

I-PETER: Modelling personalised diagnosis and material selection for an on-line English course

Timothy Read¹, Elena Bárcena², Beatriz Barros¹ & Felisa Verdejo¹

Depts.: ¹Lenguajes y Sistemas Informáticos, ²Filologías Extranjeras y sus Lingüísticas

^{1,2}UNED, Madrid, Spain

¹{tread, felisa, bbarros}@lsi.uned.es, ²mbarcena@flog.uned.es

Abstract. In this paper the underlying knowledge model and architecture of I-PETER (Intelligent Personalised English Tutoring EnviRonment) are presented. This system has been designed for the on-line distance learning of English where too many students restrict the teacher's possibilities to provide individualised guidance. I-PETER is made up of four domain models that represent linguistic and didactic knowledge: the conceptual framework related to linguistic levels and knowledge stages, and the educational content and study strategies. The student model represents the knowledge that the student has learnt, the study strategies, and his/her profile. A student's command of English is evaluated by interpreting his/her performance on specific linguistic units in terms of three related criteria, rather than by a general linguistic competence ranking. Evaluation consists of a diagnostic task model which assesses student performance, taking the form of a Bayesian network, and a selection mechanism that proposes appropriate materials and study strategies.

1 Introduction

As in other educational areas, interactive on-line environments for learning English enable students to work without having the teacher present. In a distance learning context they represent an advance over their book-based counterpart, since as well as their inherently interactive nature, they enable both the teacher to add new content as the course progresses according to the students' general needs (e.g., summaries, extra exercises, and mock exams to help them prepare for the final exam), and the students to make use of the communication facilities (e.g., e-mail, net-news) to contact the teacher or peers for help. For reasons like these, courses based upon interactive on-line environments are being adopted in a wide range of distance learning institutions.

A problem present with all taught courses is that, as the number of students grows, it becomes progressively harder for the teacher to maintain control of the overall learning process of the group: follow the progress of individual students, identify their difficulties, provide help and guidance accordingly, and introduce modifications in the way in which the material is being studied to adapt to individual needs. Due to the typically large number of students in on-line distance learning courses (in the case of the UNED, the various English courses can have between 350 and 15,000 students), and the additional communication difficulties (e.g., there is no face-to-face

contact or non-verbal feedback), the teacher becomes more of an administrator, being able to deal with only a small number of problems that certain students present via e-mail or telephone calls (not always the ones that really require the help!). A related problem present in all courses is that they do not take into account the profile, learning goals, and other features and needs of individual students.

The research presented in this paper has three objectives. Firstly, the structural decomposition of both the linguistic domain knowledge and student linguistic competence in such a way as to capture and represent the underlying conceptual content and student model. This is necessary in order to design an intelligent personalised English tutoring environment for use in a distance learning context where, as mentioned above, a very high number of students limits teacher interaction drastically. Secondly, the analysis of the results of student interactions with the system, in order to encounter study strategies (considered in terms of conceptual units and materials) which are particularly effective or ineffective, and that can be used appropriately by the system to improve student progress and learning. Thirdly, based partly on the previous objective, the automatic detection (or prediction) of students who are having (or may have) difficulties with the material in order to help the teacher know who needs help. This last objective is still work in progress and, as such, is not discussed here.

Throughout the three decades of CALL¹ research, different approaches have been developed, from the earliest vocabulary and grammar trainers to multimedia and Web-based workbenches that reflect the current interest in communicative intent and functions within Linguistics and eclectic pedagogic frameworks. Combined with these are research strategies taken from different areas, including Computational Linguistics (e.g., parsing student input and/or generating natural language), Artificial Intelligence (e.g., knowledge-based systems, student modeling), Cognitive Psychology, Psycholinguistics, Human-Computer Interaction, etc. [2], [9]. While progress has been made in the area of interactive on-line learning, the authors' experience is that it is currently unfeasible to design the perfect tutoring system. It is possible however, to solve a subset of problems present in a particular educational domain. In this paper a system called I-PETER², based to some extent upon the previously mentioned disciplines, is presented as a proposal for representing and organizing the domain knowledge and student model necessary for personalized English distance learning.

This paper starts with a description of the different pedagogic parameters that have been taken into account for the design of I-PETER. Secondly, the linguistic and didactic domain modelling undertaken for this system is presented together with details of the student model and the tasks of diagnosis and selection. Thirdly, the functionality of the system and its underlying architecture are briefly presented. Fourthly and finally, some conclusions are offered together with plans for future work.

¹ CALL: Computer Assisted Language Learning

² I-PETER: *Intelligent Personalised English Tutoring EnviRonment*

2 The pedagogic foundations of I-PETER

The design of I-PETER takes into consideration such pedagogic issues as: the influence of mother tongue in second language acquisition, the need to treat the tuition of sublanguages separately, the existence of personal preferences for language learning vs. acquisition, and the relevance of error analysis. However, there is one particular pedagogic issue that is crucial for the overall design of this system: student knowledge modelling.

On-line language courses are typically organised according to an idealised degree of quasi-native ‘linguistic competence’ that a student must reach at the end of his/her studies. To achieve this, the ‘mass of knowledge’ that the student needs to command is typically categorized into stages such as: stage 1... stage n, beginners’, intermediate, advanced, etc. These classifications reflect attempts to group together students with roughly the same level of English knowledge, and they are deeply engrained into the educational system, as well as being a motivational factor for students. However, while it is relatively easy to organize students into, for example, three, four, or even ten stages of linguistic knowledge, the reality is that the subset of English known by students that have been assigned a given stage can vary greatly within the different linguistic levels (some may have a considerable vocabulary due to extensive reading in English; others may be fluent in their production for having been immersed in the speaking community although with a poor knowledge of grammar, etc.). Inevitably, the potential progress of some students within a given group can be slowed (e.g., there are some topics that they will not need to practice), while others will fail to fully acquire certain language items due to insufficient underlying knowledge. Furthermore, the decision to move students up to a higher stage is typically based upon some sort of average mark made up from the results of test questions which measure a few linguistic skills. This practice inevitably leads to ‘holes’ in the students’ knowledge. In the authors’ experience, a more useful approach for the classification of student knowledge and progress relates three dimensions, rather than the general notion of ‘degree of linguistic competence’. These interacting dimensions are defined to be: ‘linguistic level’, ‘knowledge stage’, and ‘learning phase’, as shown in figure 1.




Fig. 1. Relationship between the three dimensions that make up the problem domain

Firstly, knowledge stage is the term used to classify the students’ linguistic knowledge (as above). The difference lies in the scope of application, because it is applied to the narrower fields of discrete linguistic conceptual units, and only from there can it be generalised to a linguistic level and the student’s general competence. In this system the stages are: beginners’, lower intermediate, upper intermediate, and

advanced. The number of stages chosen could have been larger to reflect more subtle divisions, but such fine granularity is unnecessary in this multidimensional knowledge classification.

Secondly, the linguistic level corresponds approximately to the generally accepted distinction (with some variance) amongst linguists regarding the structural composition of language. These levels are: lexicon (the level of words, their orthography and meaning, either in isolation, in locutions, and in the context of other elements), grammar (the level of morphology and syntax at phrasal, clausal and sentential level) and discourse (the level of text and composition; a supra-sentential level). I-PETER does not currently consider the oral aspect of the language (phonetics and phonology).

Thirdly, the learning phase corresponds to the extent to which the knowledge has been internalised by the student. This knowledge is made up of a set of concepts (e.g., verbs, simple sentences; see figure 3), and sub-concepts (e.g., for verbs: conjugation, internal composition, collocations, complementation), where each sub-concept is in itself a set of items (e.g., for verbal conjugation: simple present, simple past; see figure 4), and sub-items (e.g., for simple past: regular and irregular verbs). Conventional single topic-based exercises are typically used to help the student learn a particular linguistic item, and its subsequent evaluation, to check that it has been understood (learnt). However, they can be misleading because mechanical attentive practice of a particular item does not ensure that it has been really interiorised (acquired), so that the student is ready to use it correctly (and creatively) in other (non-attentive) contexts [8]. Furthermore, even though an item appears to have been acquired, it may be the case that this is a temporal phenomenon, and that it is subsequently forgotten due to lack of sufficient usage immediately after its acquisition. I-PETER attempts to avoid the dangers of monitored practice and short-term memory learning by distinguishing two learning phases: mechanical reproduction and non-attentive application, and using multiple topic-based exercises that provide extra practice and ‘secretly’ test previously learnt items as well as the current item under study.

3 Knowledge models used in the system

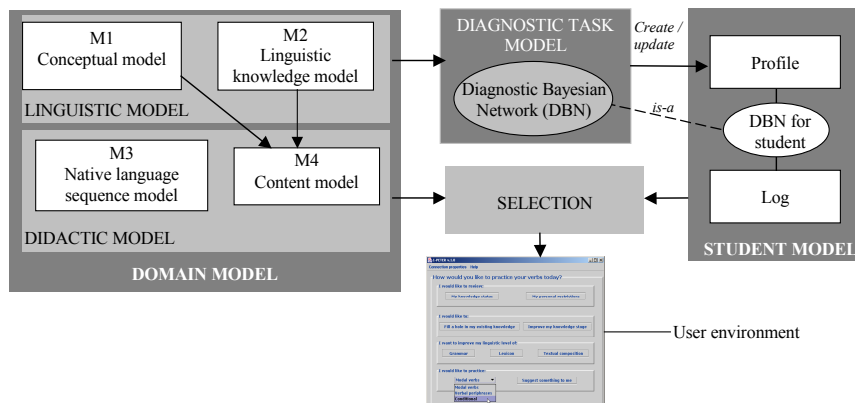


Fig. 2. Interaction of the knowledge models within I-PETER

I-PETER is made up of a set of knowledge models that include domain and student knowledge, the interactions between which can be seen in figure 2. The **domain model** of the problem is made up of two linguistic models (M1 and M2) together with two didactic models (M3 and M4):

M1: The **concepts** that make up the knowledge domain for learning English and their relation to the linguistic levels. In this model, as can be seen in figure 3, the ‘mass of knowledge’ to be learnt has been organized into a set of concepts. These have been established together with their percentage correspondences at each linguistic level, to reflect the type of knowledge that the student is progressively gaining within the field of Linguistics as s/he works on a given topic (for example: learning a list of the irregular verbs is both a vocabulary and grammar task; certain conjunctions are used both to join clauses [grammar level in this system] and sentences [text level]). Furthermore, by having concepts linked to the linguistic levels, more precise

Concepts	Linguistic levels %		
	Grammatical	Lexical	Textual
Noun	50	50	-
Determinants	80	20	-
Pronouns	80	20	-
Adjectives	50	50	-
Adverbs	65	35	-
Verbs	50	50	-
Prepositions	80	20	-
Conjunctions	80	10	10
Simple sentences	100	-	-
Complex sentences	80	-	20
Reading	-	80	20
Discourse	20	-	80

Fig. 3. Principal concepts for learning English and their relation to the linguistic levels (**M1**)

information will be available regarding the knowledge and progress of each student.

M2: The **linguistic knowledge** that relates each sub-concept to a set of items and sub-items for each stage at which they are typically taught³. In this system, a concept represents a general category of linguistic knowledge which, in turn, is made up of sub-categories of actual study elements or units (items). In figure 4 part of this knowledge model is shown for the

Sub-concept	Knowledge stage	Item	Sub-item
Verb conjugation	Beginners'	Simple present	Lexical verbs
			Primary verbs
		Present continuous	
	Lower-intermediate	Future shall/will	
		Future be going to-	
		Simple past	Regular verbs
			Irregular verbs I
		Present perfect	Regular verbs
			Irregular verbs I
		Past perfect	Regular verbs
			Irregular verbs I
		Modal verbs I	
		Conditional	
	Upper-intermediate	Simple past	Irregular verbs II
		Past continuous	
		Present perfect continuous	
		Present perfect	Irregular verbs II
		Past perfect	Irregular verbs II
		Future perfect	
		Conditional perfect	
		Modal verbs II	
		Past perfect continuous	
	Advanced	Future continuous	
		Future perfect continuous	
		Future be going to- continuous	
		Conditional continuous	
		Conditional perfect continuous	

Fig. 4. The linguistic knowledge of this domain for the sub-concept verb conjugation, relating items and sub-items to the student's knowledge stage (**part of M2**)

³ This is the general structure of linguistic content found in the majority of language courses and text books prepared without taking into account native language effects and interferences or learning preferences and restrictions.

concept verb and its sub-concept verb conjugation (including mood, tense, aspect, person, and number). The items and sub-items of this sub-concept are mapped to the four knowledge stages represented in this system. It is recommended, for example, that a typical advanced student studies such verb tenses as future perfect continuous and conditional continuous.

M3: The native language sequence model. This model represents instructional strategies in the form of the set of sequences of conceptual units, materials, and exercises most appropriate depending upon the speaker's own native language which, in the authors' experience, should be a determining factor in syllabus design.

M4: The educational content: theoretical explanations, examples, and exercises for this domain. Different types of content are stored to offer students a wide range of learning options and practice depending, among other things, upon previous interactions and results. The content of this model is classified in terms of these criteria to enable correct selection and ensure efficient retrieval. The templates used to declare the structure and content of the didactic material in this system are shown in figure 5.

CONCEPTUAL UNITS: Concept / Sub-concept / Item / Sub-item				
KNOWLEDGE STAGE				
THEORETICAL EXPLANATION				
ILLUSTRATIVE EXAMPLES				
TYPE OF EXERCISE				
CONCEPTUAL UNITS TESTED: Mechanical / Non-attentive [list of conceptual material that can be tested with this exercise and the type of test]				
DETAILS OF THE EXERCISE				
CORRECT ANSWER		EFFECT: (item/sub-item, % mod.)		EFFECT: (item/sub-item, % mod.)
INCORRECT ANSWERS:				
a	EXPLANATION a	EFFECT: (item/sub-item, % err.)	EFFECT: (item/sub-item, % err.)	ERROR TYPE: f, s or u
b	EXPLANATION b	EFFECT: (item/sub-item, % err.)	EFFECT: (item/sub-item, % err.)	ERROR TYPE: f, s or u
c	EXPLANATION c	EFFECT: (item/sub-item, % err.)	EFFECT: (item/sub-item, % err.)	ERROR TYPE: f, s or u

Fig. 5. The frame for the content model (M4)

The material in the knowledge model M4 is made up of two types of elements: firstly, a set of conceptual units (concepts, sub-concepts, items, and sub-items), together with their theoretical explanations and illustrative examples; and secondly, a set of related exercises. As can be seen in figure 5, these include information about the types of units they test, either mechanically or non-attentively. Since open production exercises (where the student is free to write) are currently beyond the scope of tutoring systems (due to the inherent intractability of natural language) [5], this system uses three generally acknowledged types of closed exercises: 'multiple-choice', 'filling in the gaps', and 'modifying the form and/or order of a word or sequence of words' [7]. The system is not only capable of interpreting the correct answers but also the erro-

PERSONAL INFORMATION			
LOGIN		PASSWORD	
NATIVE LANGUAGE: Romance / Germanic / Slavic			
TYPE OF ENGLISH: General / Scientific / Technical / Business / Legal / Formal			
OBJECTIVE: Fill knowledge gap / Improve knowledge stage / Improve linguistic level / Practise conceptual unit			
INITIATIVE: Student / System / Mixed			
DEGREE OF THEORETICAL EXPLANATION: Standard / Minimum			
OTHER PREFERENCES / RESTRICTIONS			
KNOWLEDGE STAGE:	LINGUISTIC LEVELS:	CONCEPTUAL UNITS:	LEARNING PHASE:
Beginners'	Lexical Grammatical Textual	Concept	Mechanical reproduction Non-attentive application
Lower-intermediate		Sub-concept	
Upper-intermediate		Item	
Advanced		Sub-item	

Fig. 6. The frame for the student profile

neous ones. It distinguishes between three types of error: formal (surface mistakes in the construction of a word or a larger linguistic unit, including spelling mistakes), semantic (mistakes of the meaning of a word or a larger linguistic unit), and usage (mistakes of adequacy related to communicative context), and establishes a relation between each of these and the most likely knowledge stage at which they occur.

The **student model**, the structure of which can be seen in figure 6, stores the information that pertains to each student, such as his/her preferences and restrictions (e.g., native language, learning goals) and the knowledge that the student has learnt. This profile is another determining factor for the selection of the materials within the didactic model that are most appropriate for him/her. It can be accessed and partially modified by the student at any time. There is also a log (or record) of the student activities and interactions with the system that represents the way in which the different materials within the system have been studied (e.g., sequence, time, results).

The **diagnostic task model** represents how the teacher evaluates the student, and it is implemented as a Bayesian network [11]. This approach has been selected since it appears that this probabilistic mechanism, that combines evidence to form an overall probability, is closest to the way a real teacher heuristically assesses a student [1], [3], [10] both through correct performance and error identification and analysis.

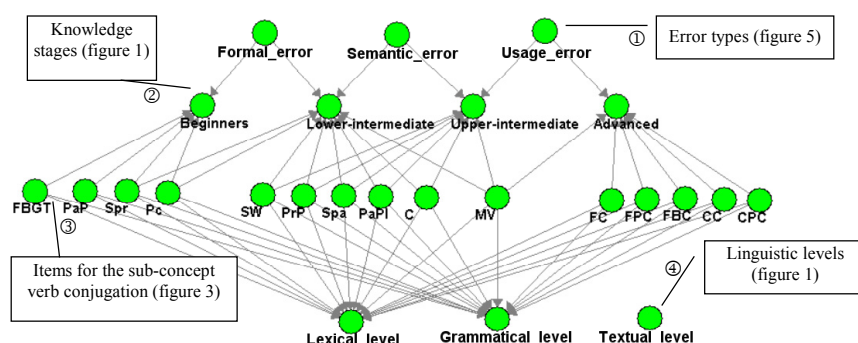


Fig. 7. Bayesian network for the diagnosis of student errors

This network is used to **diagnose** the presence of holes in the linguistic knowledge of the student in terms of the degree of membership of a particular knowledge stage and linguistic level. The domain models M1 and M2 (see figure 2) are implicitly represented in the structure of the network. As figure 7 shows, for the sub-concept verb conjugation⁴, this network is made up of four distinct types of nodes. The items that make up the sub-concept are represented as the row of nodes marked as ③ in the figure and, as can be seen, influence the nodes marked as ② (knowledge stages) and ④ (linguistic levels). Furthermore, there is another row of nodes (marked as ① in the figure), which represent the types of errors committed and influence the nodes marked as ②. In this part of the overall network, the node that corresponds to the

⁴ Only the nodes that correspond to the sub-concept verb conjugation are included in this figure for the sake of legibility. It should be noted that all the concepts represented in figure 3 would be included in the complete network.

textual level is not connected to any of the sub-concept nodes since the concept verb has no direct relation to it, as reflected in figure 3.

The student profile is created at the beginning of the course by preference selection together with an optional stage test, based upon the results generated by the diagnostic task model. Each time the student works in the environment, his/her actions and results are added to the log, and the diagnostic task model updates the student profile accordingly. Subsequently, the **selection** of appropriate material is undertaken in terms of the information contained in the profile. Firstly, a selection criterion is established by combining this information with the M3 sequence model. Secondly, a subset of the most appropriate material is proposed for the student by the application of this criterion to the M4 content model. This selection criterion also takes into account heuristic information such as the time that has gone past since the student last worked with this type of material and any other didactic criteria specified by the teacher.

4 System functionality and architecture

The main client interface to I-PETER can be seen on the right hand side of figure 8 (marked as ①). The student uses this interface to express and review his/her preferences and restrictions, choose the type of material that s/he wants to study, or simply inspect the information that the system has about his/her progress (presented visually as a set of points within the cube of figure 1 for each concept, sub-concept, etc., as selected), a feature aimed at involving and motivating the student [4].

As can be seen in figure 8 (marked as ②), the learning process can be undertaken in three ways: trying to improve or fill holes in existing knowledge stages, improve a particular knowledge stage or linguistic level, or practice specific items and sub-items, although it is also possible to let the system take the initiative. In the example in this figure, once the student selects that s/he wants to practice conditional verbs, the window shown on the left side of the figure opens, presenting the student with the theoretical explanation of the linguistic item together with a set of selected exercises (marked as ③ in the figure) that the student can use to test his/her understanding. Once the student answers a question, by selecting the answer thought to be correct, feedback can be requested, where the system presents information about the correct answer together with an explanation for each erroneous case. It should be noted, as an example of what has been discussed previously in this paper regarding the learning phases, that exercise 1 shown in the figure is interpreted by the system as a mechanical reproduction test for conditional verbs and also a non-attentive test for simple past verbs, thereby assessing also how well the student has acquired the latter.

The architecture of the system is presented below⁵ in figure 9. As can be seen, it has been developed using a JNLP-based (Java Network Loading Protocol) client-server model for three reasons. Firstly, to limit potential server load problems given the large number of students, and to make use of the processing power that is typically wasted on a standard client PC used as a Web browser platform. Secondly, to

⁵ The sequence Bayesian network is included in the architecture diagram for the sake of completeness, even though, as noted in the following section, it is still under development and, as such, not detailed in this paper.

permit a more sophisticated user environment than that which is possible in a standard HTML-based Web interface (without the use of a Java applet interface, which in itself has problems – see the third reason). Thirdly and finally, JNLP enables Java applications to be downloaded (the first time they are used) and run locally from a client computer. The advantage of JNLP over applets is that, once downloaded to the client, the next time that the user wants to run the application it is not necessary to download it again, and if there are any differences between the version on the client machine and the server, only those parts of the application that have been changed are downloaded, thereby maintaining the user version up to date.

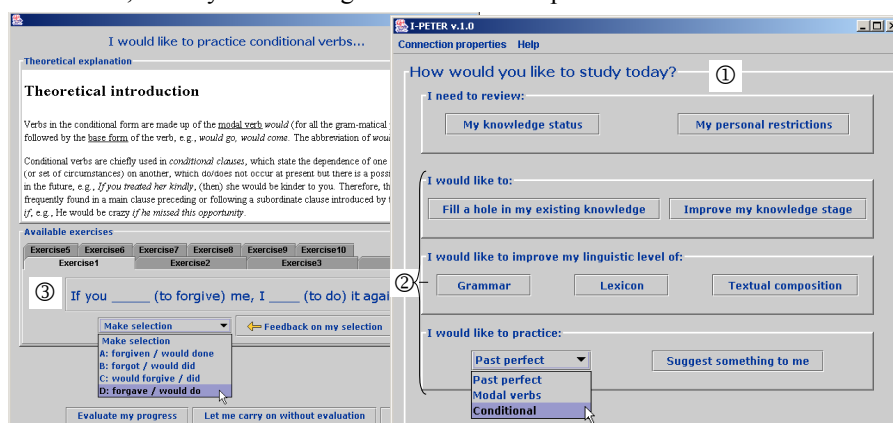


Fig. 8. Two parts of the user interface of I-PETER

5 Conclusions and future work

In this paper I-PETER has been presented, a proposal for organizing and representing the knowledge necessary for English distance learning in a personalized tutoring system. The traditional division of linguistic knowledge into simple levels has been extended to distinguish between three learning dimensions: the knowledge stage, the linguistic level, and the learning phase of the student. The domain knowledge has been structured into two parts: a linguistic model which relates units of knowledge to linguistic levels and knowledge stages, and a didactic model that contains course materials and a set of appropriate study sequences according to the student's native language. The diagnostic task model used here is a Bayesian network, which was selected because it enables the real-world heuristic evaluation of student ability to be modelled in an explicit and plausible manner. This work represents an original contribution with respect to other previous systems [7], since the three dimensional classification model provides a fine-grained representation of the knowledge levels and needs of a student beyond the general and vague concept of linguistic competence typically used. The system uses this diagnostic model to interpret student performance and subsequently selects and organises the material for him/her.

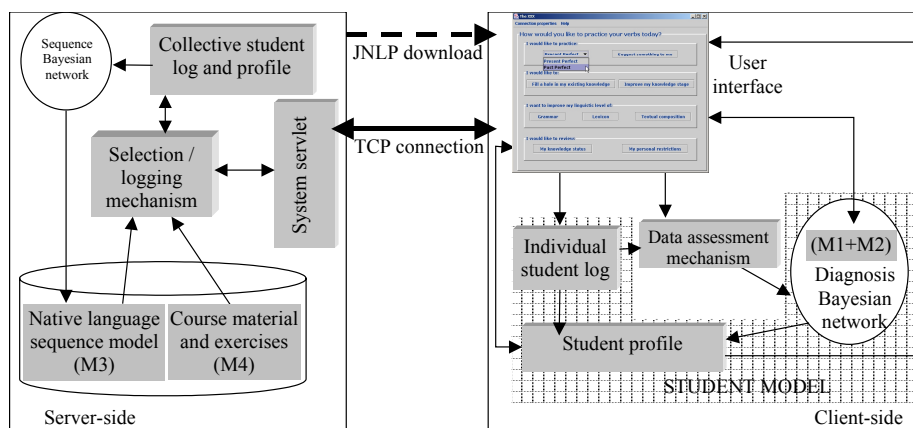


Fig. 9. The architecture of I-PETER

In the ongoing development of this system, work is being undertaken to explore the way in which a second Bayesian network could be used to detect particular domain features, such as effective learning sequences (to be incorporated in M3) and potential problems that students might encounter with the system. This work looks promising, but no conclusions can be drawn yet because more student data is required.

6 References

1. Bunt, A. & Conati, C. "Assesing Effective Exploration in Open Learning Environments using Bayesian Networks". *ITS'2002* (2002)
2. Chapelle, C. "CALL in the year 2000: Still in search of research paradigms?". *Language Learning and Technology*, vol. 1, no. 1 (1997)
3. Conati C., VanLehn K. "POLA: a Student Modeling Framework for Probabilistic On-Line Assessment of Problem Solving Performance". In *Proc. of UM-96*, Univ. Hawaii (1996)
4. Dimitrova, V. & Dicheva, D. "'Who is who': the roles in an intelligent system for foreign language terminology learning". *British Journal of Edu. Technology*, vol. 29, no. 1 (1998)
5. Harrell, W. "Language learning at a distance via computer". *International Journal of Instructional Media*, vol. 23, no. 3 (1999)
6. Heift, T. & Nicholson, D. "Web Delivery of Adaptive and Interactive Language Tutoring". *International Journal of Artificial Intelligence in Education*, vol. 12 (to appear in 2002)
7. Kohn, K. "Distributive language learning in a computer-based multilingual communication environment". In H. Jung & R. Vanderplank (eds.) *Barriers and Bridges: Media Technology in Language Learning*. Frankfurt: Peter Lang (1994)
8. Krashen, S.D. *Second Language Acquisition and Second Language Learning*. London: Prentice-Hall (1988)
9. Levy, M. *Computer-Assisted Language Learning*. Oxford: Clarendon Press (1997)
10. Millán, E. *Sistema bayesiano para modelado del alumno*. Ph.D. Thesis, U. Málaga (2000)
11. Pearl, J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan-Kaufmann (1988)