

Filtered Time Series Analyses of Student Problem-Solving Behaviors in Game-based Learning

Robert Sawyer

Department of Computer Science
North Carolina State University
rssawyer@ncsu.edu

Jonathan Rowe

Department of Computer Science
North Carolina State University
jprowe@ncsu.edu

Roger Azevedo

Department of Psychology
North Carolina State University
razeved@ncsu.edu

James Lester

Department of Computer Science
North Carolina State University
lester@ncsu.edu

ABSTRACT

Student interactions with game-based learning environments produce a wide range of in-game problem-solving sequences. These sequences can be viewed as trajectories through a game's problem-solving space. In this paper, we present a general framework for analyzing students' problem-solving behavior in game-based learning environments by filtering their gameplay action sequences into time series representing trajectories through the game's problem-solving space. This framework was investigated with data from a laboratory study conducted with 68 college students tasked with solving the problem scenario in a game-based learning environment for microbiology education, CRYSTAL ISLAND. Using this representation of student problem solving, we derive the slope of the problem-solving trajectories and lock-step Euclidean distance to an expert problem-solving trajectory. Analyses indicate that the trajectory slope and temporal distance to an expert path are both correlated with students' normalized learning gains, as well as a complementary measure of in-game problem-solving performance. The results suggest that the filtered time series framework for analyzing student problem-solving behavior shows significant promise for assessing the temporal nature of student problem solving during game-based learning.

Keywords

Game-based learning, Problem solving, Time series, Dynamic analysis

1. INTRODUCTION

Game-based learning has shown considerable promise for motivating and engaging students in learning [8]. Game-based learning environments engage students by populating game worlds with believable characters and narrative-driven learning experiences. These environments often feature problem-solving scenarios that give students a high degree of agency and freedom.

While engaging for students, this freedom also allows different problem-solving strategies to be pursued to varying degrees of effectiveness. Providing adaptive scaffolding to guide students in following effective problem-solving processes is key to creating effective game-based learning experiences. However, determining how to best scaffold student problem solving in game-based learning environments remains an open research question. Scaffolding effectively requires insight into students' problem-solving processes as well as their individual student characteristics.

In order to devise effective models for adaptive scaffolding in game-based learning environments, it is important to consider how the scaffolds will influence students. The models not only need to account for what support to provide, but also when to provide that support. In other words, the dynamic nature of student problem solving within game-based learning environments should be considered when analyzing the problem-solving behaviors of students. Thus, considering the overall sequence of a student's actions in a game-based learning environment is fundamental to making effective scaffolding decisions, including what a student has done thus far, what their general approach has been, and what cognitive and metacognitive strategies they have been using.

The space of possible problem-solving behaviors within a game-based learning environment can be vast, as students explore, inquire, gather information, and attempt to leverage their knowledge and skills to solve the problem scenario over an extended interaction. In these open environments, providing an exemplar solution path that is known to be effective can serve as a useful reference for students. Domain experts solve complex problems more efficiently than novices [12], and their solutions can serve as valuable points of comparison by students who lack relevant problem-solving expertise. The similarity between an expert solution path and a student solution path can be used to draw inferences regarding the student's trajectory through the open problem-solving space of the game-based learning environment.

In this paper, we present a general framework for analyzing the temporal sequence of student problem-solving behaviors in comparison to expert solution paths in game-based learning environments. The framework consists of filtering student problem-solving actions in a game-based learning environment into a time series representing a student's trajectory through the problem-solving space. We investigate the framework with data collected from student interactions with CRYSTAL ISLAND, a game-based learning environment for microbiology education. To

evaluate the framework, we compare several key characteristics of the time series, including a comparison between student trajectories and an expert trajectory, with measures of learning and engagement in game-based learning.

2. RELATED WORK

A growing research base focuses on analyzing problem-solving behaviors of students using summary statistics of student interactions with learning environments. Toth and colleagues clustered summary statistics of students' interactions with a computer-based educational assessment to discriminate between students with different proficiency levels in problem solving [32]. Sawyer et al. used rates of emotions and action units during student interactions with a game-based learning environment to model learning and engagement outcomes [28], while Lalle et al. used eye-gaze measures during student trials with ValueChart, an interactive visualization for preferential choice, to predict student confusion [18]. While successful in using student data to model outcomes important for adaptive learning technologies, these methods did not leverage the sequential structure inherent in student problem solving in advanced learning technologies.

Modelling sequences of student actions has important implications for adaptive learning environments, and it has been approached using both supervised and unsupervised learning methods. Kock et al. modeled sequences of user activities in an e-learning tutor as discrete Markov models, detecting problem-solving styles and learning dimensions about learners by clustering on the trained parameters of the models [17]. They subsequently investigated how these data-driven insights about students can be incorporated into an adaptive learning environment by supporting both individual users and groups of collaborative users. Hidden Markov models (HMMs) have been used widely for sequential student behavior modelling [6, 14]. Beal et al. used HMMs to model the actions of high school students [4]. After fitting HMM parameters for each student, they performed clustering based on the transition matrices of individual students to gain insight into differences in behavior and achievement of the clusters. Hansen and colleagues modeled student session log data by modeling student behaviors as distributions of Markov chains [13]. Bayesian knowledge tracing models use sequences of observations of student performance to create hidden Markov models with binary latent states representing student knowledge [9, 15]. All of this work shares the common approach of modeling student action sequences in terms of probabilistic state transitions. In contrast, our work uses characteristics of student problem-solving sequences encoded as trajectories within the game-based learning environment to predict student learning outcomes measured through pre and post-testing.

Sequence mining techniques have been used to investigate student activity sequences in adaptive learning environments to identify frequent behavior patterns and their evolution over time [16]. Martinez et al. used sequence mining on logs of a collaborative tabletop problem-solving application to examine frequently occurring problem-solving strategies in high and low achieving groups [21]. Perera et al. used trace logs of a collaborative software engineering environment to extract frequent patterns and cluster students using k-means clustering [22]. Another widely used approach is applying pattern mining techniques to logs of user behaviors in web-based learning environments [11, 23]. Our work differs from these approaches by analyzing the paths of student behaviors over full gameplay episodes rather than specific subsequences of behaviors. This approach is taken because a full trajectory and segments of the trajectory provide a comprehensive

view of a student's problem-solving process, which is composed of a very long sequence of problem-solving behaviors taken to solve the open-ended game-based learning environment.

Bauer et al. devised solution tree visualizations of user interactions with an open-ended puzzle solving game about protein folding, *Foldit* [3]. They used the visualizations to identify key patterns in problem-solving behavior among high and low performers. Others have used visual data mining on player behavior states, projecting visual representations into a more interpretable visual space [2, 19]. Notably, Liu et al. used state features to collapse complex visualizations and interpret key moments of player behaviors [19]. Our work similarly uses dimensionality reduction to create more interpretable visualizations of player behaviors over time. The primary focus of our work is quantifying the problem-solving trajectories of students in game-based learning environments, and the filtering approach we apply to student action sequences supports creating useful visualizations of the students' solution paths through the problem-solving space. While the calculated slopes and distances are quantities, their geometric interpretation with regard to the problem-solving space are also informative.

Snow et al. used a random walk analysis based on student interactions within a game-based system, iSTART-ME, to plot student trajectories and slopes [30]. They later extended this work through comparisons of student behavior patterns against random walks, revealing that students who behaved in a more deterministic manner exhibited higher quality self-explanations [31]. Our work similarly aims to dynamically analyze student trajectories based on interactions within a game-based learning environment, but it differs in several key aspects. First, our work creates student trajectories of problem-solving behaviors within an open-world game-based learning environment, a more complex space, which requires filtering through dimensionality reduction. Second, our work compares student trajectories to an expert solution path as opposed to a random walk. This comparison is particularly useful for informing the design of adaptive scaffolding functionalities in game-based learning environments. Experts and novices solve problems differently [12, 20], and our work provides an automated framework for characterizing how expert and novice problem-solving paths differ from one another.

3. GAME-BASED LEARNING TESTBED

In this work, CRYSTAL ISLAND, a game-based learning environment for microbiology education, was used as a testbed to explore the problem-solving behavior paths of students and an expert. Students who participated in the study played CRYSTAL ISLAND and completed a pre-test and post-test assessing microbiology content knowledge.

3.1 Crystal Island

CRYSTAL ISLAND integrates science problem solving in a game-based learning environment designed for microbiology education. Students adopt the role of a medical field agent tasked with discovering the source and identity of a mysterious epidemic on a remote island. In order to diagnose the illness, students gather information through conversing with a cast of non-player characters. Reading scientific books, articles, and posters scattered throughout the island provides crucial sources of information about microbiology that students need to diagnose the illness. Students test their hypotheses for the epidemic's source by scanning objects for contamination in the virtual laboratory. Students record findings regarding symptoms and contaminated objects on a diagnosis worksheet. The mystery is solved by submitting a completed diagnosis worksheet with the correct illness, source, and treatment




Figure 1. Overview of CRYSTAL ISLAND with expert solution path in gold.

plan to the camp nurse. Throughout solving the mystery, students explore an expansive 3D virtual game environment that includes a beach, infirmary, laboratory, dining hall, and residences.

There are many possible solution paths to solving the mystery successfully. An expert created an expert playthrough for a solution representing a thorough but efficient solution path for the problem-solving scenario. In a related study, a recording of this expert playthrough was used as a *No Agency* condition [7, 29], where students watched the narrated video of the expert solving the CRYSTAL ISLAND problem scenario. The expert visited each building, interacting with each of the virtual characters and reading each of the scientific texts to learn the information needed to solve the mystery (Figure 1). Although it is possible for a student to solve the mystery more quickly by skipping content in the game, the expert playthrough is intended to represent a comprehensive, efficient problem-solving path that any student could implement regardless of prior knowledge. In this work, we analyze students from the *Full Agency* condition of the study, which allowed students to freely explore the game environment after a brief tutorial introducing basic game mechanics. The expert playthrough is used for a comparison of problem-solving behaviors over the course of the gameplay interaction.

The CRYSTAL ISLAND problem scenario consisted of three phases of gameplay: (1) Tutorial, (2) Information Gathering, and (3) Diagnosis. In the Tutorial phase, students learned the basic game controls and mechanics upon arriving on the island's beach. After completing the tutorial, students moved to the main area of the game, beginning the Information Gathering phase. Students gather information through books, posters, and conversing with non-player characters such as the camp nurse, who initiates the game's problem-solving scenario narrative. Students also converse with a range of domain experts and sick patients in the game. Students transition into the Diagnosis phase when they perform their first test with the virtual laboratory scanning equipment. The Diagnosis phase and overall game are solved when students successfully submit their diagnosis worksheet to the camp nurse with the correct illness, contamination source, and treatment plan.

Outside of the Tutorial, the phases do not restrict any aspect of a student's experience within the game-based learning environment. The phases are used to segment a student's gameplay for an analysis of problem-solving behavior in different intervals of the scenario.

3.2 Study Participants

The study involved 68 participants from a large mid-Atlantic university who played CRYSTAL ISLAND in a lab setting. After removing students with corrupted data there was a total of 63 students ($M = 20.1$ years old, $SD = 1.55$) of which 42 (66.7%) were female. Prior to interaction with Crystal Island, students completed a 21-question multiple choice pre-test assessing microbiology knowledge ($M = 11.5$ (54.8%), $SD = 2.7$ (13.0%)). Students played for a range of 26.4 to 159.8 minutes ($M = 68.0$ min, $SD = 22.4$ min) while the expert playthrough lasted 91 minutes. After completion of the game, students completed the same microbiology assessment as a post-test ($M = 13.3$ (63.5%), $SD = 2.7$ (13.0%)).

3.3 Measures of Learning Performance

A primary goal of CRYSTAL ISLAND is learning relevant microbiology content. We measure student learning in CRYSTAL ISLAND in terms of normalized learning gain, which is the difference between pre and post-test score standardized by the total amount of improvement or decline possible from the pre-test. This calculation uses percentage of questions correct on the pre-test and post-test to calculate learning gain. Students demonstrated positive normalized learning gains with an average normalized learning gain of 0.19 ($SD = 0.26$).

A previously used indicator for in-game student engagement assessing progress and efficiency in the problem-solving scenario is given by *final game score* [25]. This measure was designed to allot points to students for efficient problem-solving behaviors such as talking to key virtual characters and solving the mystery in a short duration while subtracting points for inefficient behaviors such as scanning incorrect items in the virtual laboratory or submitting an incorrect solution. *Final game score* has been shown to be significantly associated with post-test score, independent of

pre-test score [25]. Scores varied widely among students with a range of -1543 to 1502 and an average of 673.7 ($SD = 616$). Both learning, as measured by normalized learning gain, and in-game student engagement, as measured by *final game score*, are target learning objectives of game-based learning environments. We therefore investigate how learning and in-game student engagement are related to student problem-solving trajectories in order to evaluate the utility of the filtered time series analysis framework.

4. TIME SERIES ANALYSIS

The similarity of two students over their entire gameplay can be defined as the distance between their trajectories through the game. First, we define student trajectories as filtered cumulative actions over time. Then, we define the temporal distance as the average Euclidean distance between trajectories over each time step, which is known as the *lock-step Euclidean distance* [10]. Distances between students and the expert playthrough are calculated. The slope of the trajectory is calculated as the ordinary least squares regression line through data points of each student’s time series, roughly measuring the problem-solving behavior of a student through an adjusted gameplay pace. This distance representing student gameplay similarity to the expert path and regression slope are compared to established measures of learning performance in CRYSTAL ISLAND: normalized learning gain (NLG) and final game score [25].

4.1 Filtering Process

Students perform several different problem-solving behaviors while interacting with CRYSTAL ISLAND. The cumulative counts of student in-game actions are recorded during gameplay, including conversing with virtual characters, reading books and articles, editing the diagnosis worksheet, completing a plot point, submitting a worksheet, and scanning an item in the virtual laboratory. A dimensionality reduction technique to convert the six cumulative counts of actions into a single value describing student progress until a particular moment in time reduces noise in distance measurements by lowering the dimensions used in calculating Euclidean distance. Filtering a multivariate time series to a univariate time series is used in sequential distance methods to reduce the effect of noise on the distance [5].

Due to the correlations between cumulative action counts at specific time intervals, principal component analysis is used for




Figure 2. Filtering process from action sequence to time series.

dimensionality reduction [1]. Specifically, the first principal component is used to filter a vector of cumulative action counts at a point in time to a single value (Figure 2). The principal components are calculated on the final action counts of each student (not including the expert counts), and the first principal component (variance explained = 37%) projects the cumulative action vectors onto a single dimension. The first principal component used to filter the cumulative action counts to one dimension is reported in Table 1, along with the means and standard deviations of the final action counts. Table 1 also indicates that the expert solution (“Gold Path”) is efficient in terms of the number of in-game actions performed.

Table 1. Summary statistics of the principal component used for filtering student problem-solving behaviors.

Gameplay Action	First Principal Component	Mean (SD)	Gold Path
Conversation	0.334	18.7 (5.9)	13
Reading	0.554	22.9 (8.0)	21
Worksheet	0.261	24.3 (12.5)	7
Plot Point	0.285	18.7 (1.6)	20
Worksheet Submission	0.444	2.29 (2.6)	1
Scan	0.484	26.0 (16.6)	3

By using this first principal component for filtering, the projection of the cumulative action count vector onto one dimension is guaranteed to be positive and nondecreasing because each element of the principal component is positive, and cumulative action counts are nondecreasing as students play through the game, i.e., as time in game progresses. For example, the transformed gold path final value would be 25.4, and any earlier time has at most the action counts in the final column of Table 1, and would thus have a smaller or equal transformed value. More generally, the filtration can be viewed as a function, f , converting the multi-dimensional action vector to a single value, c , using the first principal component, \mathbf{p} . This function is shown in Equation 1 for cumulative action vector \mathbf{x} of student i at time t .




Figure 3. Trajectories of students’ interactions in CRYSTAL ISLAND.

$$f(\mathbf{x}_i^t) = (\mathbf{x}_i^t)^T \mathbf{p} = c_i^t \quad (1)$$

A student trajectory is the time series of c values, where the time intervals represented by the upper index t are flexible. In this work they are calculated for every 10 seconds of gameplay. Figure 3 displays each student trajectory colored by normalized learning gain.

4.2 Trajectory Distance

Once each sequence of cumulative action vectors has been converted to the filtered time series, the lock-step Euclidean distance over the full gameplay session can be calculated. Since students played the game for varying amounts of time, the lengths of each time series may differ. In such cases, when calculating the distance between two series of unequal length, the shorter series is padded to the length of the longer series by repeating the final filtered value. The padding of the shorter sequence prevents violations of the triangle inequality from divergences of two longer sequences with a shorter sequence after the shorter sequence has ended.

The distance between two students is the average Euclidean distance between their filtered time series over all time steps. The average is taken to allow the distances to be compared from different numbers of time intervals. More specifically, the distance, d , between students i and j , can be calculated according to Equation 2, where n is the number of time intervals in the longer series. Note that while Minkowski distance of any order would yield equivalent results in this particular case of one dimension, the Euclidean norm is specifically mentioned to generalize to filters with multivariate outputs.

$$d_{ij} = \frac{1}{n} \sum_{t=1}^n \|c_i^t - c_j^t\|_2 \quad (2)$$

The distance between a student's trajectory and the golden path can be calculated by using the golden path as one of the students in Equation 2. The temporal distance calculated by Equation 2 to the golden path for student i is denoted g_i . To assess the advantage of taking the trajectory distance, or the average distance over time, a useful comparison is to the final point distance of filtered values, i.e. using only the final time step's filtered value to calculate the distance between students and the golden path. This will allow comparison between similarity measures that take into account the full gameplay over time (Equation 2) and a baseline measure (Equation 3) that does not use the full gameplay session, but instead uses a summary of gameplay. Figure 4 depicts examples of the baseline (a) and temporal distance (b) from one student trajectory to the expert solution path.

$$b_{ij} = \|c_i^n - c_j^n\|_2 \quad (3)$$

4.2.1 Trajectory Distance per Interval

Since the distance is calculated used a fixed mapping between points in time, the measurement is sensitive to misalignments in time. In other words, *local time shifting*, or similar segments that are out of place, will not be handled by the distance measure [10]. In order to account for similar segments of student trajectories out of place within CRYSTAL ISLAND, the distance over each gameplay phase is calculated. This procedure matches two students' time series from a specific phase to the same start time interval when calculating the distance over that phase, and it uses the same padding procedure described for students with differing phase lengths. Essentially each phase is treated as a "similar segment" and

distances are calculated over each phase, matching the start of one student's phase to the start of the other student's similar phase. Figure 4(d) depicts where phases end for two example trajectories, which demonstrate the start points that are matched to calculate phase-based measures.

4.3 Slope of Trajectory

The slope of a trajectory gives important insights regarding the style of problem-solving behavior of students over the course of their interaction with the game-based learning environment. Since the x-axis in this case is time, and the y-axis a filtered measure of cumulative actions, the slope represents the change of the filtered measure of cumulative actions over time. The student's slope can be viewed as a "pace of problem-solving actions," where each problem-solving action's contribution to the pace is weighted by the principal component used to project the cumulative action vector to a single dimension. For example, a student who scans many objects over a specific time span will have a steeper slope in their trajectory than a student who opens their worksheet the same amount of times over that same time interval because scans contribute more to the filtered value than worksheet opens.



A student trajectory's slope is estimated by fitting a simple linear regression with time (in minutes) as the single predictor of filtered cumulative action value. This is done by using the pairs of points (t, c^t) that create each trajectory of Figure 3 to estimate a line of best fit per student. When fitting the line of best fit over the entire gameplay or Tutorial phase, the intercept is set to 0, since students enter the game with no actions taken. In these cases, the line of best fit is given by $c = \beta t$ where c is the filtered cumulative action value, t is time in minutes, and β is the slope of the student's trajectory. In the Information Gathering and Diagnosis phase, in which a student enters with actions previously taken, the regression line includes an intercept term, $c = \beta t + b$, but the slope is the quantity of interest, which has a semantic interpretation as the pace of problem-solving behavior over that time interval.

5. RESULTS

This section analyzes key relationships between students' time series and measures from CRYSTAL ISLAND. First, the relationship between the slope of a trajectory and learning is demonstrated at both a full gameplay level and gameplay phase level. Second, the distance between the gold path and students is analyzed and compared to learning performance in CRYSTAL ISLAND. Third, an analysis of the measures against duration of gameplay is performed to evaluate the independence of the time series analysis against the length of the series. All reported correlations are Pearson product-moment correlations.


5.1 Trajectory Slope Relationship with Learning

A line of best fit through the pairs of time and filtered values were fit to each trajectory as described in Section 4.3. In addition to the line of best fit over the full trajectory (All), lines of best fit were calculated for each gameplay phase (Tutorial, Information Gathering, and Diagnosis). Since the filtered action value is calculated as a weighted sum of cumulative actions, the slope of the line of best fit can be viewed as an estimate of the pace of play of a student within the game-based learning environment with certain actions counting towards the pace more than others. It is also important to note that these slopes are independent of the golden path, but could be compared with cosine similarity as a measure independent of the duration of play. The slopes are found to be marginally significantly correlated with normalized learning gain




a. The dark green dashed line represents the padded portion of the student’s series to match the length of the golden path. The vertical blue line represents the baseline distance.

b. Each vertical gray line is averaged to calculate the final distance. There is a vertical gray line every 10 seconds, making this appear as an area between trajectories.



c. The slopes over the full gameplay episode for a student (green) and expert solution (gold).



d. Gameplay phase endpoints plotted in grayscale along a student’s trajectory (green) and expert trajectory (gold), illustrating the potential for local time shift issues in calculating distance.

Figure 4. Visual summaries of each time series characteristic calculated for comparison with measures of learning and engagement.

and have a positive cross validation R^2 indicating the generalizability of the results. The results by gameplay phase are reported in Table 2.

When analyzing the simple linear regression leave one out cross-validation R^2 measures, it is important to consider the difficulty of predicting normalized learning gain from in-game actions. More concretely, a baseline using a multiple linear regression using each cumulative action count with game duration (the features used in extracting the trajectory and slope) gives a leave-one-out cross validation R^2 of -0.089 . Note that a negative cross-validation R^2 indicates the model predictions on the held-out points have a higher mean squared error than using the variance of the data and are an indicator of poor fit.

Table 2 indicates a relationship between the slope of a trajectory and normalized learning gain. The Tutorial phase is a notable exception here, which indicates that the pace of actions during the Tutorial is not predictive of normalized learning gain. A marginally significant negative correlation between Information Gathering, Diagnosis and slope over the full gameplay session (All) with normalized learning gain demonstrates that as a trajectory slope becomes steeper, the normalized learning gain decreases. This relationship is further exemplified by the positive cross-validation R^2 results, especially relative to the baseline using the cumulative game actions and duration. Thus, a slower pace (lower slope) of students’ problem-solving behaviors measured by the filtered cumulative actions in phases beyond the Tutorial are indicative of positive learning outcomes in CRYSTAL ISLAND.

The slope of the expert solution path is the lowest observed slope of any trajectory in the dataset (0.27, next lowest = 0.31). The low slope indicates a relatively slow pace of play in terms of the number of actions taken within the game, which reflects the expert’s deliberate and efficient on-task problem-solving path. The deliberate play demonstrates positive problem-solving strategies, such as reading texts thoroughly and planning the next action.

Table 2. Summary of the relationship between trajectory slopes and normalized learning gain.

Gameplay Phase	Average Slope (SD)	Correlation with NLG (p-value)	Simple Linear Regression CV R ²
All	0.51 (0.11)	-0.22 (0.09)	0.0172
Tutorial	0.12 (0.08)	-0.063 (0.62)	-0.0362
Information Gathering	0.49 (0.11)	-0.22 (0.08)	0.0165
Diagnosis	0.58 (0.11)	-0.24 (0.05)	0.0275

5.2 Golden Path Distance Relationship with Learning

The temporal distance between the expert solution path and student trajectories was calculated as in Equation 2. There appears to be a relationship between learning, as measured by normalized learning gain, and similarity of a student trajectory with the golden path. The correlations by gameplay phase between normalized learning gain and gold path distance are given in Table 3. The leave-one-out cross-validation R² from a simple linear model using the distance as the lone predictor of normalized learning gain is also given for a measure of generalization of the correlational relationship.

Table 3. Summary of temporal distance between students and expert with normalized learning gain.

Gameplay Phase	Average Distance (SD)	Correlation with NLG (p-value)	Simple Linear Regression CV R ²
All	9.98 (4.0)	-0.23 (0.07)	0.0202
Tutorial	0.76 (0.22)	0.0061 (0.96)	-0.0781
Information Gathering	10.5 (4.8)	-0.20 (0.11)	0.0021
Diagnosis	17.3 (12.0)	-0.13 (0.42)	-0.0206

As seen from Table 3, the negative correlation between distance and normalized learning gain indicates that as student trajectories become farther from the golden path (the distance over time increases), their normalized learning gains decrease. The difference between phases is interesting to note, as the Tutorial phase and Diagnosis phase are not significantly correlated with normalized learning gain, while the Information Gathering phase demonstrates a correlation approaching significance and positive cross-validation R² superior to the baseline. The superiority of using the full gameplay for the distance calculation in Table 3

indicates that the time warping problem common among time series analysis may not be an issue in game-based learning. This is likely due to the freedom that game-based learning environments provide students, making recalibration of time intervals difficult to compare amongst students’ actions.

5.2.1 Comparison with Baseline Distance

While the relationship between the distance measure incorporating the full gameplay from the gold path and normalized learning gain is encouraging, the necessity of using temporal distance can be assessed by comparing the gold path baseline distance from Equation 3 with normalized learning gain. No significant correlation is observed between the baseline distance from the gold path with normalized learning gain ($r(61) = -0.153, p = 0.23$). A baseline comparison using each student’s final filtered cumulative action value as a single predictor in an ordinary least squares regression evaluated using leave-one-out cross-validation gives an R² of -0.0075. The lack of relationships demonstrated with the baseline distance compared to the correlation of the temporal distance indicates that using the distance from the expert solution over the full gameplay session provides valuable information for predicting normalized learning gain.

5.3 Comparison with Final Game Score

The *final game score* is an in-game measure designed by domain experts specifically for the CRYSTAL ISLAND game-based learning environment to assess student engagement [25]. Thus, comparisons with the *final game score* provide a complementary comparison to normalized learning gain from the actions in CRYSTAL ISLAND to gauge a student’s experience. First, it should be noted that a marginally significant positive correlation was observed between normalized learning gain and *final game score* ($r(61) = 0.25, p = 0.05$), indicating that students with a high *final game score* have higher normalized learning gains. The magnitudes of the correlations observed with the slope and expert solution distance are similar to the correlation observed between *final game score* and normalized learning gain, despite *final game score* being a hand-crafted measure of performance in CRYSTAL ISLAND while the trajectories are automatically created from student data. This is also seen when comparing the leave-one-out cross-validation R² of using *final game score* as the sole predictor in a simple linear regression model, which yields a 0.0265 value when predicting normalized learning gain.

Table 4. Summary of time series characteristics with final game score.

Condition	Slope-based Linear Regression CV R ²	Distance-based Linear Regression CV R ²
All	0.0091	0.28
Tutorial	0.030	0.028
Information Gathering	0.021	0.28
Diagnosis	0.064	0.51

The golden path reflects a trajectory with desirable problem-solving behaviors according to the *final game score* as the expert takes an efficient solution path. For example, the expert uses far less scans of irrelevant virtual objects and incorrect worksheet submissions than the average student, both of which are penalized

by the *final game score* for being indicative of guess-and-check behavior. This can be observed by the strong predictive power of the temporal distance to the expert solution path over the *final game score* given in Table 4. These results are notably strong when compared with the slope of the trajectories, which has weaker predictive power over student in-game engagement as measured by *final game score*. The relationship between distance to the expert solution and *final game score* increases as students progress through the phases of CRYSTAL ISLAND. This is likely because students perform actions that more directly impact the *final game score* (scans and worksheet submissions) during the final Diagnosis phase, which is captured by taking the distance over this interval.

6. DISCUSSION

In this work, students' problem-solving behaviors in Crystal Island were transformed into time series representing their trajectories through the problem-solving space. This section provides explanations, considerations, and implications of the results from comparing characteristics of these trajectories with learning outcomes.

6.1 Trajectory Slope

The results suggest that the slope of a student's problem-solving trajectory contains valuable information about their approach to problem solving in the game-based learning environment. The slope of a student's problem-solving trajectory was found to be marginally predictive of normalized learning gain using the full gameplay, Information Gathering phase, and Diagnosis phase. Negative slopes were found to be predictive of higher learning gains, indicating that students who performed more problem-solving actions (weighted through the principal component) per minute had worse learning outcomes.

While the slopes were calculated independently of the expert solution, it is interesting to note that the expert solution had the most gradual slope of any problem-solving trajectory. Therefore, the cosine similarity of best fit lines through trajectories would yield similar results to the current analysis of the slopes, which is independent of the expert solution because steeper slopes would be more dissimilar. Thus, the cosine similarity in this particular context would be analogous to subtracting a constant from each slope, which would not affect the measures used for the analysis in this work. Since these slopes are based on univariate time series, there is no additional information that an analysis of the cosine similarity would provide over an analysis of the slopes themselves. However, the current expert path is only one possible problem-solving solution through this space, and in future work it would be informative to conduct an analysis using solution paths that vary by problem-solving strategy, including negative solution paths, such as a guess-and-check methodology.

The slope during the Information Gathering phase was negatively correlated with learning outcomes. This is an interesting observation given the nature of the Information Gathering phase, where students do not perform any scans in the virtual laboratory. (If they had performed scans, they would be considered to be in the Diagnosis phase). While the steeper slopes indicate a problem-solving strategy more in line with a guess-and-check method, this phase by definition does not include guesses through the scanner. This indicates that the slope of the trajectory includes additional information over identifying potential guess-and-check strategies. A more gradual slope in the Information Gathering phase could be caused by students who are more deliberate in fully reading and comprehending their conversations and reading materials, which would contribute to the negative relationship between trajectory

slope and learning outcomes in this phase. This observation is in line with previous research on CRYSTAL ISLAND, which found that information gathering prior to hypothesis generation was correlated with improved problem-solving efficiency [26].

The weak relation between slope trajectory and *final game score* is surprising given the way *final game score* and the filtered cumulative action counts are calculated. *Final game score* penalizes incorrect scans in the virtual laboratory and incorrect worksheet submissions, which are both actions weighted heavily in the filtered cumulative action count. Thus, one would expect a steeper slope to indicate a lower *final game score* since the steep slope indicates problem-solving behaviors likely to have a negative impact on *final game score* being performed at a quicker rate than other students. However, this may be offset by the *final game score* rewarding problem-solving efficiency, which would be indicated by a steeper slope.

6.2 Distance from Expert Solution

The results have important implications regarding the temporal distance of a student's problem-solving trajectory and the expert solution problem-solving trajectory. Since this distance represents the dissimilarity of the student's problem-solving path over time relative to an expert's, the negative correlations between dissimilarity and learning outcomes are as one would expect: as a student's problem-solving path becomes more similar to the expert solution, the student's learning outcomes are expected to be higher. Thus, the results suggest that analyzing a student's problem-solving path in game-based learning with respect to an expert's problem-solving path can yield insight into student learning outcomes, which are measured outside of the game-based learning environment. Interaction with CRYSTAL ISLAND centers on solving a complex problem with multiple solution paths, and the expert solution represents one of many possible paths. Further work should be done in evaluating student solution paths in the context of multiple expert solution paths.

The differences between the temporal distance measure and baseline measure indicate that the temporal distance incorporates additional information regarding the problem-solving behavior path. The baseline distance does not capture information regarding intermediate steps of the problem-solving path, which are critical to learning. This is analogous to only checking if a student obtained the correct answer to a problem without considering the steps the student took to solving the problem. In the context of an ill-structured problem, the temporal distance supports a comparison between the steps students took over the course of gameplay with an expert solution rather than merely considering the final summary statistics of a student.

6.3 Heteroskedasticity of Trajectories

The current filtered cumulative action count provides several benefits such as its interpretability as a nondecreasing measure of weighted problem-solving behaviors performed. However, the trajectories become more dispersed as students follow different problem-solving paths through the game. The wide dispersion is a consequence of the open-ended nature of CRYSTAL ISLAND, which has many valid solution paths defined by trajectories. While this dispersion of trajectories is important for revealing the divergence of problem-solving paths among different students, the dispersion as time increases indicates heteroskedasticity in the filtered values, or an increase in variance among the filtered cumulative action values per time step.

This can be observed in Table 3, where the standard deviation of the distance from the expert solution increases per gameplay phase.

For example, in the Information Gathering phase, the standard deviation of the 63 student trajectory distances from the expert solution is 4.8, and this more than doubles to 12.0 in the Diagnosis phase. Future work should address whether this heteroskedasticity is desired in calculating similarities from distances or whether a variance-adjusted distance would be more appropriate to account for how the population of trajectories become more dispersed as time progresses. For example, the increased variance of distance in later phases may be the cause of the expert distance during Information Gathering being significantly predictive of normalized learning gain while the Diagnosis phase has no predictive ability over normalized learning gain. On the other hand, the distance between students and the expert path in the Diagnosis phase explains more the variance of the *final game score* than the Information Gathering phase, indicating that the wide dispersion of filtered values does not have a negative impact on the relationship between expert distance and *final game score*.

6.4 Implications of Time Series Analysis

The primary result of this work is that the trajectory of a student through the problem-solving space of a game-based learning environment has a relationship with the measured learning outcomes of normalized learning gain and significant relationship with *final game score*. The framework for creating these trajectories is generalizable to game-based learning environments tracking cumulative game actions of students as well as a broad range of advanced learning technologies that support multiple problem-solving paths. Importantly, this includes transforming an expert problem-solving solution path into the same problem-solving space as student paths, and quantifying the similarity of a student solution path relative to the expert solution. While this one expert path represents only one possible solution path through the problem-solving space, this similarity predicts normalized learning gain, indicating the potential for evaluating a student's entire problem-solving path in an open-ended game-based learning environment. The measures used here were shown to be predictive of learning outcomes, but further analysis should be done to determine qualitative characteristics related to learning and self-regulatory processes.

These observations have important design implications for adaptive learning environments. For example, the results suggest that one approach to improving student learning would involve an adaptive learning environment scaffolding a student's problem solving to increase the probability that the student follows a trajectory more closely related to an expert problem-solving path. In the context of a reinforcement learning-based tutorial planner [24, 27], characteristics of the trajectory defined by the filtered cumulative action value could be used as continuous state variables. This work has shown the problem-solving trajectory slope and distance to an expert solution are related to learning and in-game student engagement, suggesting that problem-solving trajectory slope and distance to an expert solution are useful variables to include in a state representation for a tutorial planner. The impact of decisions made by the tutorial planner on the student's trajectory in terms of its slope and distance from an expert solution could thereby be used as estimates for the transitions of a decision in a model-based reinforcement learning framework.

These results also have another key implication for the design of adaptive learning environments. In a recent study with the CRYSTAL ISLAND game-based learning environment, students who followed a predetermined path achieved significantly higher normalized learning gains than students who had freedom of control [29]. These results suggest a possible explanation for the higher observed

learning gains: students on the predetermined path followed a problem-solving trajectory more similar to the expert solution path than students who were given freedom to explore. Therefore, the effectiveness of an expert solution path could be measured using this framework for time series analysis of problem-solving behaviors, and the solution path could be considered for a limited agency design of a game-based learning environment.

7. CONCLUSION

Open-ended game-based learning environments allow a wide range of problem-solving behaviors. Analyzing student actions within a game-based learning environment can thus provide insight into students' learning processes. Incorporating the sequential nature of student actions within the game-based learning environment is important because of the complexities of the problem-solving process. This work addresses these issues by examining the dynamics of problem-solving behavior of students within a game-based learning environment through a filtered time series analysis.

A general framework for filtering problem-solving behaviors into a gameplay trajectory was presented using a dimensionality reduction filter. The slope of this trajectory, representing the pace of problem-solving behaviors, was shown to be negatively correlated with learning, indicating that students who were more deliberate in the rate of problem-solving behaviors achieved higher learning gains. The similarity of student problem-solving trajectories with an expert solution was shown to be correlated with learning, indicating students who took a similar solution path to the expert demonstrated higher learning gains. A comparison of temporal distance, using the sequential nature of the problem-solving process, and a baseline distance, using a final summary of student problem-solving process, demonstrated the utility of incorporating the temporal nature of interactions within a game-based learning environment. The results demonstrate the value of analyzing the characteristics of a student's path through the problem-solving space in the context of an expert path. In future work, it will be important to investigate how the results of time series analyses can most effectively inform runtime learning environment adaptations.

8. ACKNOWLEDGMENTS

We would like to thank our collaborators in the Center for Educational Informatics and the SMART Lab at North Carolina State University. This study was supported by funding from the Social Sciences and Humanities Research Council of Canada. Any conclusions in this material do not necessarily reflect the views of the SSHRC.

9. REFERENCES

- [1] Abdi, H. and Williams, L. 2010. Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*. 2, 4, 433–470.
- [2] Andersen, E., Liu, Y., Apter, E., Boucher-Genesse, F. and Popović, Z. 2010. Gameplay analysis through state projection. *Proceedings of the 5th Int. Conference on the Foundations of Digital Games*, 1–8.
- [3] Bauer, A., Flatten, J. and Popovic, Z. 2017. Analysis of problem-solving behavior in open-ended scientific-discovery game challenges. *Proceedings of the 10th Int. Conference on Educational Data Mining*, 32–39.
- [4] Beal, C., Mitra, S. and Cohen, P. 2007. Modeling learning patterns of students with a tutoring system using hidden

- Markov models. *Proceedings of the 13th Int. Conference on Artificial Intelligence in Education*, 238–245.
- [5] Box, G., Jenkins, G. and Reinsel, G. 1994. Time Series Analysis - Forecasting and Control. *Prentice Hall New Jersey 1994*. SFB 373, Chapter 5, 837–900.
- [6] Boyer, K., Phillips, R., Ingram, A., Ha, E., Wallis, M., Vouk, M. and Lester, J. 2011. Investigating the relationship between dialogue structure and tutoring effectiveness: A hidden markov modeling approach. *International Journal of Artificial Intelligence in Education*. 21, 1–2, 65–81.
- [7] Bradbury, A., Taub, M. and Azevedo, R. 2017. The effects of autonomy on emotions and learning in game-based learning environments. *39th Annual Meeting of the Cognitive Science Society*, 1666–1671.
- [8] Clark, D., Tanner-Smith, E. and Killingsworth, S. 2016. Digital games, design, and learning: A systematic review and Meta-Analysis. *Review of Educational Research*. 86, 1, 79–122.
- [9] Corbett, A. and Anderson, J. 1995. Knowledge-tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User Adopted Interaction*. 4, 253–278.
- [10] Ding, H., Trajcevski, G., Scheuermann, P., Wang, X., Keogh, E. 2008. Querying and mining of time series data: experimental comparison of representations and distance measures. *Proceedings of the VLDB Endowment*. 1, 2, 1542–1552.
- [11] Ha, S., Bae, S. and Park, S. 2000. Web mining for distance education. *Proceedings of the IEEE Int. Conference on Management of Innovation and Technology*, 715–719.
- [12] Hambrick, D., Burgoyne, A., Macnamara, B. and Ullén, F. 2018. Toward a multifactorial model of expertise: beyond born versus made. *Annals of the New York Academy of Sciences*, 1–12.
- [13] Hansen, C., Hansen, C., Hjulær, N., Alstrup, S. and Lioma, C. 2017. Sequence modelling for analysing student interaction with educational systems. *Proceedings of the 10th Int. Conference on Educational Data Mining*, 232–237.
- [14] Jeong, H. and Biswas, G. 2008. Mining Student Behavior Models in Learning by Teaching Environments. *Proceedings of the 1st Int. Conference on Educational Data Mining*, 127–136.
- [15] Khajah, M., Lindsey, R., and Mozer, M. 2016. How deep is knowledge tracing? *Proceedings of the 9th Int. Conference on Educational Data Mining*, 94–101.
- [16] Kinnebrew, J., Segedy, J. and Biswas, G. 2014. Analyzing the temporal evolution of students’ behaviors in open-ended learning environments. *Metacognition and Learning*. 9, 2, 187–215.
- [17] Köck, M. and Paramythis, A. 2011. Activity sequence modelling and dynamic clustering for personalized e-learning. *User Modeling and User-Adapted Interaction*. 21, 1–2, 51–97.
- [18] Lallé, S., Conati, C. and Carenini, G. 2016. Predicting confusion in information visualization from eye tracking and interaction data. *Proceedings of the 25th Int. Joint Conference on Artificial Intelligence*, 2529–2535.
- [19] Liu, Y., Andersen, E. and Snider, R. 2011. Feature-based projections for effective playtrace analysis. *Int. Conference on Foundations of Digital Games*, 69–76.
- [20] Malkiewich, L., Baker, R., Shute, V., Kai, S. and Paquette, L. 2016. Classifying behavior to elucidate elegant problem solving in an educational game. *Proceedings of the 9th Int. Conference on Educational Data Mining*, 448–453.
- [21] Martinez, R., Yacef, K. and Kay, J. 2011. Analysing frequent sequential patterns of collaborative learning activity around an interactive tabletop. *Proceedings of the 4th Int. Educational Data Mining*, 111–120.
- [22] Perera, D., Kay, J., Koprinska, I., Yacef, K. and Zaïane, O.R. 2009. Clustering and sequential pattern mining of online collaborative learning data. *Knowledge and Data Engineering, IEEE Transactions on*. 21, 6, 759–772.
- [23] Romero, C. and Ventura, S. 2007. Educational data mining: A survey from 1995 to 2005. *Expert Systems with Applications*. 33, 1, 135–146.
- [24] Rowe, J. and Lester, J. 2015. Improving student problem solving in narrative-centered learning environments: A modular reinforcement learning framework. *Proceedings of the 17th Int. Conference on Artificial Intelligence in Education*, 419–428.
- [25] Rowe, J., Shores, L., Mott, B. and Lester, J. 2011. Integrating learning, problem solving, and engagement in narrative-centered learning environments. *Int. Journal of Artificial Intelligence in Education*. 21, 1–2, 115–133.
- [26] Sabourin, J., Rowe, J., Mott, B., Lester, J., Carolina, N. and Carolina, N. 2012. Exploring inquiry-based problem-solving strategies in game-based learning environments. *Proceedings of the 11th Int. Conference on Intelligent Tutoring Systems*, 470–475.
- [27] Sawyer, R., Rowe, J. and Lester, J. 2017. Balancing learning and engagement in game-based learning environments with multi-objective reinforcement learning. *Proceedings of the 18th Int. Conference on Artificial Intelligence in Education*, 323–334.
- [28] Sawyer, R., Smith, A., Rowe, J., Azevedo, R. and Lester, J. 2017. Enhancing student models in game-based learning with facial expression recognition. *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, 192–201.
- [29] Sawyer, R., Smith, A., Rowe, J., Azevedo, R. and Lester, J. 2017. Is more agency better? The impact of student agency on game-based learning. *Proceedings of the 18th Int. Conference on Artificial Intelligence in Education*, 335–346.
- [30] Snow, E., Likens, A., Jackson, G. and McNamara, D. 2013. Students’ Walk through tutoring: Using a random walk analysis to profile students. *Proceedings of the 6th Int. Conference on Educational Data Mining*, 276–279.
- [31] Snow, E., Likens, A., Allen, L. and McNamara, D. 2016. Taking control: Stealth assessment of deterministic behaviors within a game-based system. *Int. Journal of Artificial Intelligence in Education*. 26, 4, 1011–1032.
- [32] Tóth, K., Greiff, S., Kalergi, C. and Wüstenberg, S. 2014. Discovering students’ complex problem solving strategies in educational assessment. *Proceedings of the 7th Int. Conference on Educational Data Mining*, 225–228.