

## Case-based Tutoring in Virtual Education Environments

Patrick M Regan and Brian M Slator  
Computer Science Department  
North Dakota State University  
Fargo, ND 58105

### Abstract

Virtual education environments are gaining popularity as tools to enhance student learning. These environments are often used to allow students to experience situations that would be difficult, costly, or impossible in the physical world. North Dakota State University provides students with environments to enhance their understanding of geology (Planet Oit), cellular biology (Virtual Cell), retailing (DollarBay), and history (Blackwood). In order to maximize the learning potential of each individual student, an ideal environment needs to provide customized lessons to that student based on his or her individual performance. One method to address this requirement is the use of case-based reasoning software. This software is used to monitor student performance, track progress throughout an environment, compare the student to other students in the same environment, and create customized tutor dialogs to communicate this information to the student in the form of individual tutor lessons. An example of case-based lesson building software that meets the above requirements can be observed in the current DollarBay retailing environment.

**Key Words:** Case-based Tutoring, Computers in Education, Intelligent Systems, Multimedia Applications

### Introduction

The traditional teaching environment is usually thought to be that of a classroom: a single teacher giving lectures to a group of students who are expected to use their notes and textbook to prepare for periodic examinations and demonstrate that they have learned. Technology provides an alternative to this scenario. One of the ways technology can be used to supplement learning is through the construction of virtual education environments to simulate scenarios that may be difficult for students to experience in the physical world. Using the Internet, students can access these worlds remotely, be it in a classroom or in the solitude of their own dwelling. At North Dakota State University (NDSU), the World Wide Web Instructional Committee (WWVIC) is engaged in research aimed at developing virtual education environments to assist in the education and growth of students. Some of the key factors that lead to success of these environments at NDSU are the use of graduate and undergraduate students in the development process, the use of the environments in actual classes, and the application of knowledge from one environment to the others. One of the major goals of WWVIC research is to find ways to provide tutoring agents to communicate "expert stories" to students as they progress through the environment. These agents should monitor the student and send advice on an "as needed" basis while being careful to never insist upon or block any course of action [Slator et al. 99].

Although several of the virtual education environments at NDSU provide some basic tutoring functionality to students, none had implemented a completely functional case-based system except DollarBay. This system provides complete analysis of student behavior based on selected attributes and a message delivery mechanism. This paper describes the design and functionality of the case-based tutoring system implemented in DollarBay. This system

provides the means to generate personalized lessons for each student participating in the DollarBay environment. In addition, it provides a framework that may be used to implement similar functionality in the other virtual education environments at NDSU with a minimum of coding.

The DollarBay simulation is based upon a client/server paradigm. The server side of DollarBay consists of a server program and a database. It is permanently connected to the Internet and allows other users to connect at any time, from any location, to the DollarBay environment. The database for DollarBay is implemented in LambdaMOO [Curtis 97]. The database contains representations of all of the objects in the DollarBay environment, including the MOO programs necessary to give the objects their specific behavior. On the client side, the student interacts in the DollarBay environment by using a graphical user interface (GUI). The GUI window is generated by a Java applet and serves as the connection to the LambdaMOO server. This allows the student to interact with the environment by pointing, clicking, and selecting objects as well as typing text.

### Playing the DollarBay Game

As players engage themselves in the DollarBay game, they are assigned a location and must decide what to sell, what level of service to offer, how much to spend on advertising, how much to stock, who to buy from, and what price to set to appear attractive to the customer agents [Slator & Farooque 98]. In order to simulate an economic environment, time is divided into "virtual weeks". Each week the customer agents are given a shopping list representing a weeks worth of demand for various products. After each week has concluded and the shopping lists are exhausted, each agent assigns new attractiveness ratings to each store based upon the past week's experience and new shopping lists are created for the upcoming week [Slator & Farooque 98].

At the end of each virtual week the weekly calculation charges players for their weekly expenses, recalculates the customer agent motivations as described above, and updates each of the player cases. At the end of a player's life, they are retired to the Hall of Fame. The Hall of Fame is a place where players are moved when they graduate from the game by reaching a profit goal or are inactive for a long period of time. Players are moved to the Hall of Fame by the reaper, who is sent out periodically to retire graduated and inactive players. The reaper is responsible for recycling all of the objects related to the player, such as store, company, ads, and products. Recycling makes object numbers available for reuse at a future date as new objects are created. The reaper also moves the player's active case to historical cases for future reference by case-based tutors.

### Overview of LambdaMOO

MOO is an abbreviation for a Multi User Dungeon Object Oriented or Multiuse Object Oriented system. It is designed to be a network accessible, programmable, interactive, multi-user system that is well suited to the construction of text-based collaborative software most commonly used as multi-participant low bandwidth virtual reality [Curtis 97]. It contains a small, simple language that is designed to be easy to learn and supports expressions, looping, control structures, and built-in functions. LambdaMOO models virtual reality by representing virtual world entities (tutors, cases, stores, players, etc) as objects. It supports other value types as well (integers, strings, etc) but the objects are the backbone of any LambdaMOO database [Zelenak 99]. A property is used to store an

arbitrary MOO value. Verbs are LambdaMOO programs associated with a particular object.

In addition to the functionality described above, LambdaMOO also includes a built in command parser. Any time the command parser is invoked, the first word is always taken to be a verb. After identifying the form of the command, the parser may then proceed to execute it using one of the built-in commands, or by looking up the MOO objects representing the direct and indirect objects and then executing the command [Curtis 97]. LambdaMOO is also an object-oriented language.

### Case-based Reasoning

“Case-based reasoning means reasoning based on previous cases or experiences” [Kolodner &Leake 96]. The purpose of this approach is to form a judgment about a situation based on cases that are already classified. In order to perform case-based tutoring, one must understand the definition of a case. A case is an example of a past problem and its solution. In DollarBay, cases are sets of both abstract and concrete attributes that represent the status of a player’s business strategy. Case-based tutoring is intended to enhance the effectiveness of student learning via educational simulation environments.

The design of an educational simulation environment is based upon several core principles, the first being that the environment is role based. This allows the student to fully assume the role of a character trying to accomplish some predetermined goal. The second principle is that the environment itself must be immersive. As the student assumes a particular role in the environment, he or she becomes immersed in it and starts to think and make decisions like that character would. The simulation must also be highly interactive]. Environments must have a goal or objective to accomplish. Finally, it must be game like [Slator 99].

### Virtual Educational Environments at NDSU

There are several environments currently implemented at NDSU. The Virtual Cell [White et al. 99] is a 3-D rendering of an interactive biological cell. Students are given assignments by a virtual lab instructor, and sent out to explore cells in tiny submarines. Planet Oit [Schwert et al. 99], is a part of the Geology Exploration Project and is used by geology students at NDSU. Students, playing the role of geologists on a field exploration of the mythical Planet Oit, are asked to acquire a set of field instruments. Blackwood [Slator et al. 01], is a simulation of a 19<sup>th</sup> century western town populated with intelligent agents who simulate the economic environment representative of the time. Players accept a role in the environment and are forced to compete with other players and agents holding the same role [Slator 99]. The DollarBay world allows students to simulate owning and operating their own retail store in the city of DollarBay. Students who quickly figure out how to best serve the needs of the shoppers rapidly rise to be profitable.

The overall goal of intelligent tutoring is to focus on developing and employing intelligent agents within multi-user distributed simulations to help provide effective learning experiences [Slator 99]. Examples of diagnostic tutoring may be seen in Planet Oit. For example, the science tutor looks at the decision making process that a player follows while trying to properly identify a material, and what experiments were performed on the material, the tutor is able to identify students who have made “lucky guesses” and let them know that they did not follow the proper process in getting to their answer.

Rule-based tutoring at NDSU was at one time functional in DollarBay. A rule-based tutor functions by knowing a set of rules about a domain, monitoring student action for any indication of breaking one of the rules, and then visiting the student to present a warning [Slator, Brian M. 99]. For example, one of the rules that the now defunct rule-based tutor in DollarBay monitored was if a student had set their prices to an excessive markup. In such an instance, the tutor would send a message to the student informing them that they may be setting their prices too high [Slator, Brian M. 99].

An example of case-based tutoring outside the NDSU realm is the Georgia Tech Case-Based Intelligent Tutoring System (GT-CBITS). This system is used to demonstrate the importance of critical information to airplane pilots by using stories of difficult situations or incidents encountered by other pilots. The topics presented can range from problems arising from the complex nature of aircraft, the dynamic nature of the aviation environment, or new/changing features of an aircraft [Palmer 02]. Case-based tutoring at NDSU has presently only been implemented in the DollarBay virtual world.

### Case-based Tutoring in DollarBay

The DollarBay case builder is centered on a set of verbs that are all defined on an object named \$g.gencase. These verbs control the creation, initialization, comparison, and updating of all cases. The children of the \$g.gencase object are the cases themselves. Case updates are carried out periodically during the “weekend calculations” in the game as well as prior to any tutor message generation. In addition to the \$g.gencase object, a \$g.tutor\_agent object exists for the purpose of generating tutor messages. This object owns the verbs that handle the construction of tutor messages and the scheduling of tutor visits. Tutor visits are scheduled during the execution of a player’s \$g.retail\_player:confunc verb, which runs each time the player connects to the game.

### Architecture of Case-based Lesson Building System

The backbone of any case based tutor system is an effective case matching procedure coupled with a substantial case library. The fundamental building block of any case-based system is a case. All cases in DollarBay divided into four distinct groups:

1. Active Cases (\$g.active\_cases) – the current cases for each active player’s store
2. Historical Cases (\$g.historical\_cases) – a copy of the final case for each store when the player is reaped and moved to the Hall of Fame.
3. Case Library (\$g.case\_library) – a copy of each unique case that has occurred in the history of DollarBay. Each of these cases represents a unique combination of the eight abstract properties of a case
4. Prototypical Cases (\$g.protocase) – the five prototypical cases that are used by the tutor to determine abstract case matches. The descriptions of these prototypes currently are:
  - a. Overcautious – player has plenty of money and is not using it to stock rapidly or expand product lines
  - b. Overspender – player is too liberal with money (expensive employee, expensive ads, loans) and has no clear price advantage or plan
  - c. Best Player – player who is profitable and not doing anything obviously wrong
  - d. Barking up the Wrong Tree – player products and services do not match ads and the prices are not appealing to target group

- e. Foolish Squanderer – player is bleeding money and has no clue about how to run a business

In DollarBay, each case shares a set of similar abstract attributes that define it. A list of the similarity measures, their acceptable values, and a brief description of how each is calculated follow:

1. product\_focus {"not evaluated", "tight", "moderate", "wide"} – looks at a list of the current product families and number of products stocked for each store. Computes a “focus factor” based on ratio of number of types to total quantity stocked
2. restock\_plan {"not evaluated", "understocking", "acceptable", "overstocking"} – estimates annual sales based on accounting records and annual revenue from the price of items currently on order. Creates a ratio of projected revenue to projected sales
3. ad\_cost {"not evaluated", "low", "affordable", "high"} – uses the liquid assets and advertising costs of a company to compute the average advertising cost per day. Determines the percent of available cash used to purchase advertising
4. ad\_target {"not evaluated", "fuzzy", "okay", "focused"} – gets a set of cluster groups based on what ads are running, looks at the sales records, and determines the effectiveness of each ad based on the attractiveness of the ad to each cluster group from the sales
5. staff\_level {"not evaluated", "underskilled", "appropriate", "overskilled"} – generates a ratio of current staff cost to total sales
6. cash\_reserves {"not evaluated", "low", "moderate", "high"} – gets the current cash on hand, the total of all expenses incurred, and calculates the average daily expenses
7. price\_margin {"not evaluated", "low", "moderate", "high"} – obtains the current cost and price for each product, computes the markup for each
8. liability {"not evaluated", "good", "risky", "very risky"} – gets the daily loan payments and generates a ratio of payments to income

In addition to the abstract similarity attributes of a case listed above, each case contains additional attributes that are used to assist in case-based lesson building. An explanation for each follows:

1. active\_similarity – contains a ranked list of active player cases sorted by the sum of their abstract similarity values. This attribute is used to determine the most similar active case to the current case
2. prototype\_similarity – contains a ranked list of prototypical cases sorted by the sum of their abstract similarity values. This attribute is used to determine the most similar prototypical case to a given case
3. product\_list – contains a list of products that are currently in stock for an active player store
4. prototype\_family – contains a pair of elements stored in a list. The first element is the object number of the prototypical case that a case matches the most. The second element is a descriptor of how closely the case matches the prototypical case contained in the first element
5. active\_tutor\_dialog – used during similarity calculations to store a list of information that is later used to generate a tutor message
6. last\_tutor\_visit – displays a list containing the turn and timestamp of the last tutor visit a player received
7. library\_match\_history – a list of lists. Each sub list contains a game turn and library case object number for the first time the current case matched the given library case. This attribute is

used to track which library cases an active case has matched over time

8. prototype\_tutor\_dialog - used during similarity calculations to store a list of information that is later used to generate a tutor message

#### **Prior Work**

At the outset of the implementation of this project, a portion of the work had already been completed. The current case-based message builder for DollarBay is based upon the original work.

#### **Weekly Case Update**

At the end of each virtual week in DollarBay, a set of weekly calculations is run. As a part of this process the weekly case update is performed, which brings the active cases for all of the valid player stores up to date by calculating the current values of the eight similarity measures described above. If an active case for each player store cannot be retrieved, a new case is created for that store. The case library is also updated with each active case.

#### **Update Case Library**

During a case library update the player’s case is first compared to every case in the library. During the comparison, the eight abstract similarity values of the player case are compared to the respective values of each library case. If an exact match is identified, the player’s case is added to the match\_history property of the library case. The property contains a list of the first time any player case has matched a specific library case. If the player’s case has already been matched to the current library case at some point in the past, no new entry is made to the match\_history of the case. If the player’s case has never been recorded in the match\_history of the matching library case, a new entry is generated using the generate\_history\_entry algorithm. The entry generated by this algorithm contains the current turn, the store owner object number, the store owner name, and the product families that are currently stocked in the store.

If, after comparing the player’s case to the entire case library, no exact abstract similarity value match is found, a new library case is created. The eight abstract similarity properties for the new library case are then set to the same values as the player case and the first entry, which contains the player’s information, is made to the match\_history list for the new library case.

The final step in updating the case library is to update the library match history for the player case. This property contains a list that keeps track of every library case the player case has ever matched. It is updated by first looking to see if the current entry will be the first. If so, a new entry, which consists of the turn number and the object number of the matching library case, is created. If the entry will not be the first, the last entry in the list is examined to see if its library case is the same as the library case that currently matches the player case. If the last entry does contain the same library case as the current match, no new entry is created. If the last match is not the same as the current library case match, a new entry is generated and added to the end of the list.

#### **Determining Case Similarity**

In order to construct a meaningful tutor message, cases must be compared. This comparison is completed in the hopes of determining similar cases from which information is drawn to construct a message. To accomplish this, the tutor messages depend upon comparisons of the player case to the prototypical cases and the other active cases. The compute\_similarity\_value verb is a general-purpose comparison engine used to determine the similarity of two

cases based upon their eight abstract similarity properties. The similarity is constructed by comparing each of the abstract properties for the two cases. A direct match results in 10 points, thus creating in a maximum possible similarity value of 80 for two identical cases.

The algorithm begins by looking at each abstract attribute on the two cases. If either case is “not evaluated” for the attribute currently being examined, the attribute is skipped and the next attribute examined. As the comparison is carried out, lists containing matched attributes and unmatched attributes are constructed as is the total similarity point value. The algorithm is able to handle comparison of prototypical cases with multiple values for an attribute by keeping track of the closest comparison between the other case’s value and the multiple attribute values of the prototypical case in such an instance. The verb then checks to see if the comparison involves prototypical cases. If it does, the `prototype_tutor_dialog` of the active case being compared is set to a list containing the prototypical case object number, a list of the matched attributes, and a list of the unmatched attributes. This data is important for future tutor dialog construction. The verb completes by returning the similarity value that is the result of the comparison.

### **Prototypical Case Similarity**

An active case’s prototypical similarity is stored as a list in its `prototype_similarity` property. In addition, each active case contains a `prototype_family` property that describes how strongly it matches the prototypical case that it is most similar to. In order to set the values of these properties, the `eval_prototype_similarity_and_sort` verb is called on an active case.

The first step in determining an active case’s prototypical similarity is to clear the `tutor_dialog` property of the case. This property is used to store information generated during the similarity check for later use in constructing a tutor message. The active case is then compared to each prototypical case by calling the `compute_similarity_value` verb (explained in the previous section).

As the active case is compared to each of the prototypical cases, the result of each comparison is stored in a similarity list as a list containing the object number of the prototypical case and the similarity value. Upon completion of the comparisons, this list is sorted by the similarity values from highest to lowest, and the resulting list is saved as the `prototype_similarity` property of the active case. The final step in determining prototypical similarity is to set the `prototype_family` of the active case.

The algorithm that determines the `prototype_family` of the active case begins by generating a list of the highest ranked prototypical cases. This list is called the family list and is generated in case there is more than one prototypical case with the highest similarity in the active case’s `prototype_similarity` list. In the event that there is a single entry in the family list, the `prototype_family` of the active case is set to the prototypical case of that entry. If there is more than one prototypical case in the family list, the algorithm compares the entries to determine which should be set as the `prototype_family` for the active case.

The final step in setting the `prototype_family` of an active case occurs by determining the strength of the match between the active case and the prototypical case. If the similarity value for the matched prototype family is high (in the top 25% of the similarity value range), the match is classified as “strong”. If the similarity value is between 50% and 75% of the maximum value, the match is classified as “weak”. Any `prototype_family` based upon a similarity value of less than 50% of the maximum, of which there should be very few, is classified as “very weak”. After the strength of the match is

determined, it is stored in the active case’s `prototype_family` property along with the object number of the most similar prototypical case.

### **Active Case Similarity**

An active case’s similarity to the other active cases is stored as a list in its `active_similarity` property. The value of this property is set by the `eval_active_similarity_and_sort` verb.

The first thing the `eval_active_similarity_and_sort` algorithm does is compares the player’s active case against every other active case by calling the `compute_similarity_value` verb. The results of each similarity computation are stored as a list of ordered pairs containing the object number of the compared active case and the similarity value. Upon completion of the comparisons, this list is sorted by the similarity values from highest to lowest, and the result is saved as the `active_similarity` property of the active case. Each entry in the list contains the object number of the active case that was compared, the similarity of the case to the player’s case, and two zeros that are used as placeholders for the concrete similarity results.

### **Concrete Active Case Similarity**

In addition to the active similarity value based on the eight abstract similarity properties it is useful to understand how a player compares to other players in terms that are more concrete. The intent of the `eval_concrete_active_similarity` verb is to examine active players of the highest abstract correlation in these concrete terms and provide additional comparison criteria to further refine the measure of similarity between them.

The algorithm begins by first determining the similarity value of the highest ranked active case in the player’s `active_similarity` list (the first case in the list, since it has been sorted). It then proceeds to compute a concrete similarity value for every entry in the player’s `active_similarity` list that has a similarity value equal to the first entry. The concrete similarity value is returned from a call to the `compute_concrete_similarity_value` verb. It should be noted that with ten or more players in the game there are often multiple active cases with the same active similarity in a player’s `active_similarity` list, and thus the need for this concrete comparison.

The `compute_concrete_similarity_value` verb determines a concrete similarity value by comparing several concrete attributes of the player’s case to the active case being compared. It keeps a running total of similarity points based on the results of each concrete comparison and returns it upon completion. The concrete similarity measures are:

1. `number_of_product_families_stocked` - examines the product list of each case and generates a count of product families for each by examining the ancestors of each product in the product lists.
2. `product_families` - examines the product list of each case and generates a list of product families for each
3. `number_of_products_stocked` - computes the length of the product list for each case
4. `average_markup` - looks at all of the `owned_objects` of a store’s owner, determines if each is a shipment, calculates the percent markup of the shipment, adds the markup to a running total, and increments a counter. When the verb is finished looking at all of the `owned_objects`, it then divides the total markup by the counter to get the average markup percentage for products stocked
5. `cash` - looks at the cash property of the company for each case
6. `player_age` - looks at the `first_connect_time` property for the owner of each player case

7. items\_sold\_last\_week - if an employee is found for both cases, the comparison begins by examining the lw\_sold property on each employee. A total of items sold is then computed for each case using this property
8. sales\_lost\_for\_not\_stocked - if both cases have valid employees, it examines the lw\_not\_stocked property of each. A running total of lost sales for each employee is then computed
9. employee\_type - looks at the parent of the employee for each case to determine the employee type
10. ad\_type - looks at the types of ads run by the owner of each case
11. net - looks at the net\_worth property of each case's company

There is a verb to handle each of the concrete comparisons and each returns a similarity score for that comparison based upon the strength of the match. The score returned by each verb is currently three points for a strong match, one point for a weak match, and no points for an insignificant match. Anytime a strong match is identified, a string describing the situation is appended to the active\_tutor\_dialog property of the player's case for later use in tutor message generation.

After the compute\_concrete\_similarity\_value verb returns a concrete similarity value, the eval\_concrete\_similarity verb finishes by updating the active\_similarity property of the case with new data. It was previously explained that the eval\_active\_similarity\_and\_sort verb placed two zeros as placeholders in each entry of a player's active\_similarity property of their case. The first of these placeholders is now filled with the concrete similarity score for the case, and the second with the sum of the active similarity and the new concrete active similarity scores.

### Tutoring

This tutoring process is currently triggered whenever a player logs in to DollarBay. During execution of the confunc (a function called at connect time), several criteria are checked to see if the player should receive a tutor. The player must meet all of the criteria to receive a tutor visit. The criteria are:

- Is the player a valid player?
- Is the tutor\_enable property of the player set?
- Is it true that the player has received no tutor visits yet this week?
- Has the player been playing for more than twelve hours?

If all of the criteria are met, the process begins by calling the schedule\_tutor\_visits verb. This verb coordinates the scheduling and sending of tutors to players that have just connected to DollarBay and limits the number of tutors that are sent at one time in order to avoid overloading the server. It begins by adding the connecting player to the list contained in the needs\_tutor property of the tutor agent. It then forks a process to handle the sending of the tutor to the player(s) in that list. The process looks to see if players are in the needs\_tutor list (at a minimum, the player that just connected should be in the list). It will continue processing all players in this list until it is empty. The forked process then looks to see how many tutor processes are currently running. If less than four tutors are currently active the process continues, otherwise it returns and lets the four active tutor processes handle the player who has just connected.

If there are less than three tutors running, the new process begins by taking a snapshot of the players in the needs\_tutor list. The process then finds the player in the list with the lowest net worth, operating on the theory that the player who is doing the worst needs

tutoring the most. That player is then removed from the list and the send\_tutor verb is called to dispatch a tutor to that player. Upon returning from the send\_tutor verb, the number\_of\_tutor\_visits\_this\_week property of the player is updated to document that the player has been visited by a tutor this week and the needs\_tutor list is examined to see if any players needing tutoring remain in the list. If so, the process repeats itself. Otherwise, the forked process terminates.

The first thing the send\_tutor verb does is increment the active\_tutors property of the tutor\_agent. It then proceeds to call update\_active\_case, eval\_active\_similarity\_and\_sort, and eval\_concrete\_active\_similarity (each described above) to completely update the player's case. The player is then examined to see if they are a perfect match to the best player prototypical case. If it is found that they are indeed in the "best player" class, the active\_tutors counter is decremented and the verb returns without doing anything.

If the player is not found to be a best player, a message is constructed for transmission to the player. The message begins with an introduction from the tutor. Following the introduction from the tutor, the build\_active\_message, build\_prototype\_message, and build\_historical\_message verbs are called to generate the personalized lessons for the player. These customized lessons are appended to the greeting and sent to the player as a tutor message via the show\_note verb. The show\_note verb pops up a new window containing the tutor message on the player's screen. The player may choose to read or ignore the message, and may keep the window open for reference throughout their playing session. The process completes by updating the last\_tutor\_visit property of the player's case with the turn and time that the tutor message was sent. It also decrements the active\_tutors property before it completes the process of sending the tutor to the player.

### Building Messages

The build\_active\_message verb creates a customized tutor lesson for the player based upon the store that is currently most similar to the player's store. It begins by looking at the active\_similarity list on the player case to find the store that is most similar to the player. A list of similar traits is then grabbed from the player's active\_tutor\_dialog, which was generated during the case update. The verb then constructs a message of the following nature:

*"If I were to pick a store that you are most similar to, it would be **XXX'S** which is located in **XXX'S\_NEIGHBORHOOD**. The things you have most in common with **XXX'S** are that you both:*  
- **SOME SPECIFIC TRAIT MATCH INFO**  
- **MORE SPECIFIC TRAIT MATCH INFO**  
- **ETC.**  
*Keep in mind that this may be good or bad, depending on how **XXX'S** is performing."*

The items shown in bold are replaced with information specific to the matched player store, and the lesson is returned to the send\_tutor verb where it is appended to the tutor message.

The build\_prototype\_message verb creates a custom lesson for the player based upon the prototypical case that they are most similar to. The lesson is pieced together to look something like this:

*"You **STRONGLY OR WEAKLY** remind me of a player who was **TYPE**. In general this means that:*  
- **GENERAL\_TYPE\_MESSAGE**  
*My advice for you is to:*

- **SPECIFIC\_MESSAGE**
- **ANOTHER\_SPECIFIC\_MESSAGE**
- **ETC.**”

The types and corresponding general messages are:

1. overcautious - you are not using your cash to expand or keep your stock level where it should be
2. overspending - you are far too liberal with your money
3. one of the best players - you are doing a great job
4. barking up the wrong tree - your products and service do not match your advertising, and your price is not appealing to your target group
5. a foolish squanderer - you are breaking the bank and are not using ANY sound business practices

Each of these messages is intended to give the player a general sense of their behavior pattern. The message for each player is based upon which prototype family they belong to.

The specific messages added to the end of the lesson are described in Table 1 below. The player will receive eight of these messages (one for each abstract similarity measure). Each message provides a detailed explanation on how a player should alter their behavior to better succeed or an affirmation of good behavior for each of the abstract similarity properties. The addition of these specific messages completes the construction of the prototypical tutor lesson, which is then returned to the send\_tutor verb.

	overskilled	- hire an employee with less skill and stop wasting your money on the one you have
cash_reserves	low	- try to keep more cash on hand
	moderate	- keep the same amount of cash on hand... good job!
	high	- try to spend more of your cash to improve your situation
price_margin	low	- keep your prices competitive... good job!
	moderate	- improve your prices as they are only marginally competitive
	high	- lower your prices to get competitive
liability	good	- keep your liability at an acceptable level... good job!
	risky	- consider reducing your liability
	very risky	- reduce your liability as soon as you can

Table 1 – Specific Prototypical Case Tutor Messages

### Historical Messages

The final lesson to be appended to the tutor message is based upon the history of the player. This lesson is constructed by going through the player’s library\_match\_history and examining all of the prototype case families a player has matched over time. The families are chronologically ordered, the duration of time spent at each is calculated, and a lesson of the following format is constructed:

*“Going back in time from present to distant past, it seems that you have been:*

- **TYPE STRONGLY\_OR\_WEAKLY** for **X\_NUMBER\_OF weeks**
- **TYPE STRONGLY\_OR\_WEAKLY** for **X\_NUMBER\_OF weeks**
- **ETC.**”

where the type correlates to one of those described above. If the player is currently “one of the best players”, a message stating, “*You are doing a great job! Keep it up!*” is appended to the end of the lesson. If the player was previously “one of the best players”, but has slipped into another prototypical family type, a message stating, “*You were doing a great job! What happened???*” is appended to the end of the lesson.

### Evaluation

Two "experiments" were carried out to study the effectiveness of the tutor. The first was a pilot study to demonstrate the functionality of the case library and populate it with cases. A group of senior level computer science students were assigned to play and comment on the game as part of their classroom activities during the period of 27-April-2001 to 4-May-2001. Instructions were sent to each student in the class via email. The students were divided into 5 groups. Each group member was assigned a player, given a specific behavior pattern to follow in the hopes of mimicking one of the five prototype cases, and asked to log on several times a day in order to keep their player focused on the appropriate behavior. In order to demonstrate the functionality of the case library construction software component of the case-based tutor, an analysis was performed on the cases in the case library.

Attribute	Value	Message
product_focus	tight	- stay focused on the types of products you are selling... good job!
	moderate	- don't change to any new types right now
	wide	- stick with selling the same types of products... don't change between different types
restock_plan	understocking	- stock more quantity of the products you sell
	acceptable	- keep stocking the same quantity products you sell... good job!
	overstocking	- stock less quantity of the products you sell
ad_cost	low	- spend more of your income on ads
	affordable	- keep spending the same amount of your income on ads... good job!
	high	- spend less of your income on ads
ad_target	fuzzy	- use market research to properly target your ads for the consumers you need to reach
	okay	- use market research to improve the targeting of your ads and reach the right consumer groups
	focused	- keep targeting your ads to reach the right consumer groups... good job!
staff_level	underskilled	- hire a more skilled employee to meet the needs of your customers
	appropriate	- keep your current employee... good job!

The results displayed strong evidence that the prototype similarity and family algorithms need to be tuned to provide a more accurate picture of player behavior. It is also likely that additional prototype cases are needed to fully represent the range of player behavior. Although the results of the analysis did not clearly show each player being classified as expected, they have provided valuable insight about weaknesses in player behavior coverage of the current case-based tutoring prototypical cases.

The second experiment conducted as a part of this research was aimed at capturing the effect of tutoring on players participating in a tournament in DollarBay. After a first tournament with no tutoring, tutors were activated evenly across the range of student performance in an attempt to avoid favoring any particular group and a second tournament was held. Any data that has been used has been completely stripped of any information that may identify the students by name in order to protect their identities.

During the second tournament, the tutor was activated for ten of the eighteen eligible players. These players received a customized tutor lesson once per turn as they played the second tournament.

While the small sample size and other social factors resulted in confounded results, certain measurable effects were seen. For example, the average number of case library matches for tutored students was 8.4 versus 8 for untutored students. The average number of prototype family changes for tutored students was 1.4 versus 0.5 for untutored students. Both measures show an increase in student behavior changes for the tutored students, and it is encouraging to see tutored students showing signs of behavior changes over untutored students.

### Conclusion

As virtual educational environments grow in popularity, students will have more opportunities to learn subjects at their own pace. However, one of the challenges of teaching via virtual educational environments is the customization of individualized lessons for each student. One approach to providing this critical feedback is the case-based tutoring implemented in DollarBay.

One of the current weaknesses of the system is that a better set of prototype cases is needed to more accurately model player behavior. With the case library that has been generated to date, future research may be able to use data mining to determine exactly what prototypical cases need to be implemented to improve this situation. Upon successfully incorporating new prototypical cases to the system, the tutor would then be able to provide more specific and accurate advice to players matching each prototypical case.

### References

- [Bouzeghoub 97] Bouzeghoub, Mokrane, George Gardarin, and Patrick Vaduriez, "Object Technology: Concepts and Methods", International Thompson Computer Press, Boston, MA, 1997.
- [Comer 99] Comer, Douglas E., "Computer Networks and Internets", Chapter 23, Prentice Hall, Upper Saddle River, NJ, 1999.
- [Curtis 97] Curtis, Pavel. "LambdaMOO Programmer's Manual for LambdaMOO Version 1.8.0p6", Xerox, San Francisco, CA, March 1997.
- [Headington 97] Headington, Mark R. and David D. Riley, "Data Abstraction and Structures Using C++", Chapter 10, Jones and Bartlett, Sudbury, MA, 1997.
- [Kolodner & Leake 96] Kolodner, J. L. and Leake, D. B., "Case-based Reasoning: Experiences, Lessons, and Future Directions", Chapter 2, AAAI Press/MIT Press, Menlo Park, CA.
- [Leake 96] Leake, David B., "Case-based Reasoning: Experiences, Lessons, and Future Directions", Chapter 1, AAAI Press/MIT Press, Menlo Park, CA, 1996.
- [Luger 98] Luger, G. F., and W. A. Stubblefield, "Artificial Intelligence – Structures and Strategies for Complex Problem Solving", Ch. 6, Addison Wesley Longman, Reading, MA, 1998.
- [Palmer 02] Palmer, E., NASA Ames, "Case-based Intelligent Tutoring for Pilots", <http://human-factors.arc.nasa.gov/projects/ih/casebased.html>
- [Schwert, Slator, Saini-Eidukat 99] Schwert, D.P., B.M. Slator, B. Saini-Eidukat, (1999). "A Virtual World for Earth Science Education in Secondary and Post-Secondary Environments: The Geology Explorer", Proc. Int. Conf. on Mathematics/Science Education & Technology, March 1-4, San Antonio, TX, pp. 519-525.
- [Slator, Brian M. 99] Slator, Brian M., "Intelligent Tutors in Virtual Worlds", 8th Annual Conference on Intelligent Systems (ICIS-99), Denver, CO, June 24-26, 1999.
- [Slator & Farooque 98] Slator, Brian M. and Golam Farooque, "The Agents in an Agent-based Economic Simulation Model", Proc. 11th Int. Conf. on Computer Applications in Industry and Engineering (CAINE-98), Las Vegas, NV, November 11-13, 1998.
- [Slator 99] Slator, Brian M. et al., "Research and Development of Virtual Worlds for Immersive Instruction", Proceedings of the Small College Computing Symposium (SCCS99), La Crosse, WI, April 15-17, 1999.
- [Slator et al. 99] Slator, B. M., P. Juell, P.E. McClean, B. Saini-Eidukat, D.P. Schwert, A.R. White, and C. Hill, "Virtual Environments for Education at NDSU", ED-MEDIA99, Seattle, WA, June 19-24, 1999.
- [Slator et al. 01] Slator, Brian M. with the members of CsCI345 (2001). "Rushing Headlong into the Past: the Blackwood Simulation", Proc. of the 5th IASTED Int. Conf. on Internet and Multimedia Systems and Applications (IMSA 2001), Honolulu, HI, August 13-16, pp. 318-323. Complete author list at <http://lions.cs.ndsu.nodak.edu/~mooadmin/papers/imsa-final.htm>
- [Stansifer 95] Stansifer, Ryan, "The Study of Programming Languages", Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [White, McClean, Slator 99]. White, Alan R., Phillip E. McClean, and Brian M. Slator (1999). The Virtual Cell: An Interactive, Virtual Environment for Cell Biology. ED-MEDIA 99, June 19-24, Seattle, WA, pp. 1444-1445.
- [Zelenak 99] Zelenak, Jozef. "Conversation Construction for Agents in Educational Simulation Environments", Master's Thesis, Computer Science Department, North Dakota State University, Fargo, ND, 1999.