Chapter 14
Computational Aspects of the Intelligent Tutoring System MetaTutor

Mihai Lintean
The University of Memphis, USA

Vasile Rus
The University of Memphis, USA

Zhiqiang Cai
The University of Memphis, USA

Amy Witherspoon-Johnson
Arizona State University, USA

Arthur C. Graesser
The University of Memphis, USA

Roger Azevedo
McGill University, Canada

ABSTRACT

We present in this chapter the architecture of the intelligent tutoring system MetaTutor that trains students to use metacognitive strategies while learning about complex science topics. The emphasis of this chapter is on the natural language components. In particular, we present in detail the natural language input assessment component used to detect students' mental models during prior knowledge activation, a metacognitive strategy, and the micro-dialogue component used during sub-goal generation, another metacognitive strategy in MetaTutor. Sub-goal generation involves sub-goal assessment and feedback provided by the system. For mental model detection from prior knowledge activation paragraphs, we have experimented with three benchmark methods and six machine learning algorithms. Bayes Nets, in combination with a word-weighting method, provided the best accuracy (76.31%) and best human-computer agreement scores (kappa=0.63). For sub-goal assessment and feedback, a taxonomy-driven micro-dialogue mechanism yields very good to excellent human-computer agreement scores for sub-goal assessment (average kappa=0.77).

DOI: 10.4018/978-1-60960-741-8.ch014

INTRODUCTION

We describe in this chapter the architecture of the intelligent tutoring system MetaTutor with an emphasis on two components that rely on natural language processing (NLP) techniques: (1) *detection of students' mental models* during prior knowledge activation (PKA), a metacognitive strategy, based on student-generated PKA paragraphs, and (2) the micro-dialogue component which handles *sub-goal assessment and feedback generation* during sub-goal generation (SG), another metacognitive strategy in MetaTutor.

The current MetaTutor is a complex system that consists of nine major logical components: pre-planning, planning, student model, multi-modal interface (includes agents), feedback, scaffolding, assessment, authoring, and a system manager that coordinates the activity of all components. We present details about the role of each of these components and how they are implemented with various underlying technologies including dialogue processing, machine learning methods, and animated agents technology. We will describe in-depth the two NLP-based tasks of mental model detection and sub-goal generation which are part of the mental model and planning modules, respectively.

During prior knowledge activation, which occurs at the beginning of a student-system session, students are asked to write a paragraph describing their prior knowledge with respect to the learning goal. The task is to infer from these PKA paragraphs the students' initial *mental model* that best characterizes their level of understanding of the subject matter. We regard this problem as a text categorization problem. The general approach is to combine textual features with supervised machine learning algorithms to automatically derive classifiers from expert-annotated data. The parameters of the classifiers were derived using six different algorithms: naive Bayes (NB), Bayes Nets (BNets), Support Vector Machines (SVM), Logistic Regression (LR), and two variants of Decision Trees (J48 and J48graft, an improved version of J48). These algorithms were chosen because of their diversity in terms of patterns in the data they are most suited for. For instance, naive Bayes are best for problems where independent assumptions can be made among the features describing the data. The diversity of the selected learning algorithms allows us to cover a wide range of patterns that may be hidden in the data.

The role of the sub-goal generation strategy in MetaTutor is to have students split the overall learning goal (e.g. *learn about the human circulatory system*) into smaller learning units called *sub-goals*. The sub-goals must be specified at the ideal level of specification (i.e. not too broad/general or too narrow/specific). If student-generated sub-goals are too specific or too general the system must provide appropriate feedback in natural language such that students will be able to re-state the sub-goal in a form closer, if not identical, to the ideal form. The system uses a set of ideal sub-goals, generated by subject matter experts, to assess the student-generated sub-goals. In our work reported here, we have seven ideal sub-goals associated with the general goal of learning about the human circulatory system. The sub-goals can be seen in the second level of nodes in Figure 1. A taxonomy of goals/sub-goals and concepts related to the sub-goals was chosen as the underlying scaffold for the sub-goal assessment and feedback mechanism (see Figure 1). This taxonomy captures general/specific relations among concepts and thus can help us drive the feedback mechanism. For instance, a student-generated sub-goal can be deemed too general if the sub-goal contains concepts above the ideal level in the taxonomy. Similarly, a sub-goal can be deemed too specific if it contains concepts below the ideal level in the taxonomy. We present the details of our taxonomy-driven sub-goal assessment and feedback model and report results on how well the system can assess student-articulated sub-goals.