



PERGAMON

Computers & Education 37 (2001) 297–315

---

---

**COMPUTERS &  
EDUCATION**

---

---

www.elsevier.com/locate/compedu

## A Web-based EFL writing environment: integrating information for learners, teachers, and researchers

David Wible<sup>a,\*</sup>, Chin-Hwa Kuo<sup>b</sup>, Feng-yi Chien<sup>a</sup>,  
Anne Liu<sup>a</sup>, Nai-Lung Tsao<sup>b</sup>

<sup>a</sup>*Graduate Institute of Western Languages and Literature, Tamkang University, Tamsui, Taipei Hsien, Taiwan*

<sup>b</sup>*Computer and Networks (CAN) Laboratory, Department of Computer Science and Information Engineering, Tamkang University, Tamsui, Taipei Hsien, Taiwan*

---

### Abstract

With the rise in the popularity of web-based education, there is a pressing need for the design of web-based systems that are domain-specific. This need is particularly acute for the domain of second language education, where generic web-based systems fall short of fulfilling the potential of the Internet for meeting the particular challenges faced by language learners and teachers. A novel interactive online environment is described which integrates the potential of computers, Internet, and linguistic analysis to address the highly specific needs of second language composition classes. The system accommodates learners, teachers, and researchers. A crucial consequence of the interactive nature of this system is that users actually create information through their use, and this information enables the system to improve with use. Specifically, the essays written by users and the comments given by teachers are archived in a searchable online database. Learners can do pinpoint searches of this data to understand their individual persistent difficulties. Teachers can do the same in order to discover these difficulties for individual learners and for a class as a whole. The architecture of the system makes possible a novel approach to corpus analysis of learner errors. Teachers' annotations of learner errors are exploited as bootstraps which can lead to the incremental uncovering of errors without the need to heavily error tag the learner corpus. Error analysis then feeds the design of online help content. © 2001 Elsevier Science Ltd. All rights reserved.

*Keywords:* Interactive learning environments; Applications in subject areas; Architectures for educational technology system

---

---

\* Corresponding author. Fax: +886-23620-9912.  
E-mail address: dwible@mail.tku.edu.tw (D. Wible).

## 1. Introduction

The resources needed for developing web-based learning environments are often categorized into two types: system and content. This division underlies, for example, the distinction between ISPs (Internet Service Providers) and ICPs (Internet Content Providers). The simplest relation between these two is one of near-complete independence: ISPs provide their systems and ICPs prepare their domain content in relative mutual autonomy. A consequence of such a loose relation is that content providers often end up accommodating their content to existing systems rather than imagining first how the technology should be designed to accommodate the needs of the content and the learners. Hence these content providers have to make use of existing Internet technology as-is, offering, for example, linked webpages with traditional classroom content such as lecture notes and syllabi, readings that students can access by ftp or gopher, e-mail and discussion boards to communicate with and among learners. MUDs and MOOs offer further existing interactive environments for certain domains of learning.

If the Internet is to fulfill its potential for providing learning environments, however, the common gap between system and content providers must be bridged. Specifically, systems must be developed expressly to meet the unique needs of particular learning domains in ways that traditional classrooms can not. The need for domain-specific systems is not simply a theoretical topic for abstract discussions. With the growing popularity of web-based education, schools are asking teachers to use ‘one-size-fits-all’ web-based systems designed generically to serve virtually every academic field from math and physics to history and language arts, with no special considerations taken in the system design for the unique needs of the different content areas.<sup>1</sup>

A domain where the problem can be seen particularly clearly is second language education. It is a truism by now that gaining facility in a second language is not simply a matter of obtaining a body of knowledge and that language education, therefore, entails something qualitatively different from passing on knowledge from teachers (and books) to learners (Krashen, 1994). Mastery of a language is a skill. Granted, it is a skill that depends on mastering a certain kind of knowledge, but it is a skill nevertheless—one which transcends the simple mastery of facts. Consequently, much as someone learning to swim needs access to water and not just facts about swimming, a language learner requires more than access to lecture notes and readings with traditional objective tests and quizzes. Certainly e-mail, discussion boards, and chat rooms offer learners a forum for using the target language as a communicative tool.

One of the motivating assumptions of our project, however, is that a web-based system can be designed to meet more precisely the needs of not just a learner, but specifically a language learner. The system we describe here focuses on a narrowly circumscribed domain within second language learning: EFL composition.<sup>2</sup> In addition to describing the architecture and functionality of the system, we point out novel ways in which the system design supports a range of research into learners’ grasp of the target language. Specifically, we borrow the notion of “affordance” from research on visual perception and situated cognition (Gibson, 1979) and illustrate how the

---

<sup>1</sup> At a recent conference on foreign language teaching in Taiwan, a presenter pointed out that this ‘one-size-fits-all’ type of platform was implemented at a national university. For ESL courses, a serious obstacle to the system’s success was the generic nature of the system which ignored the unique requirements of language learners (Chang, 2000).

<sup>2</sup> The system described here for EFL composition is in fact one component in a larger system that integrates various language skills. See Wible, Kuo, Tsao, and Liu (2001) for a description of the larger system and other modules in it.

tracking and storage of both learner and teacher input to the online system yields up affordances or handles for bootstrapping into knowledge about these users.

The system that we have designed, called IWiLL (Intelligent Web-based Interactive Language Learning), has been designed based upon certain general assumptions. First, the ideal online system for EFL writing classes should be interactive. Second, the system should exploit the computers' capacity to track the content of the interactions between users, that is, the essays written by the learners, the comments given by the teacher on these essays, and the response of the learner to specific comments. Moreover, the users involved should be enabled to do pinpoint searches of the record of these interactions. This feature makes available invaluable information that can serve as a cumulative source of insight for both learners and teachers, information which in traditional writing classrooms remains scattered, ephemeral, and highly inaccessible. Third, while the system described here is designed for second language writing classes, it is more accurately seen as one component in an integrated language learning environment that includes other skills, such as reading and listening (Wible, Kuo, Tsao, Liu, Sung, & Chio, 2000). The modularized and integrated design is intended to accommodate recent trends in language pedagogy which view language skills as best learned in an integrated whole rather than as a set of separate, independent skills. Fourth, we have designed not a course or a curriculum, but a platform. While the platform is designed specifically to meet the needs of a certain type of language course (second language composition) it is intended to provide as much freedom as possible for teachers within this domain to use their own approaches and materials of their choice. Finally, we propose a novel approach to one of the central research challenges in the development of intelligent CALL systems: when the system permits a high degree of freedom to users with respect to the form and content of their language output, how can that unwieldy data be harnessed to uncover insights into the users.

## 2. A brief overview of the system

During the 2000–2001 academic year, IWiLL was used in Taiwan for EFL writing instruction by 3625 students and 31 teachers from 11 schools—seven high schools and four universities. The students produced over one million words of English writing which was automatically archived in a learner corpus called *EnglishTLC* (Taiwan Learners Corpus). Approximately one fourth of the texts in the corpus were annotated by teachers with corrections and comments over the online system.

Based on the idea that the use of an interactive web-based learning environment creates valuable information, the project reported here distinguishes the following sorts of such information that are worth making available to the appropriate users: (1) the target language production by the language learners themselves, in this case, English essays written by students in university composition classes in Taiwan; (2) the comments that teachers make on these essays for the benefit of the student writers who submit them; (3) the revisions that students make based on these teacher comments. Our system tracks other sorts of information, but the scope of this paper will be confined to these three.

The IWiLL writing component is composed of a set of modules, their interconnections, and various interfaces appropriate to students, teachers, and researchers using the system. The basic motivation for developing a single integrated system for these different users is that we assume

these groups to be interdependent. The idea, then, is to develop a system which supports (and exploits) this interdependence. A brief sketch of the system might illustrate this notion in more concrete terms.

The most obvious pairing of interdependent users is that between teacher and student. Once a teacher and her students are registered on the system, students can compose essays (or import essays that they have composed off-line) and hand these in to their teacher over the Internet. The teacher can access these essays, mark them with comments, and return them to the students through the computer network as well. Students can then revise these returned essays using a screen display that allows them to toggle between the original version with teacher's comments and a duplicate version where they can make their revisions.

Another relation that this system aims to support is that between language learner and researcher. Specifically, researchers in second language acquisition and pedagogy can benefit from access to a searchable corpora of learner essays (Granger 1998; Meunier, 1998). This sort of data makes it possible to investigate pervasive patterns of difficulty in the learners' English (that is, to investigate what some applied linguists call the 'interlanguage' of learners). Moreover, based upon this sort of data, researchers can improve the design of teaching and reference materials. Ultimately it is intended that the findings of such research will be distilled into finely honed online help for learners using this system in the future (Section 4).

To support this mutually beneficial relationship between learner and researcher, with the students' permission, the essays that they turn in to their instructor through this system are automatically archived in a corpus of student compositions. The archive can be searched by researchers (and, in fact, by teachers and even learners themselves). The essays enter the corpus already annotated with the date of writing, the essay title, and relevant information provided by the authors when they register (age, sex, academic major, years of studying English, native language/dialect, etc.). While the authors of essays remain anonymous in the database, each author is identified by a randomly assigned identification number, making it possible to determine which essays were written by the same learner and to access relevant metadata about that learner without revealing who it is. This database of learner essays can be searched online. Currently, the query tools can perform key-word-in-context (KWIC) searches and the results can be saved and edited. This function enables a user to retrieve from the entire set of student essays a listing of every sentence that contains a particular word or phrase, for example, every sentence containing the word "ever" or the phrase "look into." (Fig. 1)

### 3. The organization of the system

The preceding section gives a rough sketch of the overall system. In what follows, we consider specific aspects of the system in more detail and the directions for future developments that it suggests. We begin by considering the point of view of a teacher.

#### 3.1. *The teacher's view*

A registered teacher who logs onto the system is presented a display screen of various links to components within the system. To correct student essays, the teacher links to a page which displays

their student roster (Fig. 2). Beside each student's name on the roster is a row of colored buttons, each button representing an essay that the student has turned in. The color of the button reflects the current status of that essay: red indicates, for example, that the essay has been turned in but not yet marked; yellow, that the essay has been marked (perhaps partially) but not yet returned to the student; green, that the essay has been marked and already returned to the student. (Note: If students have turned in rewrites, these are represented by buttons that appear aligned beneath the button representing the original. Teachers can toggle between a view of a student's rewrite and an original in order to compare the two versions point by point.)

From this roster screen, the teacher can select an essay for marking. The teacher retrieves the essay by clicking on the button that represents that essay on the roster page. A screen then appears displaying the essay in an environment that provides various tools for marking it (Fig. 3). Once in this environment, marking an essay is quite straightforward. To mark the essay with a comment (e.g. to mark a run-on sentence or subject-verb agreement problem), the teacher first chooses the portion of the essay targeted for comment using the mouse. Once the relevant text has been selected, there are two ways for the teacher to provide the student with a comment on it. The first is to simply type the comment in the empty text box provided especially for the teacher's

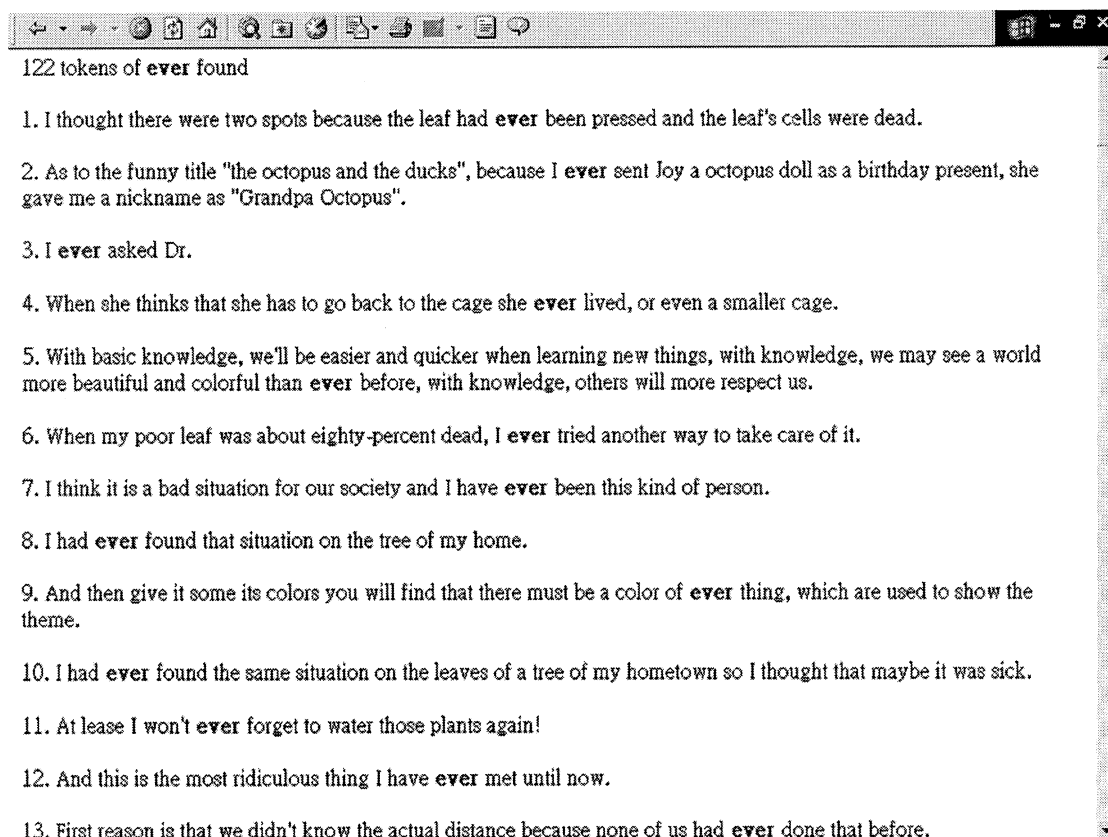


Fig. 1. Display of concordance search results for 'ever' from learner corpus.

comment and then, once the comment has been composed, append it to the intended portion of the student's text by clicking on the appropriate button ('Give the comment'). The second way of providing a comment is to choose one that has been stored in a "Comment Bank." This second way deserves some elaboration.

### 3.1.1. The Comment Bank

The Comment Bank provides each teacher with a convenient means for storing and reusing frequently used comments. To retrieve a stored comment and append it to the portion of the student essay, the teacher simply selects that comment from a drop-down menu and clicks on it (Fig. 4). If the teacher would like to revise the stored comment to better suit the particular situation at hand, an intermediate step allows for the content of the comment to be edited. Once edited, that revised version of the comment can either be used only for the current case and then discarded (retaining the original version of that comment in the bank) or it can be stored for future use in its new form. The teacher can add new comments to her Comment Bank at any time.

The Comment Bank feature circumvents the need for teachers to write the same (or nearly the same) comment from scratch each time a student problem warrants it. Moreover, once a comment has been stored, its length has virtually no effect on how easy or difficult it is for the teacher

**IWILL**  
Intelligent Web-based Interactive Language Learning  
| Main Menu | Back to | Teacher resources | Language help | Help |

● Uncorrected ● Still correcting ● Corrected and returned

**Main page**  
**What's New**  
**About IWILL**  
**Newsletter**  
**Contact us**

**Essay List**

89上學期English2A

Open all essays

Silvia Wu (488011940)	New Essay
Lily Chen (488012054)	New Essay
Bonnie Shie (488012179)	No new essay
Tina Wu (488012211)	New Essay
Kay Shih (488012328)	New Essay
Helen Wen (488012369)	New Essay
Mickey Wang (488012401)	New Essay

← Back

Fig. 2. Teacher's view of class roster with hyperlink to each student essay.

to retrieve and use that comment. Unlike the case with traditional “red-pen” corrections, with the Comment Bank it is no more time consuming for the teacher to give a clear, appropriately detailed comment than a highly abbreviated, cryptic one. Both sorts are simply retrieved from a drop-down menu. In this respect, the system encourages the teacher to provide feedback that is sufficiently clear and detailed to be useful to the student.

Crucially, the content of each teacher’s Comment Bank is entirely under the control of that teacher. A teacher can link to her Bank and add, revise, or delete stored comments at any time (though, on the suggestion of some of the teachers, we have experimented with providing a “starter kit” of prepared comments which individual teachers can then revise, add to, discard, or use “as-is”). At this stage, research is needed to understand the factors effecting how beneficial various sorts of comments are in helping students with their writing. An advantage of this system is that with it, researchers can control the crucial variables (such as the precise form and content of the teacher feedback being investigated), and it makes readily available the data needed for such research since the marked and unmarked essays are archived in forms that can be queried.<sup>3</sup> Moreover, the revised versions of an essay can be examined along side the teacher’s comments

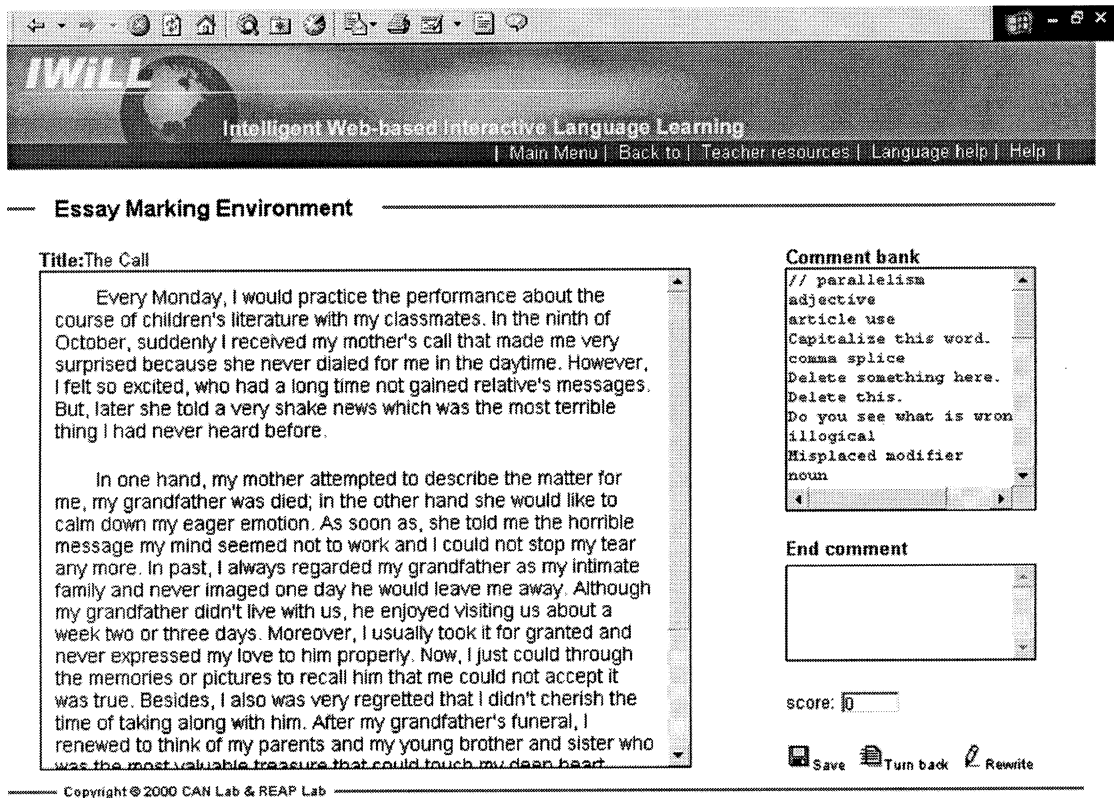


Fig. 3. Teacher’s essay correction interface.

<sup>3</sup> The research that we report on in Section 4 below takes a different tack in exploiting the teacher’s comments as a source of insight into the learners’ English difficulties.

that were given to the student on the original version of the essay, making it possible to easily track the influence of various types of teacher feedback.

Though teachers have control over the content of their own comment banks, there are advantages to having a starter kit of comments included in the bank at the outset. This set of comments can serve, for example, as an aid in teacher training. The set of comments that come bundled with the system can offer a reference for new teachers or student teachers who may have little or no exposure to the type and source of learner writing difficulties and what sorts of feedback can be offered to learners. The Comment Bank can offer them a reference point for evaluating student writing.

### 3.1.2. More on the teacher's view

Once the teacher is finished marking an essay, there are two ways to return it to the student. By selecting the appropriate "Return to Student" button, the teacher can either allow students the option of revising the returned essay or not allow the option to revise. The consequence of the teacher opting for the "not" alternative is that the student will be able to view the returned essay but not edit it. In other words, the file will be in a sense "locked." If, however, the teacher chooses to allow the option of revising, the student is provided a view of the essay which allows editing. Under this option, upon receiving an essay back from the teacher, students can toggle between the editing view where they can revise the essay on the one hand and a view that displays the

Fig. 4. Teacher's essay correction interface with comment bank drop-down.



original essay with the teacher's comments on the other. This allows the student a convenient way of referring to the teacher's suggestions as he revises. It is important to distinguish the essay correction functions of this system from the features on certain commercial word processing software. What distinguishes the annotation features offered to teachers on IWiLL is not primarily the convenience it offers to users in conveying comments on the writer's work.<sup>4</sup>

This convenience, though important, is somewhat incidental in the system design. The more important advantage is the tracking which this annotation feature supports. Specifically, since teachers' annotations are stored in separate tables and indexed to the strings in the student essay that they are intended to comment on, the annotations support a range of research and useful functions for users. Learners or teachers can retrieve all tokens of any error type that a teacher has marked and display a list of all comment types ranked in order of frequency. Moreover, the bootstrapping approach to research on learner errors which we describe in Section 4 relies crucially on the specific design of the error annotation function, specifically its integration into the online learner corpus.

A word about the scope of the teacher comments is perhaps in order here. Some teachers who have heard descriptions of this system have remarked that many of the problems with students' essays are not sentence-level grammatical or mechanical errors, but rhetorical difficulties, such as shifts in the writer's point of view or faulty argumentation. This sort of remark seems to reflect an assumption that the use of a computer will limit the types of student problems that the teacher will be able to mark or comment on. This is not the case, however. The Comment Bank is just as suitable for storing and giving comments about rhetorical concerns as for micro level difficulties such as grammar errors. Likewise, comments that a teacher composes for one-time use for a unique problem can be aimed at any level of concern under this system, from the writer's sense of audience to his spelling.

### 3.2. *The student's view*

A registered student logging onto the system is first shown a menu of links, including a link to a discussion board dedicated to the students in that class and links to helpful websites for ESL writers. To compose or turn in an essay, the student links to a page that displays a row of colored buttons, basically each button (or cluster of buttons) representing a different essay the student has written or is in the process of writing. The color of the button indicates the status of the essay, distinguishing essays that have been started but not yet turned in (red), essays that have been turned in but not yet returned from the teacher (yellow), and those that have been marked and returned from the teacher (green). If the student has previously elected to revise an essay that the teacher has returned, the revised version, as mentioned earlier, shows up on this page as a separate button directly under the button for the original version (Fig. 5). From this page, the student can opt to resume work on an unfinished essay or revision, or to submit or compose a new essay.

#### 3.2.1. *Composing and submitting essays*

To compose an essay, students can elect either to compose online by typing their essay within a designated text box on the appropriate page (Fig. 6) or to copy and paste into that box an essay composed off-line. The latter essays are imported as text files.

---

<sup>4</sup> We are grateful to an anonymous reviewer for suggesting this similarity.

From this screen, where the essay has been composed or imported, students can submit the essay to the instructor. Alternatively, they can send the essay to any number of their classmates for peer editing or commenting. They can do this by accessing a drop-down list of all of their classmates' names and selecting the names of those they wish to send the essay to for peer comments. A student who has been sent an essay from a classmate over the network in this way can view it in a display quite similar to the teacher's view and return it to the author with their comments added. The major difference is that, unlike teachers, students as peer editors do not have the Comment Bank and so must compose all of their comments from scratch. The methods of selecting portions of text for comment and for submitting comments are otherwise the same as under the teacher's view described above.

### 3.2.2. Reviewing the teacher's comments

When a student views an essay that has been marked by the teacher, the essay itself appears almost identical to the student's original, unmarked essay. None of the teacher's comments are immediately visible. The only difference in the appearance of the marked and unmarked version of an essay is that in the 'corrected' version some of the student's text shows up in blue. These are

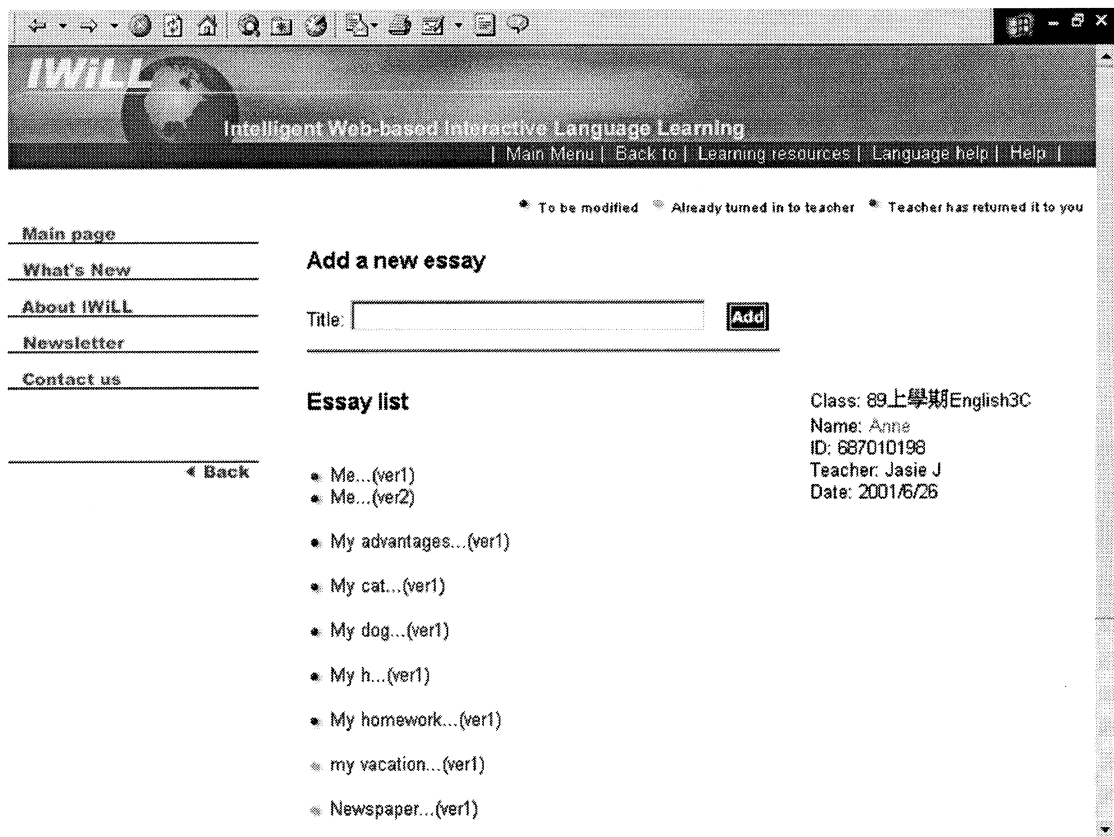


Fig. 5. Student's view of essay list with hyperlinks. Links to rewrites appear in a column under the link to the original for easy comparison.

portions of the essay that the teacher has marked for comment. To see the content of the teacher's comment, the student places the cursor on the blue text and the comment appears as a pop-up (Fig. 7).

### 3.2.3. The 'view comments' function

An important feature offered to students is a specific sort of search function which they can access through a link labeled: "Search all comments in my essays." With this function, a student can access a list of all of the comments that the teacher has marked on his essays (Fig. 8). The comments are listed in descending order of frequency as they occur in the entire set of that particular student's essays. So, for example, the display in Fig. 8 shows that this particular student received the comment "sentence fragment" more than any other comment and that the second most frequent comments were "tense shift" and word choice.

By clicking on the View button for any of these comments that appear on this display, the student retrieves a cumulative listing of all of the instances where this comment appeared in his own essays (Fig. 9). To give the minimum context that would allow the student to see the nature of the marked problem, this search function retrieves complete sentences from the student's texts even if the teacher had marked only a word or phrase or other proper subpart of that sentence for

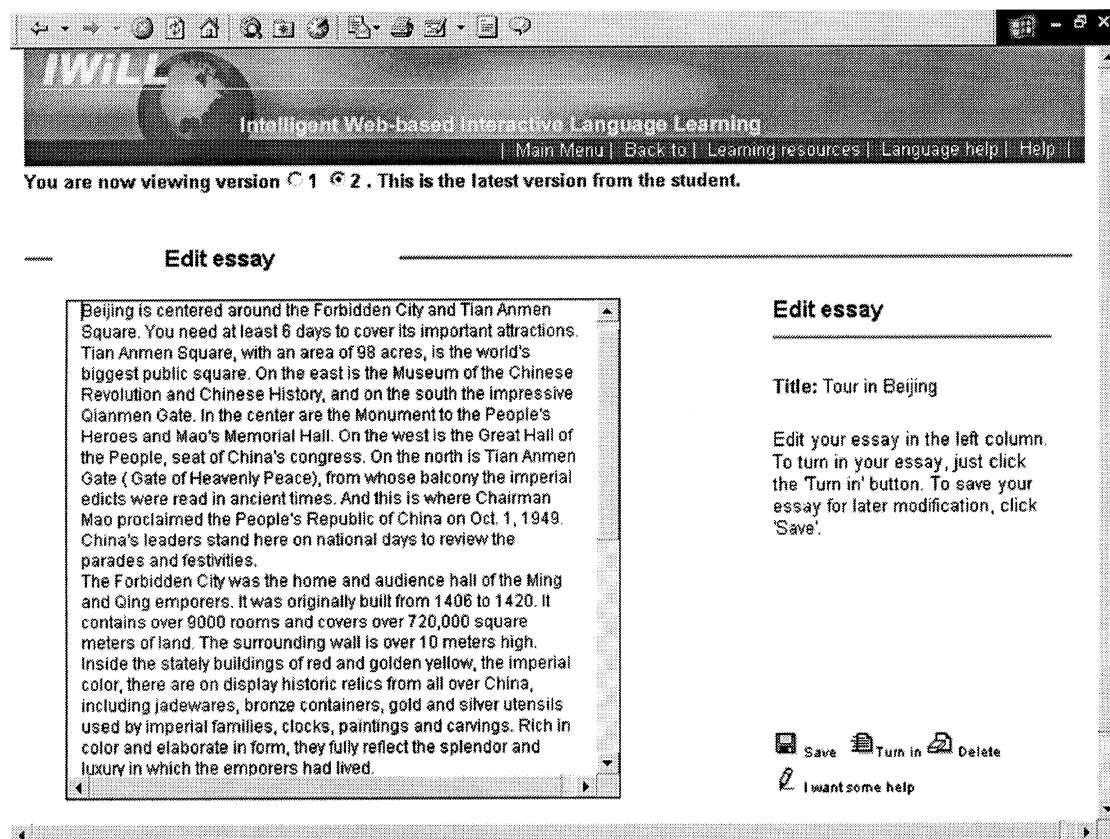


Fig. 6. Student's interface for composing, revising, and turning in essays.

comment. In instances where the teacher has marked off a chunk of text which spans a sentence boundary in the student essay, the entire text of both (or all) of those sentences is displayed. By clicking on any of the tokens that have been retrieved, the student links to the complete text of the essay from which that token was extracted, thus accessing the full context.

The assumption behind this “View Comment” feature is that in marking a student’s essay in this system, the teacher is annotating the essay with information. Moreover, these can serve as the targets of searches. Traditional “red pen” marking of hard copies of student essays can also be considered annotations of those texts, but there is a difference which our system attempts to exploit. The teacher’s annotations within our system are stored and can be retrieved in ways that information technology allows but which traditional pen-and-paper corrections do not. One consequence for the traditional method is that, even when a teacher repeatedly marks one particular type of error which appears persistently in a student’s essays over the course of, say, an academic year, that error as a specific error type may never strike the learner as salient. Each instance of that marked error appears to the student within a forest of other markings. This fact can prevent the student from noticing patterns of errors within an essay let alone patterns that persist across several essays spanning the period of a semester. When the teacher returns a “pen-marked” essay, the student may or may not even take notice of what has been marked. Add to this that over a semester the red marks come to the student as feedback on a range of distinct

The screenshot shows a web browser window with the IWILL logo and navigation links. Below the browser, the page title is "View essay". The main content area is divided into two columns. The left column contains a student essay about Beijing, with a blue font highlighting a portion of the text: "unclear" and "it isn't clear what you mean here.". The right column contains teacher comments, including the title "Tour in Beijing", the date "2001/6/26", and the comment "This is a good draft. It would be great if you watch out some details." A copyright notice "Copyright © 2000 CAN Lab & REAP Lab" is visible at the bottom.

Fig. 7. Teacher’s comments appear to student as pop-ups and the portion of student text that has been marked is indicated with a blue font.

The screenshot shows a web browser window with the IWILL logo and navigation menu. The main content area is titled "Anne's Comment List". On the left, there is a sidebar with links: Main page, What's New, About IWILL, Newsletter, and Contact us. Below these is a "Back" button. The central table lists various comment types, their frequency, and links to examples. To the right of the table is a message: "Hi, Anne! This is a list of the comments your teacher has marked on your essay(s). Click on the Link to see each example from your own writing." At the bottom right, there is a copyright notice: "Copyright © 2000 CAN Lab & REAP Lab".

Comment	Times	Link to essays
sentence fragment	10	Link
tense shift	3	Link
word choice	3	Link
subject-verb agreement	2	Link
relevance	2	Link
Reference	1	Link
source?	1	Link
comma splice	1	Link
noun	1	Link
passive	1	Link
punctuation	1	Link
unclear	1	Link

Fig. 8. Cumulative display of comments given to an individual student on all essays.

The screenshot shows the same IWILL web browser window. The main content area is titled "Essay List about 'word choice'". It contains a table with two columns: "Title" and "Sentence". The table lists five examples of essays where the teacher marked a "word choice" problem. At the bottom right, there is a copyright notice: "Copyright © 2000 CAN Lab & REAP Lab".

Title	Sentence
The Pride of Taipei	Surprise may be the dominant emotion when you look for the first time out of the wide windows of the Shin Kong Tower Observatory located in the center of Taipei
Tour in Beijing	China's leaders stand here on national days to review the parades and festivities
Tour in Beijing	You need at least 6 days to cover its important attractions
Challenges	The findings showed that students with good speaking performance were not necessarily left-brain oriented
Challenges	Regardless of the objective restrictions within which second/foreign language learner might undergo, such as learning environments or financial problems, the left and right hemispheres, exclusively, affect their success

Fig. 9. The hyperlink from the comment called 'word choice' shows this student all examples where the teacher marked a 'word choice' problem.

writing assignments separated from each other by a week or more. The chances of the learner detecting patterns of persisting difficulty fall considerably, the red marks representing simply a raw accumulation of more errors.

What the “View Comments” function provides is the opportunity for the student to see patterns of difficulty, to see in one glance a set of tokens of one type of difficulty from his own writing. Of course, what is needed here is research on the differential effects of the two approaches to providing feedback. Moreover, the effectiveness of the View Comments function will almost certainly depend not simply on the fact that the system allows searches of the essays according to teacher comments, but also upon the quality and clarity of the comments themselves. An important property designed into the system is that it can track precisely the kinds of data needed for investigating these sorts of issues.

#### 4. Some consequences for research

To reiterate one of the motivating assumptions behind the design of the IWiLL system, the teachers and students using it not only gain access to information, but they create information along the way. One way of describing our basic strategy is that we are developing ways to give users access to information that they themselves create. In the earlier section we describe how teachers create information when they mark student essays. The system treats these markings as annotations of the students’ essays and allows students (and teachers) to retrieve these annotated portions of their own writing according to the annotation tag (or, in other words, according to the comments written by the teacher). While this functionality has important implications for supporting research on writing pedagogy, in what follows, we describe some of the novel contributions for research concerning learner corpora.

##### 4.1. Learner corpora

The most basic sort of information created by the users of the system is the essays themselves that have been submitted by the students. The system has been designed to create a corpus of student essays as a byproduct of the teacher–student interaction on the system. Specifically, each essay that a student submits to his teacher over this system is, with the permission of that student, copied into a corpus of “learner English” called *English TLC*. Consequently, the corpus itself grows as the system is used by students and their teachers.

The creation and analysis of corpora of learner language data is an extremely new and promising field of research (Granger, 1998). One of the formidable obstacles in this field is a practical one of how to input the learner data. Granger (1998, p. 11) mentions three methods: (1) scanning essays from hard copies (2) keying in data manually, and (3) downloading electronic data. Granger implies that the latter refers to collecting student essays that are on disks. These methods are extremely tedious, time-consuming, expensive and the first two prone to error. Scanned texts must be painstakingly checked for scanning errors and all three methods require that metadata concerning the learner’s background and the particular essay type must be annotated by the corpus team. Our system offers another way of creating learner corpora which goes a long way toward eliminating these prohibitive drawbacks. The texts created by students enter the corpus

virtually unaffected by any intermediate steps for “inputting” them because the exact text that the student sends to the teacher over the system is copied into the corpus. Moreover, when students first register to use the system, they provide relevant metadata about their years of studying English, their gender, age, mother tongue, and the relevant fields of metadata are updated every semester. Each essay a student turns in is automatically indexed to this information and annotated with the date when that specific essay was submitted. This indexing allows for longitudinal studies of learner writing as well as cross-sectional studies that consider variables such as gender, age, or years of study. Researchers can add other fields of metadata to track other variables for specific studies.

#### 4.2. *Bootstrapping from a learner corpus*

One of the central features of the system is the relative freedom it gives to both students and teachers. Unlike many CALL systems, this platform imposes virtually no constraints on the content of the students’ written language production and of the teachers’ comments on that language output. This leads to a fundamental dilemma for the design of CALL systems: the more freedom a system offers to learners in the use of the target language, the more unwieldy the data is which the learners produce and the less able the system is to support inferences and extract knowledge about learners from that data. On the premise that such freedom for the learner is desirable, in what follows we introduce, from research on situated cognition and visual perception, the notion of *affordance* (Gibson, 1979). We apply this concept to the task of uncovering insights about learners from relatively uncontrolled target language output they produce. We illustrate the affordances offered by the system’s architecture that support a process of bootstrapping from raw language output or corpus data to potential insights into the learners’ interlanguage and gaps in their grasp of the target language.

##### 4.2.1. *Teachers’ annotations as bootstraps*

Due to the architecture of the IWiLL system, the learner corpus (English TLC) is partially annotated by the teachers who use the system to mark and return essays. We exploit this feature as a window onto the learner corpus. Typically, learner corpora are collections of text produced in the target language by learners of the target language. Such collections in themselves, however, are worth little to teachers and researchers. Their potential value lies rather in the insights such corpora can provide into the distinctive qualities of the language of the learners, into their particular difficulties in acquiring the target language, into the interlanguage or transitional grammar that lies behind the texts they produce. The corpora in themselves, however, do not yield up such insights directly. For this reason, the second great challenge beyond capturing learner corpora is annotating them in ways that make it possible for researchers to uncover these sorts of insights. One of the most fundamental aspects of learner corpus annotation is error annotation. Moreover, error annotation is one of the most tedious, subjective, time-consuming and labor-intensive aspects of corpus annotation.

We suggest that a learner corpora which has been partially annotated by teachers can support an alternative and valuable supplement to error-tagging a learner corpus by hand. The approach we develop is a process of bootstrapping which takes the teachers’ annotations as clues or bootstraps which can be exploited for more revealing pinpoint searches of the corpus. These incremental

searches can extract distinctive aspects of learner English without relying upon exhaustive error annotation. In what follows, we illustrate how we have exploited one of the most common but eclectic error types marked by teachers in their students' essays—word choice errors—to uncover novel insights into learners' second language lexical knowledge.

#### 4.2.2. Word choice and miscollocations

Word choice poses a double challenge to teachers and to researchers in second language acquisition. On the one hand, word choice is the error type which is by far the most frequently marked by teachers in the correction of students' essays, with 776 instances marked by teachers using the IWiLL platform (Table 1).

On the other hand, while word choice is by far the most common coherent class of error used by teachers, it is also perhaps the least susceptible to a unified analysis.<sup>5</sup> There could be any of a wide range of reasons that the choice of a word is inappropriate or incorrect in a particular context. In fact, information that is specific to particular words is not only one of the hardest aspects of language for learners to acquire and for teachers to teach; it is one of the most difficult aspects of linguistic knowledge to automate (Zernik, 1991). It becomes, then, one of the most important but intractable challenges to any English learning system. Let us consider in some detail how bootstrapping affordances can contribute to meeting this challenge.

Note first that due to the freedom teachers have under the IWiLL system for annotating their students' essays with feedback, the class of error tokens we obtain by retrieving every string marked with the label "word choice" borders on incoherent. Our objective, however, is to retain this freedom for learners and teachers and rely rather on the system design to yield up bootstraps that will lead eventually to the sorts of insights into learners that we are looking for. The eclectic nature of the list of tokens marked as word choice errors could be called a problem of *precision* in searching the corpus for errors. That is, in retrieving all tokens marked by teachers as word choice errors, this net that we cast also brings in errors that we would not consider to be word choice problems. Conversely, the fact that teachers mark essays selectively and certainly have left a large numbers of word choice errors untagged and have not even marked most of the essays in the corpus at all could be considered a problem of *recall*. Our approach is not to attempt to increase precision and recall of error searches directly, but rather to exploit an imperfectly tagged learner corpus and find within this data affordances for the insights that we are ultimately seeking. In this case, we decided to narrow our consideration to a specific sort of word choice error

Table 1  
Tags of the three most frequently used comments

Frequency	Error tag
776	Word choice
547	Delete this
394	Word form

<sup>5</sup> We consider error tags such as "Delete this" Or "Unclear" as not being coherent classes of errors (though useful for teachers in marking essays).



which stands as a persistent and widespread problem for learners: miscollocations, and more specifically, lexical miscollocations.<sup>6</sup>

Our first challenge in investigating miscollocations is that teachers rarely tag miscollocations as such. Therefore, the teachers' annotations do not directly uncover the error type we are interested in. Their annotations along with the system's indexing of those annotations do, however, offer crucial affordances for our purposes. First, the system can display the list of tags used by teachers and can tabulate instantly the frequency of use of each tag. Second, for each of these tags (error types) the system can display a complete list of the sentences that have been marked with that tag. With these tools, we first selected from the list of attested tags those which looked like they may have been used to mark word choice errors, specifically to mark miscollocations. The list of relevant tags appears in Table 2. From the tokens of errors extracted with this list, we were then able to isolate the tokens which were miscollocations.

By retrieving the strings tagged with these labels, 1100 tokens of word choice errors were extracted from the corpus. Among these tokens we found 177 miscollocations. Of these, 130 tokens had been tagged by teachers directly while the other 47 tokens of the same error were extracted by keyword searches of the learner corpus based upon these. In other words, from a corpus of over one million words of learner English, unrestricted error tagging by teachers afforded us an extremely narrow subset of learner writing which these teachers determined to be word choice difficulties or of some type likely to include word choice errors. This narrowly circumscribed set of errors in turn made it possible for us to quite easily uncover a narrower subset: miscollocations. Among the 177 miscollocations, there are four subtypes, classified by part-of speech, as shown in Table 3.

Table 2  
List of teachers' error tags likely to contain word choice errors

- 
1. Word choice
  2. Wrong word
  3. Wrong verb
  4. Problematic phrases
  5. Unclear
  6. Replace
  7. Correction
  8. Verb
  9. Change this expression
  10. Change this word
  11. Improper usages of words
  12. Inappropriate word or phrase (In Chinese)
  13. Word combination (In Chinese)
- 

<sup>6</sup> Miscollocations are errors in word choice where two (or more) words are used together in a way which simply violates the conventions of cooccurrence for that language, even though the miscollocation may make sense semantically. For example, in English we use the verb "take" rather than "eat" with the object "medicine." This is simply a convention of English usage. See Sinclair (1991) for detailed discussions of collocation and Gitsaki and Daigaku (1997) and Lewis (2000) on miscollocation in language learning. In the case of lexical miscollocations, the words involved are lexical or content words such as nouns, verbs, adjectives, and adverbs rather than grammatical or function words such as prepositions, articles, or conjunctions.

This list reveals a striking result. The overwhelming majority of miscollocations are VN miscollocations (145 out of 177). This result offered an empirical basis for then narrowing our investigation specifically to V N miscollocations, since this profile suggested that these constitute the most widespread difficulty. A cursory look at the 145 tokens of VN errors revealed another striking fact about the learners miscollocations: in all but three of the 145 VN miscollocations, it is the V(erb) which is incorrect, not the N(oun). This simple fact provides an extremely important clue to the nature of learners difficulties with this most common of miscollocation errors. In turn, it has led us into detailed linguistic analysis of the semantic and syntactic features of the misused verbs and the nature of the learners' representations of these verbs which lead to their errors (Wible & Liu, 2001). Further, the results have led to fruitful experiments in automating the correction of VN miscollocations which have performed 40% better than a thesaurus in locating the correct verb to replace miscollocate verbs (Wible, Kuo, Liu, & Tsao 2001).

#### 4.3. Corpus research results and online help

One of the fundamentals of our design philosophy has been to enhance the interdependence of various types of users. The last feature of the system to be described illustrates one means of exploiting this interdependence to make the system self-enhancing by design. Specifically, researchers are not only able to search the corpus of essays collected from learners. The results of the researchers' analyses of learner difficulties can be translated into the content of an active online help function for those learners. The system includes an authoring environment for content administrators (ICPs) where they simply indicate what string of text in a learners essay should trigger help, and then write the content of the help which should be displayed for that particular string. Research on the learner corpus has revealed, for example, that the word 'ever' was misused by learners in 25% of the cases where it appeared in their essays. Further analysis attributed this to negative transfer in which learners associated the English expression 'ever' with a Chinese counterpart expression (*cheng jing*). These two expressions while overlapping in use and meaning, diverge in important ways, and it is precisely in these diverging respects where students misused the English expression.<sup>7</sup>

Table 3  
Miscollocations by categorized by part-of-speech and frequency

Miscollocation type	Frequency
V N	145
Adj N	25
V Adv	5
Adv Adj	2
Total	177

<sup>7</sup> 'Ever' is in most cases what linguists term a polarity item, permitted in negative statements, interrogatives, and conditionals but rarely in simple affirmative declaratives. The Chinese counterpart *cheng jing* is not a polarity item and so is free of these restrictions. Hence a common error in the use of 'ever' is a case like "My teacher has ever told me to read more English novels."

Based on this linguistic analysis, the authoring environment for online help was used to design advice concerning the word ‘ever’ addressing precisely the difficulties it poses for Chinese learners. When learners request general help on an essay, the help function actively detects instances of ‘ever’, highlights them and creates a link to this advice.

## 5. Conclusion

The underlying goal of the project described above has been not only to create an online writing environment that connects teachers and students by way of a user-friendly interface, but also to provide ways to exploit the valuable data that is created when the environment is used. The learners’ essays themselves are stored in a growing corpus of ESL language production. The comments that teachers append to the particular segments of the learners’ texts in the course of essay correction are treated as annotations of those texts, which can be searched and retrieved. The teacher annotations also can be exploited as bootstraps for researchers, allowing them to probe learner errors without exhaustively annotating the corpus first. An authoring environment for online help permits content administrators to turn interlanguage research results into highly specific help concerning attested difficulties which traditional language education has neglected. It is hoped that increasingly sophisticated and dynamic manipulations of these sorts of data will lead to the delivery of evermore useful and useable information to learners, teachers, and researchers both online and off.

## References

- Chang, H.-C. (2000) Web-based EFL learning and teaching (unpublished manuscript).
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Gitsaki, C., & Daigaku, N.S. (1997) “English Collocations and their place in the EFL classroom,” paper presented during JALT ‘97, 9–12, October 1997. Hamamatsu ACT City, Japan.
- Granger, S. (1998) The computer learner corpus: a versatile new source of data in SLA research. In Sylviane Granger, *Learner English on computer*. London: Longman.
- Krashen, S. (1994). The input hypothesis and its rivals’. In N. Ellis (Ed.), *Implicit and explicit learning of languages*. London: Academic Press.
- Lewis, M. (2000). *Teaching collocation*. Hove, England: Language Teaching Publications.
- Meunier, F. (1998). Computer tools for the analysis of learner corpora. In S. Granger (Ed.), *Learner English on computer*. London: Longman.
- Sinclair, J. (1991). *Corpus, concordance, collocation*. Oxford: Oxford University Press.
- Wible, D., Kuo, C.-H., Tsao, N.-L., Liu, A., Sung, L.-C. and Chio, C.-L. “Putting learners first: creating an integrated online learning environment for EFL writers,” paper presented at ICEE2000, Taipei, August 2000.
- Wible, D., Kuo, C.-H., Tsao, N.-L., & Liu, A. (2001). An online writing platform for language learners. *Journal of Universal Computer Science*, 7(3), 278–289.
- Wible, D., Kuo, C.-H., Liu, A. and Tsao, N.-L. (2001) On finding a computable path from miscollocations to collocations,” unpublished paper, Tamkang University, Taipei, Taiwan.
- Wible, D. and Liu, A. (2001) A syntax-lexical semantics interface analysis of collocation errors,” unpublished manuscript to be presented at PacSLRF 2001, Hawaii.
- Zernik, U. (1991). *Lexical acquisition*. Hillsdale, NJ: Lawrence Erlbaum.