Artificial Intelligence in Education
C.-K. Looi et al. (Eds.)
IOS Press, 2005

Motivating Appropriate Challenges in a Reciprocal Tutoring System

Ari BADER-NATAL [1], Jordan POLLACK
DEMO Lab, Brandeis University

Abstract. Formalizing a student model for an educational system requires an engineering effort that is highly domain-specific. This model-specificity limits the ability to scale a tutoring system across content domains. In this work we offer an alternative, in which the task of student modeling is not performed by the system designers. We achieve this by using a reciprocal tutoring system in which peer-tutors are implicitly tasked with student modeling. Students are motivated, using the *Teacher's Dilemma*, to use these models to provide appropriately-difficult challenges. We implement this as a basic literacy game in a spelling-bee format, in which players choose words for each other to spell across the internet. We find that students are responsive to the game's motivational structure, and we examine the affect on participants' spelling accuracy, challenge difficulty, and tutoring skill.

Keywords. Teacher's Dilemma, reciprocal tutoring system

1. Introduction

Reciprocal tutoring systems offer an interactive environment for learning [2,3]. Chan and Chou define reciprocal tutoring as "a protocol of learning activity, where two or three agents (an agent is a computer or a student) take turns to play the roles of a 'tutor' and a 'tutee'" [2]. One reason that these systems are of interest is that they can potentially avoid the complex engineering effort required to formalize domain-specific student models. This can be avoided by transferring the responsibility of model-building to the peer helper, using human-in-the-loop techniques, similar to Kumar, et al. [7]. In order to realize this, however, we must motivate peers to appropriately challenge one another. This is a problem, as there is often a *motivation gap* between an activity's educational objectives and its motivational meta-structure. Such gaps are now beginning to be identified. Magnussen and Misfeldt reported on the behavior that they observed when students began using their educational multi-player game, in which players learned how to excel at the game while avoiding the educational challenges involved [8]. Baker, et al. also identified intentional subversion of tutoring systems as an observed problem [1]. In this paper, we seek to recognize and attempt to close these motivation gaps.

We present the foundation upon which this alternative can be based – the *Teacher's Dilemma* (TD). With participants taking on the task of student modelling, the tu-

[1]Corresponding Author: Ari Bader-Natal. Brandeis University, Computer Science Department MS018. Waltham MA 02454. USA. Tel.: +1 781 736 3366; Fax: +1 781 736 2741; E-mail: ari@cs.brandeis.edu.

TD_T	*Easy Task*	(E)	*Hard Task*	(H)
$+$	verification	v	joy	j
$-$	remediation	r	confirmation	c

TD_S	*Task*	
$+$	pass	p
$-$	fail	f

Figure 1. The *Teacher's Dilemma*: TD Teacher-Matrix on left, TD Student-Matrix on right. $+$ and $-$ denote correct and incorrect Student responses.

T_{points}	*Task of difficulty* (d)
$+$	$dv + (1-d)j$
$-$	$dr + (1-d)c$

S_{points}	
$+$	p
$-$	f

Figure 2. The number of Teacher-Points awarded (for challenge selection) is determined by its difficulty and the accuracy of the response. The four variables v, j, r, c are defined in the game's TD matrix (Figure 1). The number of Student-Points awarded is dependent only on Student response correctness.

toring system must provide only tutor motivation and interaction facilitation. This has been implemented as a web-based system, *Spellbee*, that was designed from the ground-up to explore these ideas. It has been publicly available for over a year, at http://www.spellbee.org. In this paper, we first examine the validity of our assumption that a player's challenge-selection strategy is influenced by the underlying motivational structure of the TD, and then examine change in player behavior over time with respect to spelling accuracy, word difficulty, typing speed, and tutoring skill.

2. Foundation: The Teacher's Dilemma

The Teacher's Dilemma presented here originates from Pollack and Blair's formulation of the *Meta-Game of Learning* [9], and has more recently been pursued by Sklar and colleagues [4,10]. The intuition behind the TD is that providing students with excessively difficult or excessively easy challenges is counter-productive, while providing appropriately challenging tasks is more valuable. The four educational extremes defining the TD are *verification* of student success at easy tasks, *joy* of student success at difficult tasks, *remediation* of student failure at easy tasks, and *confirmation* of student failure at difficult tasks. The TD provides a simple framework for describing various combinations of these educational goals. Using the TD, a teaching strategy can be described by the values a teacher attributes to each of these goals. See Figure 1.

The application of the TD to reciprocal tutoring is done by transforming the TD's representation of teaching strategy from a *model* to a *game-theoretic* formulation. Strategies in this game correspond to selecting challenges of varying levels of difficulty. The payoff values for these strategies are based on the adopted valuations (from the TD Teacher-Matrix), the level of difficulty of the challenge selected, and the accuracy of the other player's response. Figure 2 details how these payoffs are calculated for players.

The novel value of this meta-game is that players who may have no tutoring experience are effectively learning to provide the same sorts of challenges as those provided by a "model" teacher (as exemplified by the TD matrix chosen.) Improving at the TD meta-game corresponds to more closely emulating this model teacher. Given an appropriate TD Teacher-Model, pairs of students could be organized to act as tutors for one another, providing each other with increasingly appropriate challenges. Using this model, we create an entire learning community based upon participants interacting in this manner.

3. Implementation: Spellbee

In order to further explore the ideas presented above, we have built a reciprocal tutoring network for the educational domain of spelling that is based on the Teacher's Dilemma. This system, *Spellbee*, was designed for use by students in grades 3-7, and takes the form of an online educational activity[1]. Spellbee.org has been actively used for a year, during which time over 4,500 people have participated, including approximately 100 teachers and over 1,300 students of those teachers[2]. In this section, we discuss the motivational structure of the game, the mechanics of game play, and metrics for assessing challenge difficulty in this section.

3.1. Motivational Structure

The underlying motivational structure of Spellbee is derived directly from the formulation of the TD, and is presented in Figures 1 and 2. In Spellbee, each player alternates between their roles as *problem-selector* (TD's Teacher-Role) and *problem-solver* (TD's Student-Role.) When attempting to spell a word, players receive points according to a Student-Matrix in which $p = 10$ and $f = 0$ (Correct spelling is rewarded, and incorrect spelling is not.) When selecting a word for a partner, players are presented with all word-choices and corresponding $+$ and $-$ row calculated from the TD's Teacher-Matrix, given the difficulty of the word. We set the parameters of the Teacher-Matrix to $v = 0$, $j = 10$, $r = 10$, $c = 0$, in order to reward students for probing both the strengths and the weaknesses of their partner's abilities. This matrix was designed to motivate players to seek out both the hardest words that their partner might be able to correctly spell and the easiest words that their partner might not yet know.

The game itself is competitive in the sense that the partner that accrues more points (sum of Student- and Teacher-Points) wins the game. A few publicly-displayed high-score lists are maintained on the website, providing players with additional motivation to take the game-points seriously. In Section 4, we will examine the degree to which players are aware of and sensitive to the underlying motivational structure.

3.2. Game-Play

A student accesses Spellbee online at http://www.spellbee.org, and uses their pseudonym to log in. Upon entering the system, a student is placed in a *playground* of currently-available players. Mutual interest between a pair of players triggers the beginning of a new game. A game consists of a sequence of seven rounds. In each round, a player first selects a word (from a list of seven options) for their partner to attempt to spell. Each word is accompanied by a pair of point-values, determined by Figure 2. Game-play is symmetric, so both partners are concurrently selecting words. After both players select words, the word-challenges are exchanged and each attempts to spell the word that the other provided. The word-challenges are presented in a multi-modal fashion: A

[1] Spelling was selected because we recognized that for classroom adoption, our game content must coincide significantly with curricular content.

[2] This counts unique registered players who played and completed games. Teachers are self-identified upon registration, and are only counted here if they subsequently register some number of students, and those students later play and complete games.

sentence that contains the word is displayed visually, with the word-challenge blanked-out, and that sentence is also presented audibly, by playing a pre-recorded audio clip of the sentence[3]. A player is given a limited amount of time to spell the word. After spelling the word, the student first gets feedback on the accuracy of their own attempt, and then gets feedback on the accuracy of their partner's attempt. This concludes the round, and the next round begins.

3.3. Word-Difficulty Metric

In order to apply the Teacher's Dilemma to reciprocal tutoring, some measure of a challenge's level of difficulty must be available. This metric might be defined *a priori*, might be estimated using some heuristic, or might be empirically-based. In the spelling domain, we initially started with a rough heuristic[4], but quickly switched to a metric based on a particularly well-suited empirical data-set. Greene's *The New Iowa Spelling Scale* [6] aggregates data from approximately 230,000 students across the United States in grades 2-8 attempting to spell words drawn from a list of over 5,000 frequently-used words. For each word, the study calculates the percentage of students of each grade-level that correctly spelled the word. Despite being dated, this data was ideal for our needs, and so we used these grade-specific percentages as our measure of word-challenge difficulty.

4. Experiment: On Motivation

An important assumption underlying claims of Spellbee's adaptability is that players are sensitive to changes in the TD Teacher-Matrix used in the game, and this matrix can influence a player's challenge-selection strategy. We examine the validity of this assumption in a set of classroom-based experiments.

The following was done using an early Spellbee prototype in a controlled classroom setting. Students were divided into four randomly-assigned groups, and were restricted to playing games with others in the same group. Each group played using a unique TD Teacher-Matrix, as specified in Figure 3. Students in group G_1 are rewarded most for asking easy questions, independent of their partner's success or failure at responding (*Reward Easy* game). Students in G_2 are rewarded most for either asking difficult questions that their partner can correctly answer *or* easy questions that their partner cannot answer correctly (*Teacher's Dilemma* game). Students in G_3 are rewarded most for asking difficult questions, independent of their partner's success or failure at responding (*Reward Difficult* game). Students in G_4 are rewarded most for asking easy questions that their partner cannot answer correctly, and are rewarded slightly less for asking difficult questions that their partner can correctly answer (anti-collusive *Teacher's Dilemma* game)[5].

In order to compare observed player strategies, each student was characterized by the relative difficulty (among the seven options) of the majority of the challenges that they selected during their second game[6]. If the majority of words selected were among

[3]The contextual sentences were drawn from an assortment of children's literature in the public domain. Initially, sentences were read aloud and recorded, but in an attempt to rapidly expand the game's problem domain, we began generating recordings using text-to-speech software.

[4]We initially used the *Scrabble*-score of a word as an approximation of difficulty.

[5]The skew in values is meant to prevent player collusion, which is theoretically possible within G_2.

[6]The first game was ignored in order to provide an opportunity to become familiarized with the game.

T_1	E	H
$+$	10	0
$-$	10	0

T_2	E	H
$+$	0	10
$-$	10	0

T_3	E	H
$+$	0	10
$-$	0	10

T_4	E	H
$+$	0	10
$-$	20	10

Figure 3. The Teacher-Matrix used in game-play had different parameter values for each of the four groups in the motivation experiment. The values for *v, j, r,* and *c* (from Figure 1) for the groups are listed here.

	Asks Hard	*Asks Medium*	*Asks Easy*	*Asks Mixed*	*Game Description*
G_1	25%	10%	45%	20%	Reward Easy
G_2	33%	29%	0%	38%	TD
G_3	70%	9%	7%	14%	Reward Difficult
G_4	46%	27%	0%	27%	Anti-collusive TD

Figure 4. Percentages of players within each group that behaved consistent with strategies at top. Each group plays using the correspondingly-numbered TD Teacher-Matrix from Figure 3.

the most difficult two options, the player's strategy was characterized as *Asks Hard*, if the majority were among the middle three options then the player's strategy was *Asks Medium*, and if the majority were among the least difficult two options then the player's strategy was *Asks Easy*. Players without any such majority were characterized as *Asks Mixed*. Figure 4 shows the resulting distributions of observed strategies, by group.

While the resulting variations were less pronounced than expected, they were noticeable. Those playing the *Reward Easy* game chose *Asks Easy* strategies more often than any other group and, similarly, those playing the *Reward Difficult* game chose *Asks Hard* strategies more often than any other group. Those playing the *Teacher's Dilemma* game chose *Asks Mixed* strategies more often than any other groups, which reflects our expected two-pronged strategy. Players in the anti-collusive *Teacher's Dilemma* game slightly less frequently chose *Asks Mixed* strategies, as would be expected from the one-sided bias of their matrix.

After reaching these results with the Spellbee prototype, we selected the G_4 game as the basis for the production version of Spellbee. The remainder of the paper assumes the use of this matrix. While players could theoretically collude to subvert this particular game variation, no such attempt has ever been made by any partner-pairs[7].

5. Observation: On Learning

Identifying and quantifying learning in a system of this sort is inherently difficult. What follows is an admittedly crude attempt to characterize changes in player behavior over time. We examine change with respect to *accuracy, difficulty, speed,* and *teaching-value*, and characterize it based upon the slope of a linear regression of a player's corresponding data, as a crude measure of direction and rate of change. If players are improving, we would expect such slopes to primarily be positive.

[7]Collusion would take the form of both players always selecting the easiest word available and then always responding to challenges incorrectly. In the past year, no player pair has done this for an entire game, or even for a majority of rounds of a game.

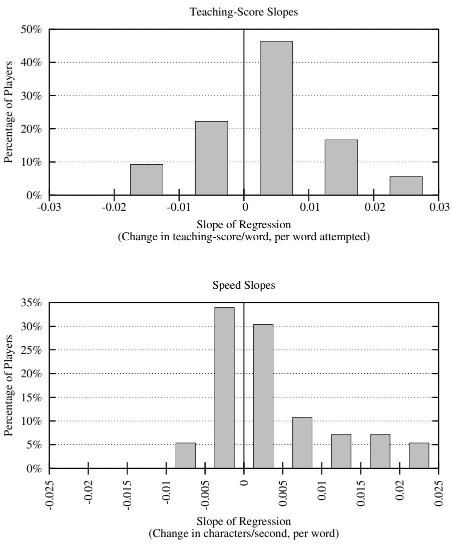

Figure 5. Graphs show players distributed according to the slope of linear regressions of data from the first 20 completed games of play. Anything to the right of the zero line indicates a positive trend.

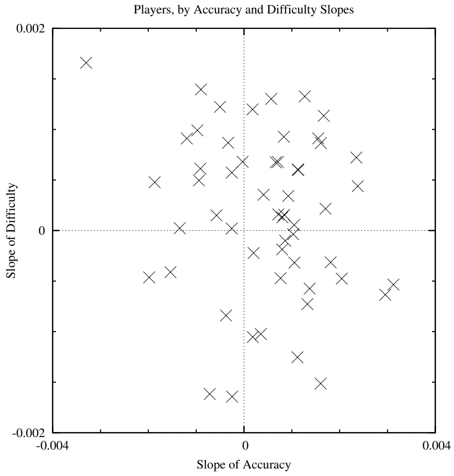

Figure 6. Players are plotted according to their difficulty slopes and accuracy slopes. Note that very few players occur in the lower-left quadrant.

For this set of experiments, we consider a refined subset of the data collected by the online Spellbee.org system[8]. Of these, we focus only on the first 20 games of players who have completed 20 or more games. Fifty-five players met all of these conditions. Given each player's sequence of 140 rounds of participation (20 games of 7 questions each), we calculate four data points at each round. In Figure 5, *speed* is measured in terms of average number of characters typed per second, and *teaching-score* is the Teacher-Points accrued in that round. In Figure 6, *difficulty* is determined by the New Iowa score for the player's grade-level, and *accuracy* is recorded as a binary correctness value.

In Figure 6, we graph each player in two dimensions: according to their difficulty and accuracy slopes. This allows us to differentiate among players in the four quadrants (players who do increasingly well or increasingly poor on problems of increasing or decreasing difficulty.) The graphs in Figures 5 and 6 indicate modest changes, but we wish to reiterate that the length of time used for these studies was relatively short. As more students participate for longer periods of time, our analysis can grow accordingly.

6. Discussion

One salient characteristic of open web-based educational systems like Spellbee is that participation is generally voluntary[9]. The non-trivially affects the dynamics of the system, in that the peer-tutoring network is only effective when it is able to retain student

[8] We consider only data recorded during a one-year period (February 1, 2004 through February 1, 2005), only considering players in grades 3-7 (inclusive), and only considering *completed* games (seven rounds finished.)

[9] The exception would be students in classrooms in which the teacher chose to have their class participate.

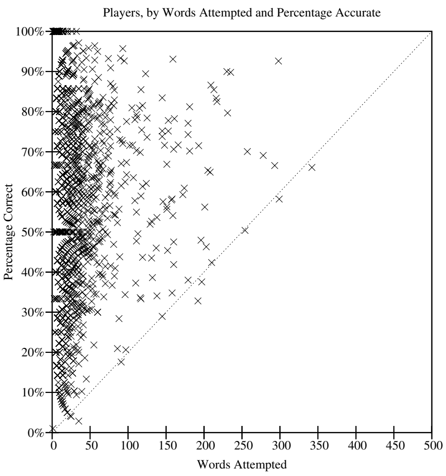

Figure 7. Players are plotted according to the number of words that they attempted while actively using the Spellbee system and the percentage of those words correctly spelled. A dotted line approximates the observed threshold at which students seemed to lose interest or motivation to continue participating.

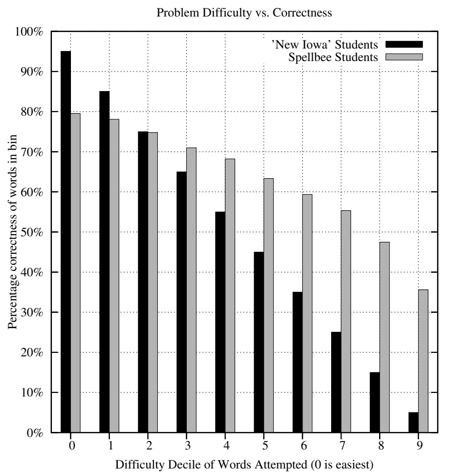

Figure 8. Words attempted by Spellbee players are classified by difficulty deciles according to the *New Iowa* scale. We then compare the percentage of these words spelled correctly by the students participating in Spellbee as compared to the students participating in the New Iowa study.

interest and participation over time. We seek to maintain this interest purely through the increasingly individualized and engaging educational interactions, rather than through extraneous means[10]. When we began exploring the return-rate data over the past year, we found that the rate of success that a student has at the game (used as an indicator of engagement) provides information about their likelihood of returning. In Figure 7, poorly engaged players (with extremely low rates of spelling accuracy) seem to have a consistent threshold for the maximum amount of repeat participation.

The spelling accuracy data that we have collected with Spellbee can yield the same type of statistics as provided by Greene's *New Iowa Spelling Scale* study [6]. In Figure 8, we compare expected student spelling-accuracy results according to the Iowa metric to the observed results from Spellbee participants. This suggests that we could theoretically stop using the Iowa data in our word-difficulty metric, and replace it with the empirical data that Spellbee have collected to date. While we have not yet taken this step, it suggests an interesting opportunity: when working with a domain for which no readily-available measure of difficulty exists, a rough heuristic can be used initially to bootstrap, and can later be replaced with a metric that uses the empirical data collected thus far.

While we have been leveraging the flexibility and openness of an internet-based system, we continue to encourage and support organized classroom participation. We recently found that one elementary school system in Michigan has over 900 students using Spellbee in school, and we hope to engage in more controlled studies with such groups in the future. This large-scale school-based participation seems particularly notable in

[10]Two frequently-requested additions to the Spellbee system are chat-functionality and a one-player version of the game. We have not implemented any extra-game communication channels due for reasons of child-safety, and we have avoided adding software players to the system in an effort to focus solely on the interpersonal nature of the peer-tutoring network.

light of work by Fishman, et al. suggesting that the adoption of research innovations by schools is often hindered by issues of system-paradigm scalability [5]. The active participation of this school district suggests that reciprocal tutoring networks like Spellbee may be as appropriate as an in-school activity as it has been as an extra-curricular activity.

The motivational layer that we have added to the reciprocal tutoring protocol enables a community of learners to learn to provide each other with the same sorts of appropriate challenges as a teacher may. As participants become more experienced at targeting the challenges that they provide, the tutoring system as a whole has improved as a learning environment. While this adaptive behavior is merely enabled and motivated by our system, this may be sufficient. Leveraging our human-in-the-loop design, we are able to envision tutoring systems that can be easily repurposed from one content domain to another.

7. Acknowledgements

The authors would like to thank members of the DEMO Lab for many useful discussions, in particular Shivakumar Viswanathan for feedback on earlier drafts of this paper. This work was made possible by funding from NSF REC #0113317, the Hewlett Foundation, and the Spencer Foundation.

References

[1] R.S. Baker, A.T. Corbett, and K.R. Koedinger. Detecting student misuse of intelligent tutoring systems. In *Proceedings of the 7th International Conference on Intelligent Tutoring Systems*, pages 43–76, 2004.

[2] Tak-Wai Chan and Chih-Yueh Chou. Exploring the design of computer supports for reciprocal tutoring. *International Journal of Artificial Intelligence in Education*, 8:1–29, 1997.

[3] Li-Jie Chang, Jie-Chi Yang, Tak-Wai Chan, and Fu-Yun Yu. Development and evaluation of multiple competitive activities in a synchronous quiz game system. *Innovations in Education and Teaching International*, 40(1):16–26, 2003.

[4] Mathew Davies and Elizabeth Sklar. Modeling human learning as a cooperative multi agent interaction. In *AAMAS Workshop on Humans and Multi-Agent Systems, at the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 2003.

[5] Barry Fishman, Ronald W. Marx, Phyllis Blumenfeld, Joseph Krajcik, and Elliot Soloway. Creating a framework for research on systemic technology innovations. *The Journal of the Learning Sciences*, 13(1):43–76, 2004.

[6] Harry A. Greene. *New Iowa Spelling Scale*. University of Iowa, Iowa City, 1954.

[7] V.S. Kumar, G.I. McCalla, and J.E. Greer. Helping the peer helper. In *Proceedings of the International Conference on AI in Education*, pages 325–332, 1999.

[8] Rikke Magnussen and Morten Misfeldt. Player transformation of educational multiplayer games. In Jonas Heide Smith and Miguel Sicart, editors, *Proceedings of the Other Players Conference*, Copenhagen, Denmark, 2004. IT University of Copenhagen.

[9] Jordan B. Pollack and Alan D. Blair. Co-evolution in the successful learning of backgammon strategy. *Machine Learning*, 32(3):225–240, 1998.

[10] Elizabeth Sklar, Mathew Davies, and Min San Tan Co. Simed: Simulating education as a multi agent system. In N. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Third International Conference of Autonomous Agents and Multi-Agent Systems*, 2004.