

# Collaborative Writing Support Tools on the Cloud

Rafael A. Calvo, *Senior Member, IEEE*, Stephen T. O'Rourke, Janet Jones, Kalina Yacef, and Peter Reimann

**Abstract**—Academic writing, individual or collaborative, is an essential skill for today's graduates. Unfortunately, managing writing activities and providing feedback to students is very labor intensive and academics often opt out of including such learning experiences in their teaching. We describe the architecture for a new collaborative writing support environment used to embed such collaborative learning activities in engineering courses. iWrite provides tools for managing collaborative and individual writing assignments in large cohorts. It outsources the writing tools and the storage of student content to third party cloud-computing vendors (i.e., Google). We further describe how using machine learning and NLP techniques, the architecture provides automated feedback, automatic question generation, and process analysis features.

**Index Terms**—Collaborative learning tools, homework support systems, intelligent tutoring systems, peer reviewing.



## 1 INTRODUCTION

WRITING is important in all knowledge-intensive professions. Engineers, for example, spend between 20 percent and 40 percent of their workday writing, a figure that increases with the responsibility of the position [1]. It is often the case that much of the writing is done collaboratively [2]. For example, Ede and Lunsford [2] showed that 85 percent of the documents produced in offices and universities had at least two authors. This result is similar to those in other studies [e.g., 3]. Collaboration and writing skills are so important that accreditation boards such as the Accreditation Board in Engineering and Technology (ABET) require evidence that graduates have the “ability to communicate effectively.”

However, motivating and helping students to learn to write effectively before they graduate, particularly in collaborative scenarios, poses many challenges, many of which can be overcome by technical means. Over the last 20 years, researchers within universities have been developing technologies for automated feedback in academic writing [4], [5], [6], [7] and for enabling collaborative writing (henceforth, CW) [8], [9], but work combining both automated feedback and CW has been scant.

Among the claimed positive effects of writing documents collaboratively are learning, socialisation, creation of new ideas, and more understandable if not more effective

documents [10]. Web 2.0 genres of CW, such as wikis and blogs, have become part of popular culture, and may support these outcomes. Yet, most of the CW performed in a professional context is done using tools such as Microsoft Word following certain patterns, many of which have not changed much in the last two decades. These patterns have been described [11] by their document control methods (centralized, relay, independent, and shared) and the writing strategies used (single and separate writers, scribe, joint writing, and consulted). Commonly used patterns, such as when the different members of a group work on different parts of the document lack concurrency, and require many files to be emailed between authors, often leading to problems in the collaboration process. One reason might be that the tools normally used were designed for individual writing. Although synchronous writing tools have been available for many years, it is only recently that they have become mainstream.

Despite their importance, the tools, the learning processes and writing practices associated with CW are rarely explicitly taught in academia. In professional disciplines, students are typically asked to write in teams on some simulated workplace task, but with little instructional or technical support. Furthermore, the information needed to understand how a team goes about its writing task is lost when using desktop tools such as MS Word.

This paper reports on an architecture for supporting CW that was designed with both pedagogical and software engineering principles in mind, and a first evaluation. The overall aim of the paper is to demonstrate how our system, called iWrite, effectively allows researchers and instructors to learn more about the students' writing activities, particularly about features of individual and group writing activities that correlate with quality outcomes. The evaluation provides data collected in general classroom activities and writing assignments (individual and collaborative), using mainstream tools yet allowing for new intelligent support tools to be integrated. These tools include

- R.A. Calvo and S.T. O'Rourke are with the School of Electrical and Information Engineering, The University of Sydney, NSW 2006, Australia. E-mail: {Rafael.Calvo, Stephen.O'Rourke}@sydney.edu.au.
- J. Jones and P. Reimann are with The University of Sydney, A35-Education Building, NSW 2006 Australia. E-mail: {janet.jones, peter.reimann}@sydney.edu.au.
- K. Yacef is with The University of Sydney, J12-School of Information Technologies Building, NSW 2006 Australia. E-mail: kalina@it.usyd.edu.au.

Manuscript received 12 Apr. 2010; revised 20 Aug. 2010; accepted 3 Oct. 2010; published online 10 Dec. 2010.

For information on obtaining reprints of this article, please send e-mail to: [lt@computer.org](mailto:lt@computer.org), and reference IEEECS Log Number TLTSI-2010-04-0048. Digital Object Identifier no. 10.1109/TLT.2010.43.

automated feedback, document visualizations, and automatically generated questions to trigger reflection. In particular, the evaluation shows how data on the *process of writing* (and its outcomes) can be collected on real-world writing tools rather than in laboratory-type scenarios, highlighting that it is the way the tools are used (not the fact that they are) that has an impact on outcomes.

Through a discussion of the evaluation data, the paper further shows how the instructional feedback provided is designed to support the conceptual understanding of both the *writing process* and *textual practices* as opposed to surface writing and grammar alone.

From a software engineering perspective, our architecture is built around new cloud-based technologies that are likely to change the way people write together. Google Docs, Microsoft Live, and Etherpad among others allow writers to work concurrently on a single document. Some of them (e.g., Google Docs) also provide Application Programming Interfaces that allow developers to build applications on top of such writing environments. Cloud-based technologies are expected to change the technological landscape in many areas, including Computer-Supported Collaborative Learning (CSCL). Systems that support CW on the web have been reviewed before [12], but production grade systems that also offer an API to access their content are very recent.

While wiki engines, blogging tools and web-service oriented document architectures such as Google Docs have simplified the task of CW, they are not designed for supporting novice writers, or for addressing the challenges of learning through CW. In particular, these technologies still fall short of supporting writers with scaffolds and feedback pertaining to the cognitively and socially challenging aspects of the writing process. Our system, described in more detail below, is intended to provide learning support by means of these innovative elements:

- Features to manage writing activities in large cohorts, particularly the management and allocation of groups, peer reviewing, and assessment.
- A combination of generic and computer generated personalized feedback. The generic feedback includes interactive multimedia animations and content. The architecture incorporates features for feedback forms such as Argument quality, features of text (such as coherence), automatic generation of questions, and feedback on the writing process.
- A combination of synchronous and asynchronous modes of CW.
- The use of computer-based text analysis methods to provide additional information on text surface level and concept level to writing groups.
- The use of computer-based process discovery methods to provide additional information on team processes. The combination of these methods with text mining ones is particularly novel, and will allow feedback about team processes based not only on events but on their semantic significance.

iWrite is currently used to support the teaching of academic writing at the Faculty of Engineering and IT, the University of Sydney, to 600 undergraduate and postgraduate students.

This paper is structured as follows: In Section 2, we discuss the literature on academic writing support systems, peer-reviewing and group work. In Section 3, we discuss the architecture of three components of iWrite. Section 4 describes how iWrite is used at the University of Sydney particularly analyzing features of the process of writing and the quality of the outcomes (i.e., grades). Section 5 concludes with a discussion on how we see this tool evolving.

## 2 BACKGROUND

### 2.1 Collaborative Writing

Computer-supported CW has received attention since computers have been used for word processing. Two areas of research are particularly relevant for our project: research that analyzes CW in terms of group work processes, focusing on issues such as process loss, productivity, and quality of the outcomes [8], [9]; and research that studies CW in terms of group learning processes, focusing on topics such as establishing common ground, knowledge building, and learning outcomes [13]. In the second line of research (CSCL), writing is seen as a means to deepen students' engagement with ideas and the literature and for knowledge building [14] by jointly developing a text or hypertext. In CSCL, in addition to knowledge building in asynchronous collaboration, synchronous collaborative development of argumentative structures and texts has received much attention (e.g., [15]).

CW—defined by Lowry et al. [8, p. 72] as “...an iterative and social process that involves a team focused on a common objective that negotiates, coordinates, and communicates during the creation of a common document”—is a cognitively and organizationally demanding process. As a distinct form of group work it involves a broad range of group activities, multiple roles, and subtasks. In addition, when performed by groups that communicate (partially or only) through communication media, the process typically involves multiple tools with different use characteristics (e.g., phone, mail, instant messaging, and document management systems). From a cognitive perspective, (individual) writing has been described as an “ill-structured” problem type, meaning that the writing task has to be clarified by the writer(s) before engaging in any more targeted problem solving [16]. When performed in an educational context, the lecturer typically provides the writing task, writing and communication tools, and group composition, so that teams can focus on team planning and document production. Both of these are typically complex, involving steps such as task decomposition, role definition, task allocation, milestone planning as components of team planning, and brainstorming, outlining, drafting, reviewing, revising, and copy editing as components of document production. These steps are not formal requirements of all collaborative writing activities and are often not managed by the instructor (or the system) but directly by the students.

Because of the complexity of the CW process, explicit support needs to be provided, in particular for novice writers. Such support generally falls into one of three classes: specialized writing and document management tools, document analysis software, and team process support. This project focuses on the latter two.

## 2.2 Writing for Learning

Writing for Learning (WfL), with variations such as Writing Across the Curriculum and Knowledge Building pedagogy [17], has attracted the interest of teachers and of researchers for more than thirty years [18], [19]. It has, for instance, seen widespread use in science education [e.g., 20], [21], [22]. We are proposing WfL not only because research has “shown that it works” (although the empirical findings are, as usual, mixed [23]), but in particular because it can be flexibly employed in formal as well as nonformal learning settings. Furthermore, writing researchers have theorized and studied the intricate relations between cognition, interest, and identity in a holistic fashion before [21], [24], which makes it particularly relevant for engineering education.

A number of reasons have been identified to explain why writing is an important tool for learning. Cognitive psychologists make the general argument that writing requires the coordination of multiple perspectives (content and audience) and the linearization thought, which might not be linear [25]. For subject matter learning, this means that writing requires deep cognitive engagement with content, which will lead to better learning. From a discourse theory perspective, it has been argued that students must learn to understand and reproduce a professional community’s traditional written discourse if they are to become members of that community [24]. And pedagogy-based arguments for the value of writing assert that writing is an important medium for reflection and, in the context of higher education, also a medium for developing epistemic orientations [21].

A specific pedagogical challenge arises when employing WfL in settings with a large number of students, i.e., for undergraduate education: How can guidance (scaffolding) and (formative) feedback on writing be provided, given the teacher:student ratio? This necessitates looking for alternative resources, in the form of self-guided learning, peer feedback, and guidance/feedback that can be provided by computational means. This brings us to automatized approaches.

## 2.3 Automated Essay Feedback and Scoring Systems

Automated feedback systems have been studied for over a decade and most of these systems focus on individual writing, not on collaborative activities. Over this period techniques of Natural Language Processing (NLP) and Machine Learning have progressed substantially and automated writing tutors have improved simultaneously. Despite this progress, the value of automated feedback and essay scoring remains contested [26]. The increasing use of automatic essay scoring (AES) in particular by many institutions has created robust debates about accuracy and pedagogical value. Two recent books discuss advances in AES, one taking a very supportive approach [27] and one providing a more critical debate [28].

Glosser [29] is an automatic feedback tool used within iWrite for selected subjects. It was designed to help students review a document and reflect on their writing [29]. Glosser uses textual data mining and computational linguistics algorithms to quantify features of the text, and produce feedback for the student [30]. This feedback is in the form of descriptive information about different aspects of the

document. For example, by analyzing the words contained in each paragraph, it can measure how thematically “close” two adjoining paragraphs are. If the paragraphs are too “far,” this can be a sign of a lack of flow, and Glosser flags a small warning sign. As a form of feedback Glosser provides trigger questions and visual representations of the document.

Other researchers have used techniques similar to those used in Glosser for Automatic Essay Assessment for building writing support tools. Criterion (by ETS Technologies), MyAccess (by Vantage Learning), and WriteToLearn by Pearson Knowledge Technologies are all commercial products increasingly used in classrooms [26]. These programs provide an editing tool with grammar, spelling, and low-level mechanical feedback. They also provide resources such as thesaurus and graphic features, many of which would be available in tools such as MS Word. To our knowledge, these tools do not have collaborative writing or process oriented support.

Other systems include *Writers Workshop*, developed by Bell Laboratories, and *Editor* [4]. Both focus on grammar and style and showed limited pedagogical benefits [5]. The main difficulty identified was that correcting surface features of student texts did not help them to improve the quality of their ideas or the knowledge of the topics that they were studying.

In *SaK* [6], which is built around a notion of voices that speak to the writer during the process of composition [31], avatars give the impression of “giving each voice a face and a personality” [6]. Students write within the environment, and then the avatars provide feedback on a different aspect of the composition, identifying strengths and weaknesses in the text but without offering corrections. Although students write individually, SaK can also analyze the topic of a sentence, identifying clusters of topics among the students so that when a new topic arises the student can be asked for an explanation or reformulation.

Some systems, such as Summary-Street [7], tend to focus on drills but their reported success is in areas such as learning to summarize, rather than driven by disciplinary concepts.

As far as we know, all the systems reported in the literature are designed as stand-alone activities, normally used outside the context of a real class scenario. This would likely affect the conceptions that students have about the activity, and therefore the way they engage in it. Evidence shows that in collaborative and in writing activities [32], [33] this significantly affects the learning outcomes. Systems like iWrite, which afford collaborative writing activities that are embedded and “constructively aligned” [34] with the assessment and the learning outcomes, are more likely to be successful.

## 3 ARCHITECTURE

The “iWrite” website provides students with information about their writing activities, tools for writing and submitting their assignments, and a complete solution for scaffolding the write-review-feedback cycle of a writing activity. Fig. 1 shows its three sections, two of which (“For Students” and “For Academics”) consist of content and interactive tutorials on developing students’ understanding of different concepts and genres of writing. These consist of

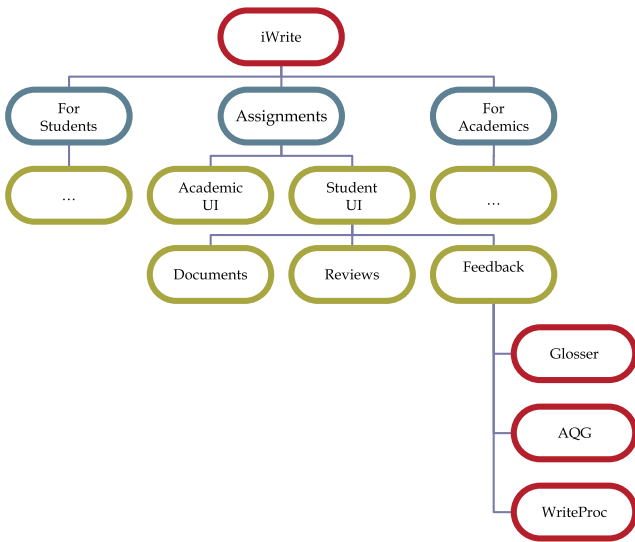


Fig. 1. The iWrite information structure diagram.

discipline specific tutorial exercises where students are introduced to writing concepts through examples written by others. Only the “Assignments” section—which supports students to contextualize these writing concepts in their own compositions—will be discussed here in detail.

**3.1 Assignment Manager**

The Assignment Manager is designed to use cloud computing applications and their APIs. This means that the writing tool and the documents themselves are managed by a third party. This significantly reduces the cost of managing a system with large number of students, and a Service Level Agreement (SLA) ensures that assignment documents are always available.

The architecture of the system is illustrated in Fig. 2. The writing tools and activities, on the left-hand side, are implemented with Google Docs, a cloud-based office suite for editing documents, presentations, and spreadsheets. The API provides programmatic access to the documents.

The right-hand side of Fig. 2 shows the Assignment Manager, Glosser, and WriteProc. Assignment Manager

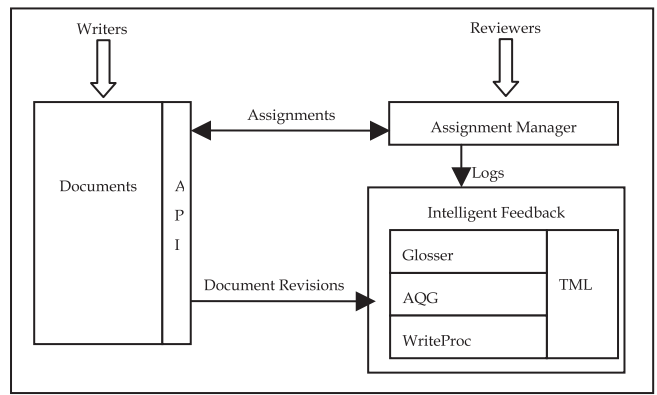


Fig. 2. The iWrite architecture diagram.

deals with the administration and scheduling of courses and writing activities. In addition, tools that analyze the documents using NLP techniques provide additional functionalities. Automatic Question Generation (AQG) generates questions from templates based on the references used in a document. WriteProc is a tool for analyzing students’ usage of iWrite in combination with the methodological process of their writing.

Through the “Assignments” tab of the website, students have access to the documents, feedback from instructors, peers and system, information about deadlines, and so forth. These are shown in the top two boxes of the screenshot of Fig. 3. If the user is identified as an instructor for a particular course, additional features are provided (e.g., downloading a zip file with all the submitted assignments) as shown in the lower box of Fig. 3.

Both Glosser and WriteProc use TML [35], a multipurpose text mining library that implements the NLP and machine learning techniques that analyze the actual content of the document revisions. TML provides a comprehensive set of text mining algorithms and scaffolds every stage of the text mining process. TML integrates the open source Apache Lucene search engine, the Stanford NLP parser and the Weka machine learning libraries, and is itself open source. TML provides functionalities for the preprocessing of documents,

**Writing Tasks**

Course	Activity file	Due date	Submitted	Feedback
ELEC3610	<a href="#">Project Presentation</a>	31 Mar 12:00 PM		
	<a href="#">PSD</a>	Draft: 12 Feb 12:00 PM Final: 12 Feb 12:00 PM	<a href="#">Draft-SNAPSHOT</a> <a href="#">Final</a>	<a href="#">Automatic Feedback</a>

**Reviewing Tasks**

Course	Document to review	Your review	Due date
ELEC3610	<a href="#">Draft-SNAPSHOT</a>	<a href="#">PSD Review</a>	12 Feb 12:00 PM

**Instructor Panel**

Course	Activity	Submitted	Due date
ELEC3610	<a href="#">Project Presentation</a>	<a href="#">Download all drafts</a>	31 Mar 12:00 PM
	<a href="#">PSD</a>	<a href="#">Download all reviews</a> <a href="#">Download all assignments</a>	Draft: 12 Feb 12:00 PM Review: 12 Feb 12:00 PM Final: 12 Feb 12:00 PM

Fig. 3. A screenshot of the Assignment Manager: the student UI displays the writing and reviewing tasks, while the academic UI also displays the instructor panels.

tokenising, stemming, and stop-word removal. It maintains three corpora, adding each new document, at the sentence, paragraph and document level. In order to reduce the lag-time, all these are stored in a repository, along with the results of the text mining operations.

Assignment Manager handles all aspects of the assignment submission, peer-reviewing, and assessment process. It uses the API provided to a Google Apps for Education account to administer user accounts and to create, share, and export documents. The APIs operate using an Atom feed to download data over HTTP. Although the Assignment Manager system is currently only integrated with the Docs service, there is the potential to incorporate other Google services into activities, such as Sites or Calendar. An abstraction layer also allows systems from other vendors to be added.

Google Apps offers a SAML-based single sign-on service that provides full control of the usernames and passwords to authenticate student accounts with Google Docs. This allows Assignment Manager to automatically authenticate logged in users with Google Docs. A user simply needs to click on a link in Assignment Manager to authenticate with Google Docs and start writing their document. Google Docs can be accessed from any browser with an Internet connection as well as offline in a limited capacity and can be synchronized when an Internet connection is available.

The Assignment Manager's academic and student User Interfaces (UI) are illustrated in Fig. 3. The student UI contains two panels tabling a student's writing tasks and reviewing tasks. The writing tasks panel includes links to the activity documents in Google Docs, information about deadlines, a link to download a PDF formatted snapshot of the document at each deadline, as well as links to view the feedback provided by instructors and peers and to automated feedback from Glosser. The reviewing tasks panel is styled in a similar fashion to the writing tasks panel, but instead contains a table of the documents that a student has been allocated to review.

Assignment Manager is administered through a web application based on Google Web Toolkit which facilitates the creation of courses and writing activities. Each course has a list of students (and their contact information), maintained in a Google Docs spreadsheet and synchronized with Assignment Manager on request. Keeping this information in a spreadsheet allows course managers to easily modify enrolment details in bulk, and assign students to groups and tutorials. Assignment Manager maintains a simple folder structure of courses and writing activities on Google Docs. The permission structure of the folder tree is such that lecturers are given permission to view all documents in the course and tutors are given permission to view all documents of the students enrolled in their respective tutorials. A writing activity can specify a document type (i.e., document, presentation, or spreadsheet), a final deadline along with optional draft and review deadlines, along with various other settings.

When a writing activity starts, a document is created for each student or group from a predefined assignment template. The document is then shared accordingly, giving students "write" permission and their lecturers and tutors

"read" permission. Students are notified via email when a new assignment is available and given instructions on how to write and submit their new assignment.

When a draft deadline passes, snapshots of the submitted documents are downloaded in PDF format and distributed for reviewing or marking through links in the iWrite interface shown in Fig. 3. A number of configurations can be used to automatically assign documents to students in peer-reviewing activities, or to tutors, and lecturers for reviewing and marking. These include assignments within a group, between groups or manual. Students are notified via email when a new document is assigned for them to review.

When a review deadline passes, links to the submitted reviews will then appear as icons in the feedback column of the writing tasks panel alongside the corresponding document they critique. Students can click on these icons to view the feedback and revise their document accordingly.

Lastly, when the final deadline of a writing activity passes, the permission of the students' documents is updated from "write" to "read," so documents can no longer be modified by students. A final copy of each submitted document is downloaded in PDF format and distributed to tutors and lecturers for marking.

The Assignment Manager system greatly simplifies many of the administrative process in managing collaborative writing assignments, making it logistically possible for educators to provide feedback to students quicker and more often. Similarly, for students there are notable benefits in the automatic submission process and the use of Google Docs, especially for collaborative assignments.

### 3.2 Intelligent Feedback: Automatic Feedback, Questions, and Process Analysis

Using the APIs the system has access to the revisions of any document. This allows new functionalities such as automatic plagiarism detection, automatic feedback, and automatic scoring systems to be integrated seamlessly with the appropriate version of the document. iWrite currently implements three such intelligent feedback tools, Glosser [30], AQG [36], and WriteProc [37], to generate automatic feedback, questions and process analysis, respectively.

#### 3.2.1 Glosser: Automatic Feedback Tool

Glosser [30] is intended to facilitate the review of academic writing by providing feedback on the textual features of a document, such as coherence. The design of Glosser provides a framework for scaffolding feedback through the use of text mining techniques to identify relevant features for analysis in combination with a set of trigger questions to help prompt reflection. The framework provides an extensible plugin architecture, which allows for new types of feedback tools to be easily developed and put into production.

Glosser provides feedback on the current revision of a document as well as feedback on how the document has collaboratively progressed to its current state. Each time Glosser is accessed, any new revisions of a document are downloaded from Google Docs for analysis.

The feedback provided by Glosser helps a student to review a document by highlighting the types of features a document uses to communicate, such as the keywords and

Home Flow Topics Topics Map Language Participation

**Topics**

- Are the ideas used in the essay relevant to the question?
- Are the ideas developed correctly?
- Does this essay simply present the academic references as facts, or does it analyse their importance and critically discuss their usefulness?
- Does this essay simply present ideas or facts, or does it analyse their importance?

To help you reflect on these questions, Glosser has identified what seem to be the most important topics or recurrent ideas in your essay. Important sentences pertaining to each topic are listed to the right.

Topic	Important sentences	Contributors
Cultures	<ul style="list-style-type: none"> <li>These two fields are two where collaboration and understanding are essential to avoid misinterpretation and cultural conflicts.</li> <li>Counterarguing this point would suggest that language and culture are correlative, so the eradication of one could also influence the cultural differences around the world that provide for a plethora of concepts and ideas.</li> <li>So it is a bit of a catch 22, where developing a lingua franca promotes an ease of communication between diverse cultures, and the possibility of losing this diversity because cultures can lose their identity.</li> </ul>	soro6056
World Language	<ul style="list-style-type: none"> <li>The idea that the English language is increasing around the world is to some a daunting prospect.</li> <li>A language such as English, which is used throughout the world, doesn't necessarily need to be deemed as 'the world language'.</li> <li>Even the presence of non-verbal way of communicating suggests that English or any other language will never be used as a world language.</li> </ul>	soro6056
People	<ul style="list-style-type: none"> <li>So to keep the diversity of the globe alive it should be stated that if people are wanting to reside in a particular area they should not expect the 'natives' to accommodate for their language.</li> <li>To conclude, it can be seen that English is a widespread language but this doesn't mean that it is or will be a universal language for the majority of people.</li> <li>This suggests that as time has progressed and with it the innovations of humans, people have felt the need for cross-cultural communication so that ideas can be shared for the benefit of all mankind.</li> </ul>	soro6056

Fig. 4. The Topics feedback tool in Glosser. The trigger questions are displayed at the top of the page and the “gloss” below.

topics it includes, and the flow of its content. The highlighted features are focused on improving a document by relating them to common problems in academic writing. Glosser is not intended to give a definitive answer on what is good or bad about a document. The feedback highlights what the writers of a document have done, but does not attempt to make any comparison to what it expects an ideal document should be. It is ultimately up to the user to decide whether the highlighted features have been appropriately used in the document.

Glosser has also been designed to support collaborative writing. By analyzing the content and author of each document revision, it is possible to determine which author contributed which sentence or paragraph and how these contribute to the overall topics of the document. These collaborative features of Glosser can help a team understand how each member is participating in the writing process. The user interface of the Topics feedback tool in Glosser is displayed in Fig. 4. The trigger questions at the top of each page are provided to help the reader focus their evaluation on different features of the document. Below the questions is the supportive content called “gloss”, to help the reader answer those questions. The “gloss” is the important feature that Glosser has highlighted in the document for reflection. A rollover window on each sentence indicates who and when wrote it.

### 3.2.2 WriteProc: Process Mining Tool

The autosave function in Google Docs acts as a version tracking functionality, saving documents every 30 seconds or so (as long as the student has written something in that period). This means that, for each single document written

by a student or team, thousands of revisions are stored. This versioning information has the potential to provide a valuable insight into the microstructure of the process students follow while writing. This information can be used to not only understand which processes are most likely to lead to successful outcomes but also to give feedback to students and instructors.

WriteProc takes advantage of these data traces to analyze the processes involved in writing the documents [38], as opposed to focusing on the end product or the snapshot of the document at a given time. It uses log files of page views generated from students using the iWrite assignment submission system and the actual content of the document revisions to build a profile of a student’s behavior. This analysis is performed using a combination of text mining and process mining techniques.

Process mining techniques are used to analyze the process that group writers follow, and how the writing process correlates to the quality and semantic features of the final composition. We have developed heuristics to extract semantic meaning of text changes during writing, and then used these to identify writing activities in the writing processes. WriteProc currently uses a taxonomy of writing activities proposed by Lowry et al. [39]. In addition, we used a model developed by Boiarsky [40] for analyzing semantic changes in writing processes. The writing activities classified include brainstorming, outlining, drafting, revising, and editing which are common categories that, besides being theoretically supported, are also well understood by writing and subject matter instructors.

### 3.2.3 Automatic Question Generation (AQG)

The iWrite architecture includes a novel AQG tool [36] that extracts citations from students’ compositions, together with key content elements. For example, if the students use the APA citation style, author, and year are extracted. Then the citations are classified using a rule-based approach. For example, based on the grammatical structure and other linguistic properties, the citations are identified as an opinion, or describing an aim, or a result, or a method, or a system. Finally, questions are generated based on a set of templates and the content elements. For example, if the citation is an opinion, the AQG could generate a question that looks like: “What evidence is provided by X to prove the opinion?”. If the citation is used in describing a system, the question could look like: “In the study of X what evaluation technique did they use?”

A study on differences in writers’ perception between questions automatically generated by iWrite and humans (Human Tutor, Lecturer, or Generic Question) found that the learners have moderate difficulties distinguishing questions generated by the proposed system from those produced by humans [36]. Moreover, further results show that our system significantly outscores Generic Question on the overall quality measures.

## 4 EVALUATION

We present here three different evaluation aspects. First, we show a traffic analysis of iWrite, which is a key to understanding how the tool is used and the writing



processes involved. Second, we analyzed further how high achieving students differed from other students with respect to the way they worked on their collaborative writing assignment. Lastly, we include some user feedback.

During the first semester 2010, iWrite was used to manage the assignments of four engineering subjects. In total, these courses consisted of 491 students who wrote 642 individual and collaborative assignments, for which we have recorded 102,538 revisions that represent over 51,000 minutes of students work. The amount of data and detail being collected by iWrite about students' learning behaviors is unprecedented. As a way of showing how the architecture can be used in real scenarios, we describe the activities as summary case-studies of iWrite. In order to gain some insights about how collaborative usage of the tool differs from the individual one, we also include here the data collected from two courses with individual assignments. Details of the four courses that participated in the project are as follows:

- **ENGG1803:** Professional Engineering. 154 first year students wrote four collaborative assignments (in groups of five students), as well as an individual one. The topics revolved around an engineering project. The class was divided into three cohorts based on their English proficiency. The evaluation aspects discussed in this section refer to the "general" cohort, which comprises 103 students, the largest number in the course. This cohort undertook two specific activities 1803A and 1803B: **ENGG1803A** consisted of a written presentation and a 4-6 pages long design report. **ENGG1803B** consisted in a 20-30 pages final project report.
- **ELEC3610:** E-Business Analyses and Design. 53 third year students wrote a project proposal (PSD1) in groups of two. Feedback was provided to students as reviews by peers and tutors, and automatically by using Glosser. After the feedback students submitted a revised version of the assignment (PSD2).

#### 4.1 Writing Process

The writing processes followed by students in different activities (or subjects) reflect their understanding of what is expected in the activity, their motivation and other educational factors. Often students' behavior during an activity can be different from what instructors expect, and this variation may raise issues on how the activity is designed. We computed a number of variables that could be expected to have an impact on the composition's quality, and therefore their grades.

- **glosserPageViews:** a measure of how much a student has used Glosser, the automatic feedback tool. This was only available to students in ELEC3610. This information may help determine the impact of automatic feedback tools.
- **iwritePageViews:** how much a student has used iWrite, while writing, reviewing or reading instructional material on aspects of writing. This information could be broken down to study the impact of different content sections. The traffic on a particular part of the site could be reported to the instructor to

indicate whether students are using content specifically designed to support the activity.

- **userRevisions:** while a student is writing, Google Docs automatically saves every 30 seconds or so. This is a measure for the amount of time a student spent typing. Writing time could include time reflecting or reading where revisions are not stored. This measure gives the instructor an accurate estimate of how much time students at the individual, group or cohort level are spending in the activity.
- **teamRevisions:** total number of revisions for the document, including those by the different team members. It is the same as **userRevisions** in individual assignments.
- **contribution:** the ratio  $teamSize \times userRevisions / teamRevisions$  indicates the relative contribution of the individual student to the assignment. This measure is close to zero if the user contributed little, less than one is if they contributed less than their fair share and more otherwise.
- **durationDays:** the number of days spanning from the first revision to the last. Often the assumption is that starting early is good. This measure can provide evidence on the circumstances under which this assumption is true.
- **sessionsWriting:** the number of sessions a student worked on a document. A new session starts if a student has not modified the document for 30 minutes.
- **revisionsPerSession:** the number of revisions is proportional to the time that the student works on a document in a single session. This information could help estimate the optimum amount of time a student can work on a document without taking a break. This measure together with **sessionsWriting** can for example describe if a student, or a group, tends to write in many brief sessions or fewer longer ones. Finer granularity could also be used to see such patterns of writing "sprints" within a session, for example when a writer puts down an idea as stream of thoughts and then goes quiet (e.g., reflects) and comes back to fix what was just written.
- **daysWriting:** the student might have started early but then did nothing until close to the deadline. This measures the number of days in which some work was done.
- **revisionsPerDay:** this measures how much work was done by a student in a single day.
- **gradeActivity:** the grade obtained by a student in the particular writing activity (out of 100).
- **gradeOverall:** the overall grade obtained by a student in the subject (out of 100).

The averages for all students, including those who did not contribute to the group submission (zero revisions), using the system for collaborative assignments are shown in Table 1 and the averages for a selected subset of activities are shown in Table 2. We can see that usage patterns change between activities in the same class. Comparing PSD1 and PSD2, for instance, on average students spent less time writing/ revising the PSD2 document (234 revisions) than the PSD1 document (517 revisions). They used the automatic feedback

**TABLE 1**  
Overall Descriptive Statistics for the Process and Outcomes Measures

	N	Min	Max	Mean	Std. Dev.
glosserPageViews	517	0	73	11.39	16.2
iwritePageViews	517	0	225	47.99	37.86
userRevisions	517	0	1409	145.79	232.92
teamRevisions	517	0	3388	480.85	605.09
contribution	517	0	6.00	1.00	1.13
durationDays	517	0	48.47	4.44	7.62
sessionsWriting	517	0	32	4.26	4.97
revisionsPerSession	517	0	441	26.40	37.24
daysWriting	517	0	12.00	2.00	2.03
revisionsPerDay	517	0	599.00	58.34	82.42
gradeActivity	517	0	100	71.59	13.89
gradeOverall	517	12	87.50	68.07	10.18

**TABLE 2**  
Average Measures for Each Writing Activity

	PSD1	PSD2	1803A	1803B
glosserPageViews	4.45	18.32	N/A	N/A
iwritePageViews	65.89	48.13	35.92	46.76
userRevisions	517.17	234.30	91.81	52.81
teamRevisions	1034.34	465.66	481.19	300.01
contribution	1.00	1.00	1.00	1.00
durationDays	7.74	3.61	2.79	2.73
sessionsWriting	9.68	4.85	3.89	2.26
revisionsPerSession	57.34	51.85	18.43	9.71
daysWriting	3.83	2.08	2.09	1.09
revisionsPerDay	139.84	117.47	39.31	22.25
gradeActivity	60.75	72.08	74.33	73.17
gradeOverall	64.55	64.55	69.34	69.34

tool more in the second assignment: 4 glosserPageViews for PSD1 and 18 for PSD2. Both of these results are what were expected by the instructors.

The total amount of time students spend working on the assignment, and how the time is distributed, can provide useful information to detect successful writing processes. For example, if all the work is done in the last few hours before the assignment is due, we would expect this process to lead to lower quality outcomes. If the project is collaborative, the need for an early start should be even higher, since team members need to find common ground and agreement on the topic, structure and other aspects of the composition. Fig. 5 shows the average number of revisions written per student per day over the period of an individual (ENGG1803) and a collaborative (ELEC3610) assignment. Both assignments were limited to 2,000 words in length.

In the case of the individual assignment, the majority of the students did not begin writing in Google Docs until the last couple of days before the deadline, when students made an average of 16 and 33 revisions. This compares to less than two revisions per student per day in the 13 days prior. Looking at the data, it was found that the majority of students made less than 20 revisions to their documents, while a minority of more active students accounted for the

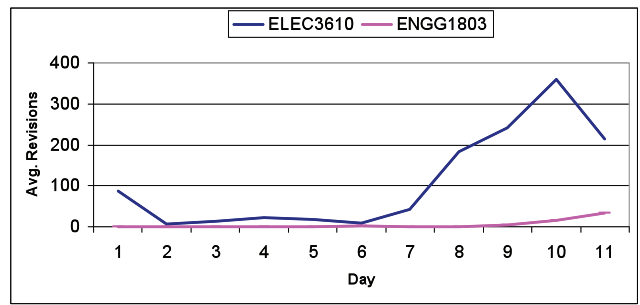


Fig. 5. A plot comparing the average number of revisions written per student per day for an individual (ENGG1803) and a collaborative (ELEC3610) assignment.

**TABLE 3**  
Grouping of Engg1803 Students According to Their Collaborative Writing Assignment Grades

Assignment Mark	N	Mean	Std. Dev.	Min	Max
Low	28	58.30	3.96	47.5	62.5
Medium	103	72.05	4.57	63.3	79.5
High	23	83.66	2.49	80.5	88.3
Total	154	71.28	8.47	47.5	88.3

majority of the total revisions committed. This is likely due to the fact that some students chose to write their assignments using an alternative word processor and then copied and pasted their work into Google Docs for submission. Reasons for this behavior include 1) the fact that they are more accustomed to desktop word processor applications, such as Microsoft Word and OpenOffice, 2) the lack in GoogleDocs of a bibliography package (e.g., Endnote) required in one of the courses, and 3) concerns about working offline on a web application.

However, contrary to the above, it was found that students working on collaborative assignments made much more consistent use of Google Docs. The reasons for this are believed to be twofold. First, the superior collaborative functionality of Google Docs allows multiple students to synchronously work on the same document at the same time. Second, accessibility of the document through the Glosser tool meant that authorship of the different keywords, sentences, paragraphs, and topics of the document could be attributed to individual students and highlighted for all to see. This made the extent and value of individual participation in the collaborative writing process more transparent to all.

#### 4.2 Relating Aspects of iWrite Use with Student Performance

We categorized students of the ENGG1803 course into three groups, according to the mark they received for their collaborative writing assignment. We considered a low mark to be below one standard deviation from the mean mark, a medium mark to be within one standard deviation from the mean, and a high mark above the mean plus one standard deviation. Table 3 summarizes the details of these three groups.

We then compared these three groups in relation to the iWrite usage variables defined in the previous section, using ANOVA. We found that the four variables which



were significantly related to grades were userRevisions ( $F = 3.146$ ,  $p = 0.049$ ), teamRevisions ( $F = 3.388$ ,  $p = 0.025$ ), sessionsWriting ( $F = 6.381$ ,  $p = 0.003$ ), and daysWriting ( $F = 7.948$ ,  $p = 0.001$ ). Post hoc analysis (using Tukey's HSD) revealed the following:

- Students with low grades did more individual and group revisions compared to those with medium grades ( $p = 0.025$  and  $p = 0.026$ , respectively)
- Students with low and medium grades engaged in fewer writing sessions compared to students with high grades ( $p = 0.036$  and  $p = 0.001$ , respectively)
- Students with low and medium grades engaged in fewer writing days compared to students with high grades ( $p = 0.003$  and  $p < 0.001$ , respectively).

These results indicate that it is not *whether*, but *how*, students used iWrite which made a difference in their CW performance. Students who obtained high grades were in teams that engaged more frequently in sustained writing sessions. Fewer bursts of document revisions were associated with lower marks.

### 4.3 User Feedback

We collected informal feedback from the course lecturers who used iWrite. They were extremely positive about the experience. One course manager commented "An online assignment submission system will save us a lot of time sorting and distributing assignments. In addition, we send copies of a portion of our assessments to the Learning Centre, so online submission really minimizes our paper usage."

User feedback also highlighted the risk of users having to learn new technologies. For example, students found the use of Google Docs and the automatic submission process confusing, "A number of students tried to upload a word document or created a new Google document, rather than cutting and pasting in to the Google document we had created for them ..." There was also criticism of the layout and lack of styling options in Google Docs, "Google Docs seems to act like one long page, so students work was not well formatted (we had assigned grades for formatting)."

## 5 CONCLUSIONS

The architecture for iWrite, a CSCL system for supporting academic writing skills has been described. The system provides features for managing assignments, group and peer-reviewing activities. It also provides the infrastructure for automatic mirroring feedback including different forms of document visualization, group activity, and automatic question generation.

The paper has focused on the theoretical framework and literature that underpins our project. Although not a complete survey of the extensive literature in the area, it highlights aspects that later supported architectural decisions.

A key design aspect was the use of cloud computing writing tools and their APIs to build tools that make it seamless for students to write collaboratively either synchronously or asynchronously.

A second design guideline was that data mining tools should have access to the document at any point in time to be able to provide real time automatic feedback.

A final design guideline is based on the principle that advanced support tools should be embedded in real learning activities to be meaningful for students. This means also incorporating scaffolding that academics can use to incorporate collaborative writing activities easily within their curricula.

We described aspects of its use with large cohorts, and comments from students and administrators. While an evaluation of the system's impact on learning and the students' perceptions of writing are outside the scope of this paper, we analyzed student use of iWrite in relation to student performance and found that the best predictors for high performance are the way students use iWrite, not necessarily whether they used the tool. This is an important finding that gives clear design guidelines for teachers as well as explicit good writing practices for students. Our future evaluation work will include showing this type of statistical information to instructors and inquire if the values are what they expected and how these data can be used to inform their pedagogical designs.

## ACKNOWLEDGMENTS

Assignment Manager was funded by a University of Sydney TIES grant. Glosser was funded by Australia Research Council DP0665064 and AQG and WriteProc were funded by DP0986873.

## REFERENCES

- [1] M.L. Kreh, "A Survey of the Co-Op Writing Experiences of Recent Engineering Graduates," *IEEE Trans. Professional Comm.*, vol. 43, no. 2, pp. 137-152, June 2000.
- [2] L.S. Ede and A.A. Lunsford, *Singular Texts/Plural Authors: Perspectives on Collaborative Writing*. Southern Illinois Univ., 1992.
- [3] G.A. Cross, *Forming the Collective Mind: A Conceptual Exploration of Large-Scale Collaborative Writing in Industry*. Hampton, 2001.
- [4] Modern Language Association, E.C. Thiesmeyer and J.E. Thiesmeyer, eds., 1990.
- [5] T.J. Beals, "Between Teachers and Computers: Does Text-Checking Software Really Improve Student Writing?" *English J. Nat'l Council of Teachers of English*, vol. 87, pp. 67-72, 1998.
- [6] P. Wiemer-Hastings and A.C. Graesser, "Select-a-Kibitzer: A Computer Tool that Gives Meaningful Feedback on Student Compositions," *Interactive Learning Environments*, vol. 8, pp. 149-169, 2000.
- [7] D. Wade-Stein and E. Kintsch, "Summary Street: Interactive Computer Support for Writing," *Cognition and Instruction*, vol. 22, pp. 333-362, 2004.
- [8] P.B. Lowry, A. Curtis, and M.R. Lowry, "Building a Taxonomy and Nomenclature of Collaborative Writing to Improve Interdisciplinary Research and Practice," *J. Business Comm.*, vol. 41, pp. 66-99, 2004.
- [9] G. Erkens, J. Jaspers, M. Prangmsma, and G. Kanselaar, "Coordination Processes in Computer Supported Collaborative Writing," *Computers in Human Behavior*, vol. 21, pp. 463-486, 2005.
- [10] N. Phillips, T.B. Lawrence, and C. Hardy, "Discourse and Institutions," *Academy of Management Rev.*, vol. 29, pp. 635-652, 2004.
- [11] I.R. Posner and R.M. Baecker, "How People Write Together," *Proc. 25th Ann. Hawaii Int'l Conf. System Sciences*, vol. 4, pp. 127-138, 1992.
- [12] S. Noël and J.-M. Robert, "How the Web Is Used to Support Collaborative Writing," *Behaviour & Information Technology*, vol. 22, pp. 245-262, 2003.
- [13] M. Scardamalia and C. Bereiter, "Higher Levels of Agency for Children in Knowledge Building: A Challenge for the Design of New Knowledge Media," *The J. Learning Sciences*, vol. 1, pp. 37-68, 1991.
- [14] M. Scardamalia and C. Bereiter, "Knowledge Building: Theory, Pedagogy, and Technology," *The Cambridge Handbook of the Learning Sciences*, R.K. Sawyer, ed., Cambridge Univ., 2006.

- [15] M. van Amalvoort, J. Andriessen, and G. Kanselaar, "Representational Tools in Computer-Supported Collaborative Argumentation-Based Learning: How Dyads Work with Constructed and Inspected Argumentative Diagrams," *J. Learning Sciences*, vol. 16, pp. 485-521, 2007.
- [16] J.R. Hayes and L.S. Flower, "Identifying the Organization of the Writing Process," *Cognitive Processes in Writing*, L.W. Gregg and E.R. Steinberg, eds., pp. 3-30, Erlbaum, 1980.
- [17] M. Scardamalia and C. Bereiter, "Knowledge Building: Theory, Pedagogy, and Technology," *The Cambridge Handbook of the Learning Sciences*, R.K. Sawyer, ed., pp. 97-115, Cambridge Univ., 2006.
- [18] C. Bereiter and M. Scardamalia, *The Psychology of Written Composition*. Lawrence Erlbaum, 1987.
- [19] D. Galbraith, "Writing as a Knowledge-Constituting Process," *Knowing What to Write: Conceptual Processes in Text Production*, T. Torrance and D. Galbraith, eds., pp. 139-150, Amsterdam Univ., 1999.
- [20] L.P. Rivard, "A Review of Writing to Learn in Science: Implications for Practice and Research," *J. Research in Science Teaching*, vol. 31, pp. 969-983, 1994.
- [21] V. Prain, "Learning from Writing in Secondary Science: Some Theoretical and Practical Implications," *Int'l J. Science Education*, vol. 28, pp. 179-201, 2006.
- [22] M.A. McDermott and B. Hand, "A Secondary Analysis of Student Perception of Non-Traditional Writing Tasks over a Ten Year Period," *J. Research in Science Teaching*, vol. 47, pp. 518-539, 2009.
- [23] M. Gunel, B. Hand, and V. Prain, "Writing for Learning in Science: A Secondary Analysis of Six Studies," *Int'l J. Science and Math. Education*, vol. 5, pp. 615-637, 2007.
- [24] J.P. Gee, "Language in the Science Classroom: Academic Social Languages as the Heart of School-Based Literacy," *Crossing Borders in Literacy and Science Instruction: Perspectives in Theory and Practice*, E.W. Saul, ed., pp. 13-32, Int'l Reading Assoc., 2004.
- [25] J.R. Hayes and L.S. Flower, "Identifying the Organization of the Writing Process," *Cognitive Processes in Writing*, L.W. Gregg and E.R. Steinberg, eds., pp. 3-30, Erlbaum, 1980.
- [26] M. Warschauer and P. Ware, "Automated Writing Evaluation: Defining the Classroom Research Agenda," *Language Teaching Research*, vol. 10, pp. 157-180, 2006.
- [27] M.D. Shermis and J. Burstein, *Automated Essay Scoring: A Cross-Disciplinary Perspective*. Lawrence Erlbaum Associates, 2003.
- [28] P.F. Ericsson and R.H. Haswell, *Machine Scoring of Student Essays: Truth and Consequences*. Utah State Univ., 2006.
- [29] R.A. Calvo and R.A. Ellis, "Students' Conceptions of Tutor and Automated Feedback in Professional Writing," *J. Eng. Education*, vol. 99, no. 4, pp. 427-438, Oct. 2010.
- [30] J. Villalon, P. Kearney, R.A. Calvo, and P. Reimann, "Glosser: Enhanced Feedback for Student Writing Tasks," *Proc. Eighth IEEE Int'l Conf. Advanced Learning Technologies (ICALT '08)*, pp. 454-458, 2008.
- [31] L. Flower, *The Construction of Negotiated Meaning: A Social Cognitive Theory of Writing*. Southern Illinois Univ., 1994.
- [32] R. Ellis and R.A. Calvo, "Discontinuities in University Student Experiences of Learning through Discussions," *British J. Educational Technology*, vol. 37, pp. 55-68, 2006.
- [33] R.A. Ellis, C. Taylor, and H. Drury, "University Student Conceptions of Learning through Writing," *Australian J. Education*, pp. 6-28, 2006.
- [34] J.B. Biggs, *Teaching for Quality Learning at University: What the Student Does*. Open Univ., 1999.
- [35] TML - Text Mining Library, <http://tml-java.sourceforge.net>, 2011.
- [36] M. Liu, R.A. Calvo, and V. Rus, "Automatic Question Generation for Literature Review Writing Support," *Intelligent Tutoring Systems, V. Alevan, J. Kay, and J. Mostow, eds.*, vol. 1, pp. 45-54, Springer, 2010.
- [37] V. Southavilay, K. Yacef, and R.A. Calvo, "WriteProc: A Framework for Exploring Collaborative Writing Processes," *Proc. Australasian Document Computing Symp. (ADCS)*, 2010.
- [38] V. Southavilay, K. Yacef, and R.A. Calvo, "Process Mining to Support Students' Collaborative Writing," *Proc. Int'l Conf. Educational Data Mining*, pp. 257-266, 2010.
- [39] P.B. Lowry and J.F. Nunamaker Jr., "Using Internet-Based, Distributed Collaborative Writing Tools to Improve Coordination and Group Awareness in Writing Teams," *IEEE Trans. Professional Comm.*, vol. 46, no. 4, pp. 277-297, Dec. 2003.
- [40] C. Boiarsky, "Model for Analyzing Revision," *J. Advanced Composition*, vol. 5, pp. 67-78, 1984.



**Rafael A. Calvo** received the PhD degree in artificial intelligence applied to automatic document classification in 2000. He is an associate professor in the School of Electrical and Information Engineering, University of Sydney, and director of the Learning and Affect Technologies Engineering (Latte) research group. He has also worked at Carnegie Mellon University, Universidad Nacional de Rosario, and as a consultant for projects worldwide. He is an author of numerous publications in the areas of affective computing, learning systems, and web engineering, the recipient of three teaching awards, and a senior member of the IEEE, the IEEE Computer Society, and the IEEE Education Society. He is an associate editor of the *IEEE Transactions on Affective Computing*.

**Stephen T. O'Rourke** received the bachelor of engineering (electrical) degree from the University of Sydney in 2007 and the master of philosophy degree in 2009. He is a researcher in the Learning and Affect Technologies Engineering (Latte) group in the School of Electrical and Information Engineering, University of Sydney. His main research interests are in text mining applied to education technologies.

**Janet Jones** received the Doctorate in Education (EdD) degree from the University of Sydney, Australia. She is a senior lecturer with the Learning Centre, University of Sydney. She has worked as a teacher in Australia, United Kingdom, and Germany, specializing in teaching English as a foreign language.

**Kalina Yacef** received the PhD degree in computer science from University of Paris V in 1999. She is a senior lecturer in the Computer Human Adaptive Interaction (CHA) research lab at the University of Sydney, Australia. Her research interests span across the fields of artificial intelligence and data mining, personalisation and computer human adapted interaction with a strong focus on education applications. Her work focuses on mining users' data for building smart, personalized solutions as well as on the creation of novel and adaptive interfaces for supporting users' tasks. She regularly serves on the program committees of international conferences in the fields of artificial intelligence in education and educational data mining and she is the editor of the new *Journal on Educational Data Mining*.



**Peter Reimann** received the PhD degree in psychology from the University of Freiburg, Germany. He is a professor of education in the Faculty of Education and Social Work, University of Sydney, and a senior researcher in the university's Centre for Research on Computer-Supported Learning and Cognition (CoCo). Current research activities comprise, among other issues, the analysis of individual and group problem solving/learning processes and possible support by means of IT, and analysis of the use of mobile IT in informal learning settings. He is among other roles associate editor for the *Journal of the Learning Sciences*.