focused on the analysis and creation of works of computational literature, with the former sponsored by the Literary Arts Department and the latter sponsored jointly by the Computer Science Department and the Rhode Island School of Design's Digital+Media program. Students backgrounds were highly varied with a diverse mix of creative writers, plastic and digital artists, computer scientists, and digital media theorists. Technical backgrounds varied from programmers with many years of experience to those with zero³. In fact, negotiating this diversity of experience in terms of lectures and assignments was one of the primary challenges of the course, one that would have been difficult to overcome without a unified toolset for all students.

Although more rigorous evaluation is needed, feedback thus far suggests that RiTa provides an effective and positive experience for the majority of students. Students appreciated the fact that they could run all programs on their home computers or laptops (rather than specifically designated lab computers.) Similarly, the ability to pursue interesting projects from the outset and to easily share work on the web were additionally noted benefits. All but the least experienced students found the level of detail in the documentation to be adequate for learning to use the toolkit and appreciated the ease with which they could combine different components to build new and creatively engaging projects. Further, it was generally agreed that having access to a large community of digital artists working with the same tools (e.g., Processing) greatly accelerated students' progress, even though relatively few existing Processing projects have focused on language-based work. Because the instructor was also the primary programmer for the library it was relatively simple to provide bug fixes and reasonable feature requests as they occurred throughout the semester. As it was our hope that RiTa would facilitate a greater degree of hands-on exploration of the possibilities for computational literature, perhaps most convincing of all is the breadth and depth of the student work represented in the RiTa project gallery⁴. It is our plan to conduct a more formal evaluation of our teaching objectives during the next iteration of the course in Spring 2009.

³ One literary arts undergraduate wrote, "looking back, I do feel as though I know SO MUCH more about computers than I used to. Just knowing that Ctrl-C and V will copy and paste has changed my life! I used to pronounce 'control' literally as 'ctrl' (citrel)."

⁴ Available at: http://www.rednoise.org/rita/rita_gallery.htm

CONCLUSIONS

RiTa provides an extensible, end-to-end programming framework for assignments, projects, and class demonstrations in computational literature. It is well-documented, easy to learn, and simple to use. While there exists several frameworks for teaching natural language and computational linguistics [refs], none of these focus on the specific requirements of an arts context. Similarly, existing tools and strategies for teaching digital art rarely address the domain of language. Further still, none have focused on addressing the needs of a highly diverse, inter-departmental student population. It is our hope that RiTa will not only provide a useful resource for instructors teaching in such a context, but may also provide some insight into how best to support creativity and procedural literacy in a range of potential domains extending well beyond the traditional borders of computer science departments.

ACKNOWLEDGMENTS

Special thanks to the Brown University Computer Science and Literary Arts departments and to the RISD Digital+Media program for their generous sponsorship of this work. Also to John Cayley, Bill Seaman, and Braxton Soderman for their insight into the design and development of the toolkit and their feedback on this paper. Most of all to our students for their continual feedback on RiTa and the amazing work they created with it.

REFERENCES

- 1) Bird, S and Loper, E. 2002. NLTK: The Natural Language Toolkit. In Proceedings of the ACL Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics.
- 2) Hartman, C. O. 1996. Virtual Muse: Experiments in Computer Poetry. Wesleyan University Press.
- 3) Howe, D. and Soderman, A. B. 2008. The Aesthetics of Generative Literature: Lessons from a Digital Writing Workshop (forthcoming).
- 4) Mateas, M. 2005. Procedural literacy: Educating the new media practitioner. In On the Horizon: Special Issue on Future Strategies for Simulations, Games and Interactive Media in Educational and Learning Contexts.
- 5) Reas, C. and Fry, B. 2007. Processing: A Programming Handbook for Visual Designers and Artists. MIT Press.
- 6) Resnick, M., Myers, B., Nakakoji, K., Shneiderman, B. Pausch, R., Selker, T., Eisenberg, M. Design Principles for Tools to Support Creative Thinking. NSF Workshop Report on Creativity Support Tools, Washington, DC, 12-14 June, 2005, 37-52.
- 7) Wing, J. M. Computational thinking, Communications of the ACM, v.49 n.3, March 2006.