Intelligent feedback polarity and timing selection in the Shufti Intelligent Tutoring System

Stuart JOHNSON & Osmar ZAIANE
Alberta Innovates Center for Machine Learning, University of Alberta, Canada
{shjohnso, zaiane}@ualberta.ca

Abstract: It is well known that the training of medical students is a long and arduous process. Students master many areas of knowledge in a relatively short amount of time in order to become experts in their chosen field. The Socratic Method used in the latter stages of medical education, where a physician directly monitors a group of students, is inherently restrictive due to the limited number of cases and length of the students' rotations. Innovative Intelligent Tutoring techniques offer a solution to this problem. This paper outlines the overall structure and design of Shufti, an Intelligent Tutoring System (ITS) focused on mammography and medical imaging. Shufti's aim is to provide medical students with an improved learning environment, exposing them to a broad range of examples supported by customized feedback and hints driven by an adaptive Reinforcement Learning system and Clustering Techniques.

Keywords: Intelligent Tutoring System, Feedback, Hints, Reinforcement Learning, Machine Learning, Data Mining, Breast Cancer, Serious Games

Introduction

Shufti is an Intelligent Tutoring System (ITS) focused on mammography designed to help medical imaging students master the complexities of producing a diagnosis based on relatively poorly defined, low contrast images. Shufti takes the form of a web-based computer educational game where learners accumulate points for correctly diagnosing images. The learners are presented with pairs of mammographic images with overlaid grids and are expected to identify Regions Of Interest (ROI) within those images.

ROI's are regions which would normally necessitate further investigation by a radiologist. Students identify what they believe to be lesions by selecting squares within a grid which has been overlaid on the mammogram (see Figure 1). Once students have completed an exercise they are then given a score, which is derived from their accuracy in identifying lesions minus points for hints they may have requested along the way. Figure 1 depicts an example of feedback given to a student, post-exercise, showing their score relative to other students as well as relative to their previous attempts at the exercise.

Producing a high quality ITS for mammography is a non-trivial task. As demonstrated in Corwley et al.'s Slide Tutor [2], which made use of Natural Language Processing to resolve this issue in the field of pathology. Many attributes of the field do not lend themselves readily to computerized instruction. Amongst these, the two most important are the lack of sufficient time on the part of medical students to explore a large number of cases, and that mammography is an ill-defined domain according to the criteria outlined by Fournier-Viger et al. [3]

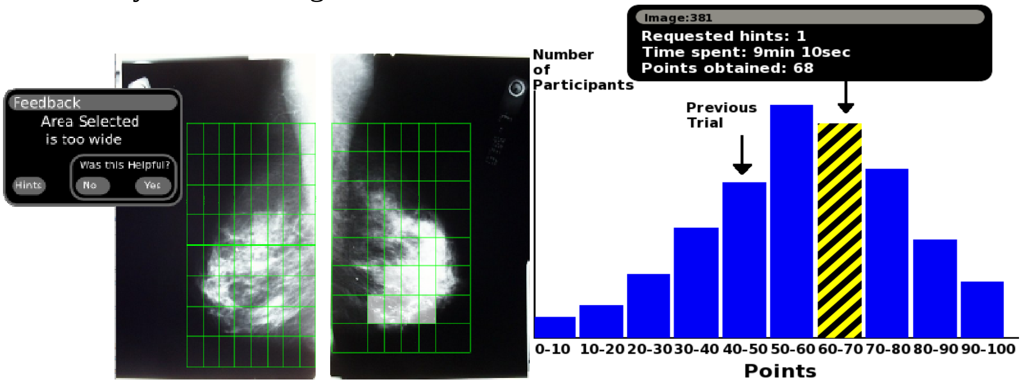

Figure 1: Shufti's feedback during and post-exercise.

To address the issue of learners viewing an insufficiently broad selection of cases, we not only make it easy for them to gain exposure to a wide array of well selected cases but we also incent them to broaden their exposure using a number of concepts from the field of competitive gaming.

The first game design convention used in Shufti is a point system. Points are a means of performance quantification that lack a direct connection with any one metric. In Shufti points are based on accuracy, time spent on an exercise, and the number and type of hints requested by the learner.

The second game design convention used is the concept of a competitive ranking system. Students are ranked based on their total accumulated scores on exercises, emphasizing volume and breadth of cases seen. Competition has been shown to be a good motivator for individuals to spend more time on tasks and can increase performance as demonstrated in Education and Crowdsourcing Schulze et al. [4].

Along with these game design conventions Shufti utilizes ideas taken from human tutoring; hints and feedback. In most literature about ITS, the terms “hint” and “feedback” are used interchangeably. This stems from the concept that all tutor-learner interactions are of the same assistive nature. In Shufti this is not the case; hints and feedback are complementary tools. Hints are much more specific than feedback in the information they impart to the user. An example of a hint would be, the number of ROIs present in the image, whereas an example of feedback, would be a message such as “Good job.”, or “Are you sure you’re done?” Additionally, Shufti provides hints in a prompted manner to the user whereas feedback is provided in an unprompted manner. More specifically, Shufti presents the user with a list of hints and associated score penalties during an exercise. Feedback, on the other hand, is provided automatically with the user having only an indirect influence on whether or not it is issued.

1. Feedback and Hints

One of the core attributes of human one-on-one tutoring is the active role which the tutor plays in the learning experience. An effective human tutor will provide hints, positive feedback and constructive/negative feedback in a strategic fashion so as to aid the learner's progress. Moreover, while some people may appreciate feedback, others may dislike it. Human tutors intuitively understand to whom feedback is beneficial. Mammography lacks the clear domain models, formal theorems, or cognitive models necessary to automatically teach mammogram diagnosis[3], consequently Shufti utilizes a variety of means to effectively simulate attributes of a human tutor.

Users are presented with a set of possible hints, each one being labeled with a description of the type of information the user will receive, and the specific score penalty which will be applied should the user accept the hint.

Shufti assigns penalties to hints to discourage gaming of the system, a phenomena where the user repeatedly requests hints until the answer is fully revealed [1]. Hint penalties may also have the interesting effect that learners will strategically select the minimum number of hints necessary for them to answer an exercise correctly. This strategic hint selection results in a form of user controlled difficulty, as, if the student selects the minimum number of hints necessary for the completion of the exercise, they have in effect adjusted the difficulty to the maximum they can successfully complete. Under normal circumstances mammography lacks an automated means with which to determine exercise difficulty. By giving users control over their own difficulty but incenting users to tackle challenges, Shufti elegantly resolves this issue.

Feedback in Shufti takes the form of both negative and positive messages – an attribute called polarity. Positive feedback is encouraging in nature such as, “Nicely done” and serves as a sign to the learner they are on the correct path to solving the exercise. Negative feedback is corrective in nature serving to steer the learner back onto the correct path with statements such as “You've missed something,” or “Look around more.”

The timing of a feedback is critical as it needs to be associated with an event or state and should not be disruptive. The selection of feedback polarity and timing is performed using two methods. First a clustering-based method is used which relies on a learner's reaction to feedback in order to determine its relevance. The timing in this method is controlled by one of the timing models described below and displayed in Figure 2. The second method relies on Reinforcement Learning (RL) to control the content, polarity, and timing of feedback delivery.

2. Approaches for Feedback

Exercises in Shufti are categorized by difficulty level, and students move from one level to the next after accumulating sufficient points and diagnosing a defined number of mammograms. Learners are modeled by retaining their current level, the total number of points, the number of images they have attempted, the average number of hints they requested per image, and the accumulated penalties due to requested hints .

Moreover, Shufti records the task state transitions during each exercise. Task State Transition are comprised of the exercise state and actions undertaken by the learners during the exercise. The Task State Transition record includes the current and past states representing the current solution, the last action taken by the learner, the proposed feedback, and the reaction to the feedback by the learner. The state is the number of grid cells selected that differ from the exercise solution (i.e. hamming distance). Actions are operations such as selecting a square, de-selecting a square, some mouse movements, or submitting the exercise for evaluation. The reaction to feedback is whether the learner explicitly found the previous feedback helpful.

Feedback polarity is based upon whether the state of the exercise improved or degraded. The degradation or improvement is determined by comparing whether the hamming distance between the past state and the solution has increased or decreased in contrast with the hamming distance between the current state and the solution.

Feedback is a critical part of the effectiveness of a tutor. Shufti contains methods for determining the content, polarity and timing of feedback. Polarity refers to whether it is a positive, encouraging message or a negative, corrective message. We propose two feedback control approaches: a clustering-based method and a technique based on RL.

2.1 Clustering-based method

The first approach Shufti uses to decide whether or not to give a feedback and the type of feedback to use, considers learners in groups of similar learners – in other words, *clustering*. Clusters of similar learners are based on their levels, points accumulated, their requested hints and the number of attempted exercises.

Feedback timing is determined by one of five models (see Figure 2): *Random* is feedback occurring randomly. *Timed* is feedback delivered after timed intervals. *After Action* feedback is issued in response to a learner's action. *Timed After Action* is triggered by an action however is delayed. *Random After Action* is again similar to After Action except it is randomly delivered (i.e. may or may not be issued).

Delivery of feedback for the timing models is based on the task state transitions of similar learners (i.e. learners in the same cluster as the current learner). The nature of the feedback delivered is chosen by examining the feedback that historically has been most appreciated by other users in the same cluster and the same task state transition.

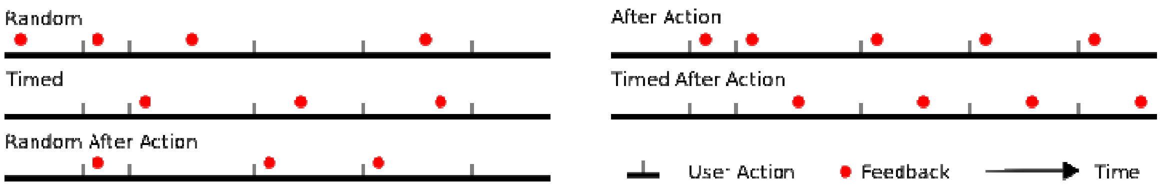

Figure 2: Patterns of feedback timing models.

2.2 Reinforcement Learning based method

Tuning to individual students is one of the ways in which a tutor can offer a superior learning experience. RL offers an automated method with which an ITS can tune its feedback delivery to individual learners and thus approximate a human tutor.

RL is a class of machine learning techniques which resolve problems of mapping situations to actions to maximize or minimize a metric[5]. RL allows Shufti to learn the most effective times to issue feedback, avoiding the use of preset timing models.

An RL system can be thought of as two components; an environment and an agent within which it acts. The environment provides state data and a reward signal to the agent which attempts to maximize the total reward over time. The agent makes use of methods such as Temporal-Difference Learning[5] or Monte Carlo Methods[5] to determine the most long-term rewarding action to take in any given state.

Shufti's environment offers task state transitions as state information to the agent seeking to minimize P in the following formula: $P = \sigma * count(\tau) + \omega * count(f) - \alpha * score$ where P is the total penalty assessed to the agent, σ is the penalty assigned over time, $count(\tau)$ is the total time passed, ω is the feedback penalty, and $count(f)$ is the total number of feedbacks given, α is the reward per score point earned, and *score* is the total score assigned for the exercise. The longer the learner takes the larger the penalty which encourages the agent to provide feedback. This is balanced by penalties assigned to the agent for giving feedback, which results in strategic feedback issuance. Rate of feedback increases with σ and decreases as ω increases. RL allows Shufti to fine tune for individual learners, thus more closely matching a human tutor, however it lacks the adaptive advantages of learning from many users (clustering).

3. Competition

One of the limitations in traditional imaging analysis training is the amount of cases students are exposed to. Shufti improves on traditional training in two ways; it has an extensive selection of exercises covering a range of scenarios unlikely to be seen during a short rotation as a student in a radiology department, and, it uses competitive techniques learned from gaming to incent students to review as broad a range of scenarios as possible.

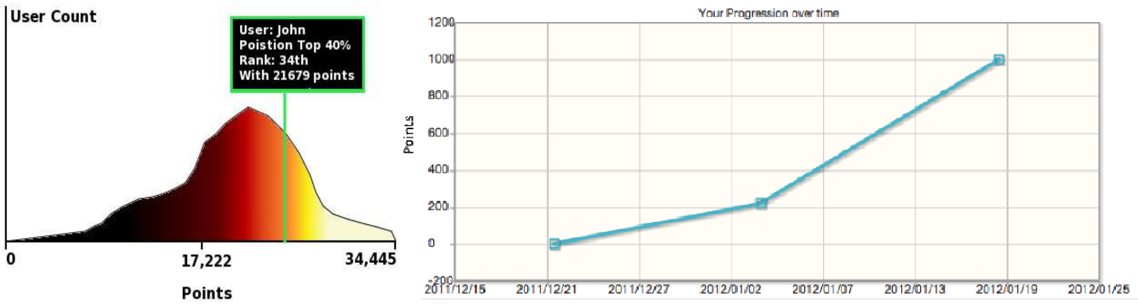

Figure 3: A representation of a user's position

In total, scores are based on problem difficulty, answer accuracy, time spent, and the hints requested. Learners are presented with a variety of means to see how they rank next to their peers, including public leader boards (similar to those used in online games), and performance distribution curves. Figure 3 shows the distribution of participants' scores and the relative position of the user "John". Overall ranking in Shufti is determined by the sum of all scores they have received, encouraging them to attempt a large number of exercises.

4. Conclusion

Shufti is an innovative solution to many of the issues involved in providing a high quality learning experience to learners in the field of mammography. We present two machine-learning techniques to provide high quality automatic feedback to learners; Clustering and Reinforcement Learning. We also make use of an interesting user controlled hint structure in an effort to not just reduce learners' attempts to game the system but to also exploit the gaming habits of learners in an effort to aid their learning experience. Along with feedback and hints, Shufti also innovates in the field of ITS UI design adopting many features from serious games in an effort to improve the learners' experience.

References

[1] Baker, R. S., Corbett, A. T., Koedinger, K. R., Lester, J. C., Vicari, R. M., & Paraguacu, F. (2004). Detecting Student Misuse of Intelligent Tutoring Systems. In *Intelligent Tutoring Systems* (Vol. 3220, pp. 54-76). Berlin/Heidelberg: Springer Press.

[2] Crowley, R. S., Legowski, E., Medvedeva, O., Tseylin, E. E., & Roh, J. D. (2007). Evaluation of an intelligent tutoring system in pathology: Effects of external representation on performance gains, metacognition, and acceptance. *JAMIA* , *74* (2), 281-296.

[3] Fournier-Viger, P., R., N., & Nguifo, E. M. (2010). Building intelligent tutoring systems for ill-defined domains. *Studies in Computational Intelligence* , 308.

[4] N., K., Schulze, T., & Veit, D. (2011). More than fun and money. Worker motivation in crowdsourcing - a study on mechanical turk. *Seventeenth Americas Conference on Information Systems.*

[5] Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction.* Cambrige, MA: MIT Press.