A Multi-View Predictive Student Modeling Framework with Interpretable Causal Graph Discovery for Collaborative Learning Analytics

Halim Acosta
North Carolina State University
Raleigh, North Carolina, USA
hacosta@ncsu.edu

Seung Lee
North Carolina State University
Raleigh, North Carolina, USA
sylee@ncsu.edu

Daeun Hong
Indiana University
Bloomington, Indiana, USA
dh37@iu.edu

Wookhee Min
North Carolina State University
Raleigh, North Carolina, USA
wmin@ncsu.edu

Bradford Mott
North Carolina State University
Raleigh, North Carolina, USA
bwmott@ncsu.edu

Cindy Hmelo-Silver
Indiana University
Bloomington, Indiana, USA
chmelosi@iu.edu

James Lester
North Carolina State University
Raleigh, North Carolina, USA
lester@ncsu.edu

ABSTRACT

Understanding the relationship between student behaviors and learning outcomes is crucial for designing effective collaborative learning environments. However, collaborative learning analytics poses significant challenges, not only due to the complex interplay between collaborative problem-solving and collaborative dialogue but also due to the need for model interpretability. To address these challenges, this paper introduces a multi-view predictive student modeling framework using causal graph discovery. We first extract interpretable behavioral features from students' collaborative dialogue data and game trace logs to predict student learning within a collaborative game-based learning environment. We then apply constraint-based sequential pattern mining to identify cognitive and social behavioral patterns in student's data to improve predictive power. We employ unified causal modeling for interpreting model outputs, using causal discovery methods to reveal key interactions among student behaviors that significantly contribute to predicting learning outcomes and identifying frequent collaborative problem-solving skills. Evaluations of the predictive student modeling framework show that combining features from dialogue and in-game behaviors improves the prediction of student learning gains. The findings highlight the potential of multi-view behavioral data and causal analysis to improve both the effectiveness and the interpretability of collaborative learning analytics.

Halim Acosta, Seung Lee, Daeun Hong, Wookhee Min, Bradford Mott, Cindy Hmelo-Silver, and James Lester. A Multi-View Predictive Student Modeling Framework with Interpretable Causal Graph Discovery for Collaborative Learning Analytics. In Caitlin Mills, Giora Alexandron, Davide Taibi, Giosuè Lo Bosco, and Luc Paquette (eds.) Proceedings of the 18th International Conference on Educational Data Mining, Palermo, Italy, July, 2025, pp. 42–54. International Educational Data Mining Society (2025).

https://doi.org/10.5281/zenodo.15870161

Keywords

Causal graph discovery, Collaborative problem solving, Predictive student modeling, Sequential pattern mining

1. INTRODUCTION

Collaborative problem solving involves a complex interplay between cognitive and social dimensions of group behavior during collaborative learning [2, 3]. Improving students' collaborative problem-solving abilities can in turn improve their critical thinking skills while fostering essential soft skills such as teamwork and communication, which are crucial for success in the 21st century workplace [15, 19, 32]. If adaptive learning environments had access to accurate predictive student models that could effectively analyze students' fine-grained problem-solving interaction data and text-mediated communication [5, 32, 33], they could provide adaptive scaffolding that could enhance student learning and support teachers in classrooms. However, developing predictive models for collaborative learning presents significant challenges due to both computational complexity and the need for model interpretability.

Multi-view machine learning can provide a powerful foundation for predictive student modeling that enables the integration of diverse data sources to capture complementary aspects of student learning behaviors [14, 30]. It can incorporate different views of the data describing the same entity from multiple perspectives to capture the interrelations among data sources best describing the construct under investigation. In addition, it can accommodate a broad range of supervised, semi-supervised, and clustering-based techniques [16, 26, 35]. By leveraging multiple perspectives, multi-view predictive student modeling can support rich representations of student engagement and collaboration and can overcome issues with single-view techniques that have difficulty incorporating features across multiple perspectives simultaneously. This allows multi-view methods to effectively improve the accuracy, robustness, and generalizability of downstream learning analytics tasks. However, these techniques pose significant challenges, such as the curse of dimensionality, view consistency, and complementarity [38]. Our work incorporates behavioral patterns extracted from dialogue and game trace log information to represent social and cognitive aspects of students' behaviors while interacting with a collaborative game-based learning environment. We explore kernel multiple canonical correlation analysis (kMCCA) and a deep learning-based approach to incorporate each view to predict student learning outcomes.

Sequential pattern mining (SPM) identifies recurring behaviors and learning trajectories in student interactions, providing insights into their learning processes and outcomes [31, 39]. In collaborative game-based learning, SPM extracts key engagement patterns, but traditional methods often generate too many irrelevant patterns. Constraint-based sequential pattern mining (CSPM) overcomes this by incorporating domain knowledge and filtering out less informative patterns [36]. Our research applies CSPM to extract meaningful behavioral patterns from dialogue and game trace data, facilitating the identification of key predictors of student learning.

Predictive student modeling aims to infer learning outcomes from behavioral data, making the extraction of salient features that effectively represent these outcomes well suited for a causal modeling approach. Causal modeling provides an effective avenue for better understanding the impact of learning behaviors on student learning outcomes [21, 34]. Traditional causal modeling approaches often rely on domain knowledge by manually specifying causal relationships in directed acyclic graphs. However, this process is often labor-intensive, requiring domain expertise and iterative refinement to develop effective representations of the underlying data. Discovery-based methods instead learn relationships directly from observational data to reveal the underlying data-generating process [10]. Causal discovery techniques have been used to perform counterfactual and interventional reasoning, but it is difficult due to structure identifiability and computational complexity [23]. Our work employs CSPM techniques to abstract low-level interaction and dialogue data to extract a smaller subset of features for causal structure discovery.

This paper introduces a multi-view predictive student modeling framework that integrates multi-view learning and causal discovery methods to create an interpretable framework for predicting learning outcomes from fine-grained interaction and dialogue information in a collaborative learning context. Building on previous work mapping behavioral process data along multiple dimensions of collaborative problem solving, we align in-game actions and dialogue with the social and cognitive dimensions of a collaborative problem-solving framework [2]. This abstraction method generates an interpretable sequence of dialogue and game trace data that can be used to extract pedagogically meaningful behavioral patterns. The multi-view predictive student modeling framework considers dialogue and game trace log information to represent two distinct views (cognitive and social) of students' collaborative behaviors. Utilizing CSPM techniques, we collect a set of candidate behavioral predictors of student learning outcomes. We evaluate the effectiveness of these features using four machine learning models and estimate causal relationships using the Peter-Clark (PC) algorithm [25]. Additionally, we link causal structures with feature importance scores to better understand how causal structures influence predictive modeling outputs, enhancing interpretability. We address the following research questions:

RQ1: What impact does the use of CSPM have on identifying candidate cognitive and social behavioral predictors of student learning outcomes in collaborative learning contexts?

RQ2: How effective is the multi-view learning approach in integrating interaction and dialogue data to predict student learning outcomes in collaborative learning contexts?

RQ3: How does combining interpretable causal discovery with multi-view predictive modeling reveal underlying cause-effect dynamics between collaborative behaviors and student learning outcomes prediction?

2. RELATED WORK

Multi-view machine learning (MVML) has emerged as a powerful approach for analyzing complex educational data by integrating information from multiple sources or perspectives [27]. These techniques have been applied to enhance student performance prediction and understanding of collaborative behaviors. Recently, Venkatachalam & Sivanraju introduced an enhanced generative adversarial network with an improved semi-automatic deep learning model based on a multi-view approach to predict student academic performance [30]. Their model effectively integrates heterogeneous student behavioral logs and assessment scores to provide a more comprehensive analysis of student performance. Other research in this area explored the use of co-training for predicting student performance by combining two distinct feature views in learning management systems [14]. This approach also enables early identification of struggling students, offering opportunities for timely interventions to improve learning outcomes.

Sequential pattern mining (SPM) [39] has been successfully applied in computer-supported collaborative learning and adaptive learning environments [7, 11, 12]. Additionally, these methods have been used to extract meaningful patterns from dialogue [28, 40, 41] and interaction log data [18, 37]. They have proven effective in analyzing students' progression through problem-solving activities, self-regulated learning processes, and capturing local learning patterns that link log-generated information to specific learning outcomes [24]. Recent studies have explored the use of SPM in online learning environments to provide valuable insights into learner behavior. This includes work that has focused on applying SPM to navigational patterns in learning management systems [20]. Research in science education has used SPM to analyze student behavior during scientific calculations, highlighting the importance of guidance usage for successful task completion [31].

Recent research has highlighted the growing importance of causal modeling in learning analytics. Causal modeling helps bridge the gap between learning analytics and learning theory by providing interpretable insights [13]. Causal discovery methods, which extract relationships among variables without predefined structures, offer a powerful approach to understanding complex systems [9]. In learning analytics, causal inference plays an important role in identifying the impact of different factors on learning outcomes. Insights derived from causal discovery methods can aid in the design of adaptive learning environments and other digital learning platforms [6]. Prior work has emphasized the distinction between causal discovery and causal inference, with both approaches proving essential for moving beyond correlation-based machine learning models [17]. In addition, causal inference has informed education policy [29]. The present work introduces a novel multi-view predictive student modeling framework that incorporates causal graph discovery to interpret complex relationships between student behaviors. Unlike prior work that has focused separately on dialogue or problem-solving behaviors, this approach offers a unified method for improving the predictive power and interpretability of collaborative learning analytics.

3. COLLABORATIVE GAME-BASED LEARNING ENVIRONMENT

The EcoJourneys collaborative game-based learning environment is designed to improve students' understanding of life science topics and enhance their collaborative problem-solving skills [22]. It uses a problem-based learning inquiry cycle that emphasizes complex problem-solving through active collaboration and dialogue. Students work in groups of up to four, with each student using their own laptop to engage with the unfolding narrative, while exploring a virtual tropical island, taking notes, watching videos, and talking to non-player characters (NPCs) who act as local experts as students work to solve a mysterious illness affecting fish on the island. Students must collect and analyze information while discussing their findings via a persistent in-game chat interface. Initially, students work independently before coming together at predefined intervals to solve the game objectives, while explaining their reasoning process collaboratively. The collaborative game-based learning experience is organized into three phases: 1) Talk & Investigate, 2) Deduce, and 3) Explain. The Talk & Investigate encompasses non-explicitly collaborative activities where students explore the virtual environment while collecting clues and information about their diagnosis. The Deduce and Explain phases are explicit collaborative activities where students must work together to answer questions and reach a consensus on their proposed hypothesis. The game-based learning environment includes four activities, including a tutorial and three quests.

During the Deduce phase, students must collaboratively answer multiple-choice questions that help them interpret the data they have collected for the final Explain phase of the game. Students share relevant information and ideas while negotiating differences in opinion. Once they have reached a consensus on their hypothesis, the game provides feedback on the validity of their answers, prompting students to revisit the task if their answers are incorrect. The Deduce phase ends with students collaboratively answering a constructed response question. In the Explain phase, students utilize a virtual whiteboard to structure the information they have gathered supporting their hypothesis claims. They argue for or against the claim and must again come to a consensus regarding the support of a particular claim for their hypothesis while explaining their reasoning. In this study, we specifically focus on the collaborative activities of students' interactions with the EcoJourneys learning environment, only using data extracted from the Deduce phase of each activity for downstream predictive modeling.

Our study uses data collected from 75 middle school students in sixth through eighth grades (ages 11-14) while they interacted with the EcoJourneys collaborative game-based learning environment. For each participant, their primary caregivers were provided informed consent forms and students completed assent forms, as approved by the university's ethics review board. These forms outlined the research objective and procedures of the study, including participation in data collection and surveys. From these interactions, student game trace log information captures a wide range of in-game actions such as interacting with NPCs, watching videos, sending/receiving chat messages, and navigating the virtual environment, as well as miscellaneous information such as user interface interactions and game progress monitoring. Students' text-based communications through the persistent in-game chat interface allow for the capture of collaborative dialogue practices, which can provide insights into group collaborative dynamics. On average, there are 663 trace log events and 33 dialogue contributions per student during the Deduce phase of their interactions with the learning environment. We conceptualize a multi-view modeling approach by considering dialogue and game trace log information as different views of students' social and cognitive collaborative problem-solving behaviors. Both cognitive and social behavioral views are analyzed for our downstream predictive modeling approaches.

4. METHODOLOGY

We now outline our methodology for transforming raw student interaction data into interpretable features for predictive and causal modeling. We begin by describing the abstraction process for dialogue and game trace data, followed by the application of constraint-based sequential pattern mining to extract meaningful behavioral patterns. Next, we detail our predictive modeling approach using multi-view learning and conclude with our causal discovery framework to identify underlying relationships between student behaviors and learning outcomes.

4.1 Game Trace & Dialogue Abstraction

The EcoJourneys learning environment collects students' game trace information through semi-structured CSV files. These files encode event information with a timestamp and key-value pairs containing mixed data types. This heterogeneous representation complicates downstream sequential pattern mining algorithms by containing a large search space of temporal states and can result in uninterpretable outputs, limiting its use in learning analytics. We introduce an abstraction method that streamlines sequential analysis by transforming fine-grained, semi-structured behavioral data into coarse-grained yet more interpretable action sequences. Each trace log event is first mapped to a high-level action intent through cluster-based analysis. Text-based representations of each game trace log are embedded using a Word2Vec model trained on reserved game trace data to capture contextual and semantic information from each trace event. K-means clustering is used to categorize each trace event, while silhouette score and latent-Dirichlet allocation are used to devise interpretable cluster names that describe their high-level intent. For each trace-log cluster, a domain-expert reviewed the extracted topics and manually created a descriptive category that encapsulates the overall intent of the corresponding trace logs. Next, each trace log event is annotated with the high-level intent category to produce an interpretable sequence of in-game actions from raw trace log data. Finally, a qualitative analysis by domain experts established a mapping from high-level intent categories to four collaborative problem solving (CPS) codes within the cognitive dimension of the collaborative problem-solving ontology [2], which serves as our guiding pedagogical framework for analyzing students' collaborative behaviors (Table 1).

Table 1 presents the mapping of trace log clusters and their high-level intent to the CPS codes within the ontology proposed by Andrews-Todd & Forsyth [2]. First, the Exploring and Understanding CPS code is linked to activities involving interacting with in-game objects, learning resources, and non-player characters as they reflect students' active engagement in seeking information and deepening their understanding of the learning environment. Second, the Representing and Formulating CPS code is associated with actions reflecting answer formulation and submission, including clusters such as "Interaction & Exploration Cessation" and "Answer Submission and Validation." Third, the Monitoring CPS code refers to actions that track progress, evaluate understanding, and/or regulate actions to achieve goals. Finally, we categorize active and passive communication under the broad Social category. This qualitative analysis facilitates the interpretation of raw game trace log information in the context of a relevant collaborative problem-solving theory, allowing for interpretable and actionable insights.

To effectively capture salient information from the active and passive communication under the Social CPS code, we further analyze the collaborative dialogue. Contributions to collaborative dialogue can be extracted from game trace log information, providing additional points for analyzing dialogue dynamics. However, raw dialogue sequences can be challenging to interpret without domain knowledge and often contain significant noise while also failing to capture pedagogically meaningful characteristics, such as how students negotiate, share ideas, and regulate group efforts.

We overcome this limitation by devising an LLM-based dialogue act recognition model to identify CPS-related dialogue acts, representing the communicative intents of utterances [4]. We extract each student's sent message from their game trace logs and apply zero-shot prompting to a Llama 3.1 8B model to assign one of six dialogue acts under the Social CPS code (Table 2). In this way, we transform a sequence of raw input utterances into a more interpretable sequence of CPS dialogue practices. Llama 3.1 was selected for its strong performance on benchmark datasets and its open-source availability, offering advantages over proprietary LLMs and legacy natural language processing models. While fine-tuning was beyond the scope of this work, it presents a promising direction for future research. Moreover, model quantization enables efficient inference on limited computational resources, making it well-suited for deployment in educational settings. In conjunction with abstracted student game trace logs, dialogue sequence information is used to extract sequential behavioral patterns along cognitive and social dimensions for downstream predictive modeling and causal discovery.

Utterances classified as "Maintaining Communication" are contributions that encourage or support others during their collaborative tasks. "Sharing Ideas" refers to utterances that attempt to share ideas, resources, or information about ongoing tasks. "Negotiating Ideas" are dialogue contributions that clarify, correct, or elaborate on ideas presented during group discussions. "Regulating" behaviors attempt to establish a shared understanding by organizing, planning, or evaluating the group's progress. Finally, "Off-task" refers to any utterances that are not on topic but still contribute to building social rapport. Any utterance that cannot be classified into one of the previous five categories is simply labeled as "Other".

These codes capture the discursive practices that align with the social dimension of the CPS ontology and, when combined with abstracted game trace log information, provide a more holistic understanding of students' collaborative behavior contributions.

4.2 CSPM Behavioral Pattern Extraction

Previous work has shown the efficacy of using text-mediated communication and process data as evidence of collaborative problem-solving practices and mapping this evidence along each dimension of the CPS framework [2, 8]. Similarly, we map collaborative dialogue contributions and low-level in-game actions to each dimension of the CPS ontology. Although this mapping generates an interpretable sequence of collaborative contributions, identifying meaningful and pedagogically relevant behavioral patterns within these sequences remains challenging due to the sheer volume of possible patterns and the presence of redundant or low-impact sequences. To address this, we employ constraint-based sequential pattern mining techniques to extract cognitively and socially relevant behavioral patterns that are most indicative of student learning outcomes. We utilize Seq2Pat [36], a CSPM technique based on multi-valued decision diagrams that find frequent patterns within a sequence database subject to specified constraints. In this context, constraints refer to domain knowledge such as specified orderings, cycles, or otherwise a priori behavioral information that can be used to restrict the search space of possible sequential patterns. Additionally, Seq2Pat enables efficient conversion of sequential patterns into binary-valued feature vectors that indicate the presence or absence of each detected pattern for individual students. These features support downstream predictive modeling and causal discovery.

Table 1: Mapping of trace log clusters to CPS codes along the cognitive and social dimensions of the CPS ontology.

CPS Code	Description	Trace Log Cluster
Social	Communication through the in-game chat	Passive Communication Receipt
		Active Communication Initiation
Exploring & Understanding	Exploring the game environment by interacting with objects and resources	Termination of Specific Interactions
		Workbook Engagement and Note-Taking
		Active Exploration and Interaction Initiation
		Tutorial Progression Engagement
Representing & Formulating	Answering and submitting multiple-choice or constructed-response questions	Answer Submission and Validation
		Interaction and Exploration Cessation
Monitoring	Receiving feedback from NPC conversations	NPC Dialogue and Conversation Initiation

Table 2: CPS codes assigned to dialogue utterances within the social dimension of the CPS ontology.

CPS Category	CPS Code	Description
Maintaining Communication	*Maintaining*	Encourages or supports others during the task.
Sharing Resources/Ideas	*Sharing*	Shares ideas, resources, or information about the task.
Negotiating Ideas	*Negotiating*	Clarifies, corrects, or elaborates on ideas.
Regulating Problem Solving Activities	*Regulating*	Organizes, plans, monitors, or evaluates progress.
Off-task	*Off-task*	Not related to the task.
Other	*Other*	Does not fit into any category.

We extract behavioral patterns by independently applying CSPM to abstracted dialogue and game trace log sequences. Dialogue patterns can help to identify effective collaboration cycles or breakdowns in engagement and may be influential in predicting student learning outcomes. To filter out spam chat messages, we apply a count constraint on the minimum (3) and maximum (20) number of words per sequence. Often, students spam chat by either sending many small messages or a small number of very large messages, typically containing non-intelligible words. We apply this count constraint to avoid identifying spurious patterns. We should note that in some cases, it may be preferable to detect these types of disruptive behaviors. However, a preliminary analysis showed that this often introduces a large amount of noise, negatively affecting downstream predictive modeling and causal discovery methods. In addition to the word count constraint, we further impose a minimum support count (2) and maximum span (8) constraint on extracted patterns to identify patterns that occur across two or more individuals and to restrict patterns to only cover up to eight utterances, both of which were likewise determined through preliminary analysis. Restricting the maximum span helps to ensure that patterns occur over a temporally related period of time rather than across students' entire dialogue contribution. After applying CSPM to student dialogue sequences, we identify 56 dialogue patterns. Using Seq2Pat's pattern-to-feature generation, we obtain a binary feature vector for each student representing the presence/absence of each extracted dialogue pattern. These extracted features provide a structured representation of dialogue patterns, enabling more interpretable and effective downstream predictive modeling and causal analysis. Figure 1 illustrates how the Seq2Pat algorithm extracts cognitive behavioral patterns from action sequences derived from game trace logs.

Similarly, we apply CSPM to in-game action sequences that

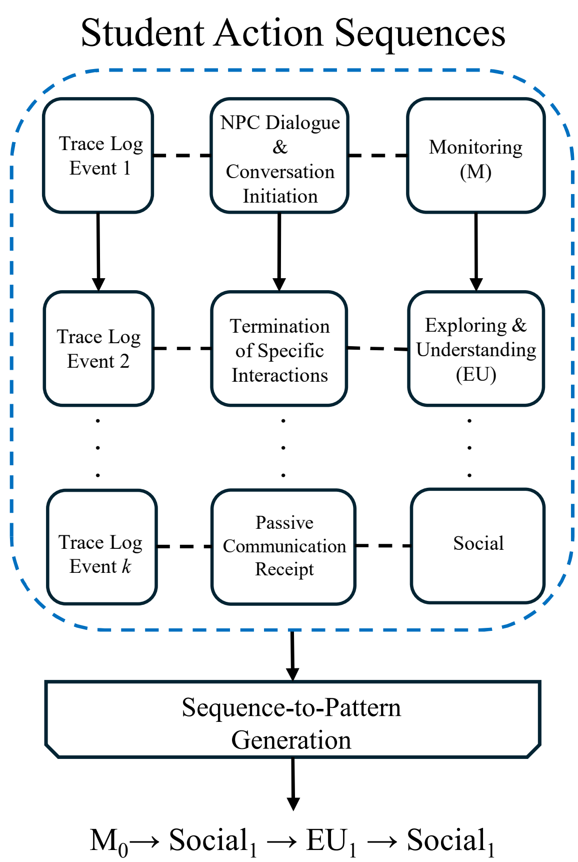

Figure 1: An illustrative example converting raw trace log information to CPS patterns. Seq2Pat's sequence-to-pattern generation extracts patterns from abstracted trace log sequences.

have been linked to the cognitive dimension of collaborative problem solving. Patterns extracted from game trace data can pinpoint knowledge exploration and problem-solving behaviors in addition to socially mediated and feedback-oriented patterns. These behavioral practices provide a lens into the cognitive behaviors predictive modeling finds informative for understanding student learning outcomes. Due to a large number of game trace log events (663 events on average) per student, Seq2Pat finds a large number of spurious and routine behaviors that do not necessarily reflect students' collaborative problem-solving behaviors. We mitigate this by introducing temporal and categorical restraints. Temporal constraints ensure that an extracted pattern spans no more than one minute, maintaining temporal consistency and minimizing overlap between disparate behavioral sequences. Categorical constraints are designed to minimize the number of cyclical actions that repeat more than three times. For example, categorical constraints will avoid the problem of frequently repeating patterns of the form "$Negotiating \rightarrow Negotiating \rightarrow Negotiating \rightarrow \ldots$". This also ensures that the algorithm is more likely to capture patterns that "break" from cyclical behaviors and move on to patterns with more action diversity. As a final measure, we filter out any patterns that occur more than one hundred times. Patterns with very large support counts often reflect routine or system-level behaviors and generally are not pertinent to the predictive modeling task. By refining the extracted patterns through these constraints, we enhance their relevance for predictive modeling, ensuring they capture meaningful collaborative problem-solving behaviors rather than routine or spurious actions.

4.3 Predictive Modeling

In order to assess the efficacy of the extracted multi-view behavioral patterns as indicators of student learning outcomes, we compare the performance of four machine learning models in their ability to disambiguate high- and low-performing students using dialogue and game trace log information. We apply ten-fold cross-validation with three-fold nested cross-validation for hyperparameter tuning. We evaluate their predictive performance using the averaged accuracy and macro-F1 scores from the outer cross-validation folds. Models were evaluated using extracted behavioral patterns from dialogue and game trace views separately, comparing single-view models against those utilizing both views (i.e., multi-view). Our work examines the predictive performance of Explainable Boosting Machines (EBM), Gaussian Process (GP) models, Tabular Prior-data Fitted Network (TabPFN), and a custom Multi-View neural network (MVNN).

Our model selection was motivated by the need to balance interpretability, adaptability, and flexibility to multi-view data while addressing some of the limitations of more traditional machine learning approaches. Explainable Boosting Machine (EBM) is a predictive modeling approach that provides high interpretability and performance while overcoming the limitations of Generalized Additive Model (GAM)-based approaches by providing iterative error correction abilities through boosting. By learning feature interactions in a data-driven manner, EBMs are adaptable to extracted behavioral patterns while maintaining transparency in feature contributions.

Due to our relatively small dataset size (75 students), Gaussian Process (GP) models are a non-parametric method that was chosen because of its ability to generalize well in low-to-moderate resource environments. Additionally, it can capture potential non-linear relationships while providing a principled way of quantifying uncertainty. TabPFN is a pretrained transformer based on the supervised classification of tabular data. It generally requires no hyperparameter tuning and is trained to approximate Bayesian inference on synthetic datasets. As a transformer-based model, TabPFN can also model feature interactions without any manual tuning. Finally, we construct a multi-view neural network that takes advantage of the attention mechanism to attend to one or more views simultaneously. It explicitly learns both shared and view-specific representations of extracted behavioral patterns while identifying complementary information highly predictive of student learning outcomes. Although neural network-based approaches lack the interpretability of other machine learning approaches, the use of integrated gradients or Shapley (SHAP) values provides a reasonable means to approximate feature contributions along the path from baseline input to model outputs. We also evaluate model performance without applying CSPM or performing multi-view learning. To this end, we utilize the UniLM-V6 BERT-based sentence transformer model to generate embeddings for dialogue and game trace data independently, serving as a naive baseline. We average each sequence of embeddings to generate a final embedding vector for each student that summarizes their trace and dialogue contributions. These representations are processed by a simple feed-forward neural network to identify high- and low-performing students. Along with a simple majority baseline, these naive embedding-based implementations serve as a baseline comparison for our chosen machine learning methods.

To systematically compare single- and multi-view modeling approaches, we generate multi-view representations using two distinct methods. For all models except the Multi-View Neural Network (MVNN), we apply Kernel Multiple Canonical Correlation Analysis (KMCCA) [1] to integrate dialogue and game trace data. This method performs dimensionality reduction using a Gaussian kernel with regularization, ensuring that the combined representations preserve high-order dependencies between views. In contrast, the MVNN directly processes raw input features, learning cross-view relationships automatically during training. Rather than relying on a precomputed joint representation, the MVNN projects each view into a shared feature space, where interactions are refined through a scaled dot-product attention mechanism. This formulation allows us to evaluate the MVNN's ability to integrate multiple views without imposing external constraints, ensuring a fair comparison between models that use engineered multi-view representations and the MVNN's learned feature integration.

4.4 Causal Discovery

To understand how extracted behavioral patterns relate to student learning outcomes, we employ a causal discovery method to reveal the underlying dependencies and potential causal structures within the data. Given that we do not already know the underlying causal structure of the data-generating process, we utilize the Peter-Clark (PC) algorithm, a constraint-based approach in causal discovery

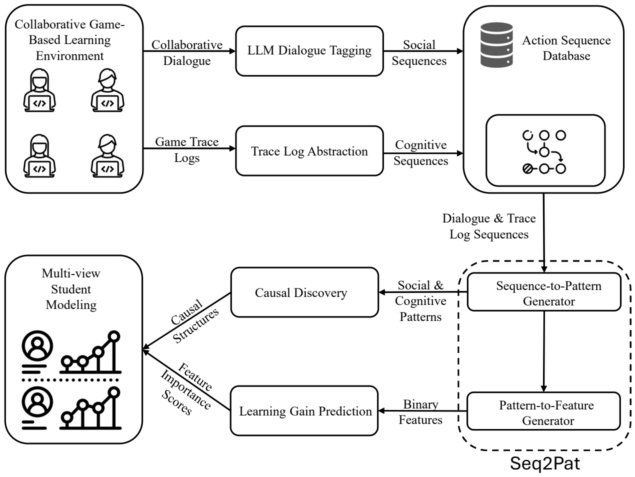

Figure 2: An overview of our multi-view student modeling framework.

[23]. The PC algorithm is a constraint-based causal discovery method that constructs a causal graph by leveraging conditional independence tests. It is well-suited for scenarios with high-dimensional data and relatively small sample sizes, as it does not require solving complex optimization problems. Additionally, PC is efficient for discrete data and is less sensitive to strong regularization constraints that may distort true causal relationships. We apply causal discovery to extracted cognitive and social behavioral patterns, identifying potential causal relationships between each aspect of students' collaborative behaviors and their learning outcomes. By analyzing conditional dependencies among extracted patterns, we can find sequential structures or co-occurring behaviors that may influence learning success, distinguishing between behaviors that are merely correlated and those that play a more direct role in shaping outcomes.

While the PC algorithm effectively reveals causal structures and the direction of dependence between behavioral patterns and learning outcomes, it does not provide information about the strength of these relationships. To estimate the magnitude of causal effects, we fit Generalized Additive Models (GAMs) to quantify how changes in one variable influence another along identified causal pathways. This approach enables us to assign edge weights to the causal graph, offering a more nuanced understanding of behavioral influences. Estimating edge weights is particularly advantageous in our research context, as it helps differentiate strong predictive relationships from weaker associations, allowing us to prioritize key behavioral patterns that significantly impact learning outcomes. By capturing nonlinear effects, this approach provides a more interpretable and data-driven foundation for understanding how different collaborative behaviors contribute to student success, and by integrating causal structure discovery with effect size estimation, we gain a more comprehensive understanding of how behavioral patterns contribute to student learning. An overview of our multi-view student modeling framework is outlined in Figure 2.

5. RESULTS

We now present our findings on predicting student learning outcomes using multi-view behavioral data. We first evaluate the effectiveness of predictive modeling approaches leveraging dialogue and game trace logs, followed by an analysis of causal discovery techniques to uncover underlying relationships between student behaviors and learning gains.

5.1 Predictive Modeling

We evaluate the performance of our predictive models, comparing single- and multi-view representations to assess their effectiveness in disambiguating high-performing (Class 1) and low-performing (Class 0) students. Table 3 shows the average results of our predictive modeling approaches using extracted dialogue patterns. The results for the dialogue view indicate that the Explainable Boosting Machine (EBM) outperforms the other models, achieving the highest accuracy (67.99%) and macro-F1 score (67.34%). EBM maintains strong performance across both classes, with balanced classification for Class 0 (68.07% macro-F1) and Class 1 (66.61% macro-F1). TabPFN and the Multi-View Neural Network (MVNN) show similar performance in accuracy and macro-F1, but MVNN has slightly better-balanced performance across both classes. The GP model, while achieving the highest Class 1 performance (68.83% macro-F1), struggles significantly with Class 0 (52.86% macro-F1), leading to the lowest macro-F1 score (60.84%). These results suggest that EBM is the most effective model for the dialogue view, providing both strong overall predictive power and class balance.

Table 3: Predictive modeling results for the dialogue view (social) of students' collaborative behaviors. Class 0 and Class 1 refer to the macro-F1 scores for each class.

Dialogue View				
	Accuracy	**Macro-F1**	**Class 0**	**Class 1**
TabPFN	64.00%	63.63%	62.20%	65.07%
EBM	**67.99%**	**67.34%**	**68.07%**	**66.61%**
GP	64.00%	60.84%	52.86%	68.83%
MVNN	65.33%	62.73%	63.41%	62.04%

For the game trace view (Table 4), both TabPFN and MVNN achieve the highest accuracy (69.33%), with MVNN slightly outperforming in macro-F1 score (68.72%vs. 68.26%). MVNN also demonstrates the best balance between Class 0 (67.06% macro-F1) and Class 1 (70.38% macro-F1) performance, suggesting it is the most robust model for this view. TabPFN performs similarly but has a lower Class 0 performance (63.35 macro-F1%). EBM follows closely with an accuracy of 66.66% and macro-F1 of 66.17%, maintaining a more balanced class performance compared to GP. The GP model, while strong in Class 1 prediction (73.51% macro-F1), struggles significantly with Class 0 (52.52% macro-F1), leading to the lowest macro-F1 score (63.02%). These results highlight MVNN as the best-performing model for the trace log view, with strong generalization across both classes.

Finally, for the combined view (Table 5), MVNN outperforms all other models by a significant margin, achieving the highest accuracy (72%) and macro-F1 score (71.62%). It also maintains the best balance between Class 0 (69.43%

Table 4: Predictive modeling results for the game trace view (cognitive) of students' collaborative behaviors. Class 0 and Class 1 refer to individual macro-F1 scores.

Game Trace View				
	Accuracy	**Macro-F1**	**Class 0**	**Class 1**
TabPFN	69.33%	68.26%	63.35%	73.17%
EBM	66.66%	66.17%	63.23%	69.11%
GP	66.60%	63.02%	52.52%	73.51%
MVNN	**69.33%**	**68.72%**	**67.06%**	**70.38%**

macro-F1) and Class 1 (73.81% macro-F1) performance, indicating strong generalization across both classes. In contrast, the other models exhibit notably lower performance. GP and EBM achieve identical accuracy (61.33%), but GP slightly edges out in macro-F1 (60.05% vs. 59.14%), with better Class 0 performance (57.59% macro-F1). EBM, while stronger in Class 1 (66.66% macro-F1), struggles with Class 0 (51.62% macro-F1). TabPFN performs the worst, with the lowest accuracy (60%) and macro-F1 (55.11%), showing a strong bias toward Class 1 (64.88% macro-F1) at the expense of Class 0 (45.83% macro-F1). These results suggest that MVNN is the most effective model when integrating multiple views, benefiting from its ability to learn feature associations across modalities.

Table 5: Predictive modeling results for combined dialogue and game trace views of students' collaborative behaviors. Class 0 and Class 1 refer to individual macro-F1 scores.

Combined View				
	Accuracy	**Macro-F1**	**Class 0**	**Class 1**
TabPFN	60.00%	55.11%	45.33%	64.88%
EBM	61.33%	59.14%	51.62%	66.66%
GP	61.33%	60.05%	57.59%	62.52%
MVNN	**72.00%**	**71.62%**	**69.43%**	**73.81%**

When compared to the raw input baselines (i.e., naive baselines), all of the proposed methods show a marked improvement in predictive performance (Table 6). For example, the dialogue view's raw input baseline has an accuracy of 49.33% and a macro-F1 of 39.64%. Yet, our predictive models, particularly the EBM, achieve much higher scores with balanced performance across classes. Similarly, the game trace data baseline's very low Class 1 performance (0% macro-F1) contrasts sharply with the robust results obtained by MVNN and other models, which deliver strong results for both classes. Finally, the combined view significantly outperforms the naive baseline (accuracy of 53.33%), with MVNN reaching an accuracy of 72% and a macro-F1 of 71.62%, underscoring the benefits of integrating multiple data views. When comparing the dialogue and game trace views, game trace views generally deliver robust predictive performance with higher overall accuracy and macro-F1 scores. While the dialogue view achieves competitive results, it tends to be slightly less consistent across both classes compared to the game trace view, where models like MVNN demonstrate a stronger balance between Class 0 and Class 1. This indicates that game trace data may offer a more stable signal for predicting student outcomes, highlighting its potential as a valuable input when used alongside dialogue data in a combined multi-view approach.

Table 6: Naive baseline implementations using either raw input features or a simple majority classifier. Class 0 and Class 1 refer to individual macro-F1 scores.

Raw Input				
	Accuracy	**Macro-F1**	**Class 0**	**Class 1**
Game Trace	46.66%	31.81%	63.63%	0.00%
Dialogue	49.33%	39.64%	59.29%	20.00%
Majority	53.33%	34.78%	0.00%	69.56%

5.2 Causal Discovery

Figures (3a) and (3b) show the discovered causal graphs for the dialogue and game trace logs, respectively. To reduce noise in the extracted causal graphs we apply thresholding to prune weak causal relationships. Thresholding edge weights in the causal graphs helps remove irrelevant or weak relationships, ensuring that only statistically significant and meaningful causal connections (black lines) are retained for clearer interpretation and more accurate modeling of the underlying causal structure. Additionally, generated structures often contain many disconnected sub-graphs, so as an additional post-processing step, we focus only on the largest connected sub-graph. The Peter-Clark (PC) algorithm is applied only to the input features and, therefore, has no means of showing the connection between the extracted behavioral patterns and predictive modeling of student learning outcomes. To bridge this gap, we utilize feature importance scores (red and blue lines) to highlight the impact of specific causal structures on predictive modeling decisions. In the figures shown, we use the feature importance scores from the MVNN model due to space limitations; however, our framework is model agnostic, meaning that the feature importance scores from any of the tested models can be substituted in place to maintain consistency and interpretability across different approaches. Feature importance scores, whether calculated directly or estimated through methods like Integrated Gradients or SHAP, provide a model-agnostic way to link discovered causal structures to predictive modeling decisions, as they offer a consistent framework for understanding which features influence the model's predictions, regardless of the underlying model or calculation method. This flexibility allows for a unified interpretation of feature relevance, bridging the gap between causal discovery and predictive analysis.

6. DISCUSSION

Our findings highlight the importance of integrating cognitive and social behavioral patterns to better understand collaborative learning processes. We will discuss how these findings address our research questions below.

6.1 Cognitive and Social Behavioral Patterns

The results demonstrate the effectiveness of constraint-based sequential pattern mining (CSPM) in extracting behavioral predictors of student learning outcomes (**RQ1**). To quantify the contribution of CSPM-derived features, we define it as the improvement in predictive performance compared to models trained on raw, unstructured features. This improvement is measured by comparing evaluation metrics (accuracy, macro F1-score) across the different feature sets. Even in the single-view setting, all evaluated models showed

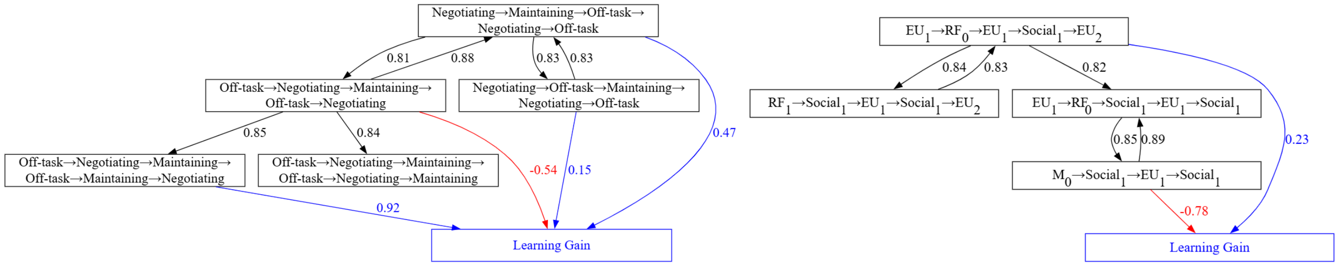

(a) Causal structure for the dialogue view (social).

(b) Causal structure for the game trace view (cognitive).

Figure 3: Causal graphs for student collaborative dialogue and collaborative problem-solving game trace logs. "EU" refers to "Exploring and Understanding", "RF" refers to "Representing & Formulating", and "M" refers to "Monitoring" behaviors. Subscripts denote the trace log type described in Table 1.

consistent performance gains over naive baselines, indicating that CSPM enables the discovery of structured behavioral patterns that contribute meaningfully to prediction. In the multi-view setting, performance improvements were even more pronounced, suggesting that integrating behavioral patterns from multiple views enhances the model's ability to capture relationships between cognitive and social CPS skills, which we further explore below (**RQ2**).

6.2 Effectiveness of Multi-View Learning

The effectiveness of multi-view learning is demonstrated by its ability to integrate diverse data sources, leading to improved predictive performance and a deeper understanding of the relationships between student behaviors and learning outcomes. As exemplified in Tables 3-6, we can see that overall, the MVNN model is the best-performing multi-view model (Accuracy 72%) when compared to all other approaches **(RQ2)**. Notably, we see a drastic drop in predictive performance for all other models when analyzing multiple views. Unlike the MVNN model, which learns both single-view and multi-view relationships during the training process, all other models require an intermediate step to combine views due to their inherent single-view nature. For this step, we utilize Kernel Multiple Canonical Correlation Analysis (KMCCA) to generate a combined representation of multiple views of data. By leveraging kernel methods, KMCCA can model non-linear relationships between views, improving the alignment and integration of diverse data sources for downstream prediction and analysis. The MVNN may outperform models that utilize KMCCA because it is able to learn complex, nonlinear relationships between different views directly from the data, offering greater flexibility in capturing interactions between features across views. Additionally, KMCCA may face challenges when combining binary matrices, as it assumes that the data within each view is linearly correlated (in kernel space), which may not hold true for binary representations where relationships are often non-linear. The inherent sparsity and lack of continuity in binary data can further complicate KMCCA's ability to find meaningful correlations, potentially limiting its effectiveness in handling such views. These limitations could be mitigated by using a neural network variant of CCA (DeepCCA). However, we argue that in such cases, employing the MVNN model simplifies the framework while enabling a data-driven, end-to-end learning approach.

The MVNN model's ability to integrate interaction and dialogue data highlights its advantage in real-world educational settings, where data often originates from heterogeneous sources (e.g., text and in-game actions). Unlike approaches that require intermediate steps for combining views, MVNN provides a direct, end-to-end framework that learns shared representations from raw inputs, simplifying the modeling process while effectively capturing complex multi-view relationships (**RQ2**). To evaluate its effectiveness, we compare the MVNN's predictive performance against baseline models that use early fusion or traditional multi-view integration methods, such as kernel multiple canonical correlation analysis (KMCCA). This comparison enables us to assess how well the multi-view approach captures complementary behavioral signals across modalities. Our results suggest that MVNN enables a more nuanced understanding of collaborative behaviors, supporting the development of adaptive learning environments that can dynamically integrate multiple data streams to deliver personalized, context-aware feedback and enhance student learning outcomes.

6.3 Causal Discovery in Multi-View Modeling

Figures (3a) and (3b) present the discovered causal graphs for dialogue and game trace log information. These structures represent causal relationships between each of the extracted behavioral patterns by the black lines connecting nodes. The PC algorithm constructs causal relationships in the form of partially directed acyclic graphs where some edges may be undirected, which we represent as two directed edges. To link the generated causal structures with predictive modeling decisions, we utilize feature importance coefficients for arcs between nodes within the causal graph and our model's predictions of student learning gains. Specifically, feature attribution can be performed in a model-agnostic way, providing a means to explore the relationship between model decisions and causal relationships. Our study focuses on feature importance scores from the MVNN model and their alignment with the extracted causal graphs; however, any model could be used, with only the specific node relationships to learning gains varying.

6.3.1 Dialogue Causal Graph

Examining the causal graph of dialogue data reveals a few key insights related to the MVNN's predictive modeling decisions. First, the model only found extracted dialogue patterns (black nodes) that only included "Maintaining Communication", "Negotiating Ideas", and "Off-task" behaviors to be particularly useful, while ignoring "Sharing Ideas" and "Regulating Activities". The exclusion of these features as strong predictors of learning gains implies that there may not be a strong direct relationship with learning gains in patterns that include these features. Further, this highlights that maintaining engagement, negotiating understanding, and off-task behaviors play a more dominant role in distinguishing student outcomes. It should be noted that sharing and regulation behaviors may contribute to predicting learning outcomes; however, their effects may be diluted or utilized more effectively by alternative modeling methods. The patterns are predominantly negotiation cycles ("Negotiating $\rightarrow \ldots \rightarrow$ Negotiating"), with three starting and two ending in off-task behavior. Additionally, two nodes exhibit communication maintenance cycles ("Maintaining $\rightarrow \ldots \rightarrow$ Maintaining"), both overlapping with negotiation cycles. These relationships highlight the role of negotiation in collaboration, the persistence of off-task behaviors, and the supporting role of communication maintenance in collaborative activities. The prevalence of negotiation cycles implies that students frequently engage in refining and clarifying their ideas, which may reflect productive collaborative engagement. The fact that these cycles often begin or end with off-task behaviors presents a more nuanced view. Negotiation cycles that start or end with off-task behaviors raise questions about whether these behaviors are always productive or if they indicate students are struggling to stay on task.

Alternatively, off-task behaviors are not always considered harmful to the collaborative process and could indicate socio-emotional contributions meant to improve social rapport during group activities. The persistence of off-task behaviors appearing in many patterns may be an indication of students having difficulty maintaining engagement during open-ended discussions. These patterns may also be capturing moments of periodic disengagement before students return to the negotiation process. Finally, the fact that "Off-task $\rightarrow$ Negotiating $\rightarrow$ Maintaining $\rightarrow$ Off-task $\rightarrow$ Maintaining $\rightarrow$ Negotiating" has the highest positive feature importance suggests that cycles incorporating both negotiation and communication maintenance cycles may be beneficial for learning outcomes. This could indicate that students who engage in off-task behavior but ultimately reinforce group cohesion through maintaining communication are able to re-engage productively. Conversely, the lowest-scoring pattern, "Off-task $\rightarrow$ Negotiating $\rightarrow$ Maintaining $\rightarrow$ Off-task $\rightarrow$ Negotiating", lacks a communication maintenance cycle, suggesting that when students disengage without a structured return to maintaining communication, the collaboration may suffer. This highlights the potential importance of mechanisms that help students re-engage meaningfully after off-task moments, reinforcing the idea that off-task behavior alone is not necessarily detrimental but must be balanced with effective social and cognitive regulation **(RQ3)**.

The structure of the causal graph suggests a strong interconnectedness among negotiation cycles, particularly those that incorporate off-task behaviors. The presence of high-weighted edges between these patterns implies that they frequently co-occur or influence one another. Additionally, the bidirectional influences among the negotiation cycles indicate that these behaviors do not occur in isolation but are part of an evolving collaborative dynamic, where off-task engagement may sometimes serve as a transition point between productive interactions rather than merely being disruptive. The presence of intermediate nodes such as "Negotiating $\rightarrow$ Maintaining $\rightarrow$ Off-task $\rightarrow$ Negotiating $\rightarrow$ Off-task" that mediate between other patterns suggests that certain behaviors may act as transition points between productive and unproductive collaboration. Designing interventions that target these intermediary stages could help redirect students toward more effective collaboration sequences.

6.3.2 Game Trace Causal Graph

Causal graphs extracted from game trace data provide several insights into students' cognitive collaborative behaviors. Many of the extracted patterns prominently feature social interactions, indicating that communication plays a central role in the learning process. This suggests that students frequently discuss their exploration and problem-solving strategies with peers while engaging in the game. We additionally see exploration as a recurring element, appearing multiple times across different behavioral patterns and often co-occurring with representation and social behaviors. This suggests that students repeatedly engage with the game environment before and after discussing their actions or answering questions. Monitoring behaviors (M) appear in one of the key sequences; however, it is inversely associated with learning gains (-0.78). This could indicate that students who rely too much on passive feedback from NPC conversations may struggle to internalize key learning concepts if they are not supplemented by actively engaging in exploration or discussion. We also see some variation in the frequency and usage of "Representing and Formulating" behaviors (RF). Some sequences include RF early in the pattern, such as "$EU_1 \rightarrow RF_0 \rightarrow Social_1 \rightarrow EU_1 \rightarrow Social_1$", while others include RF later or even omit it. This suggests differences in students' approaches where some students may engage in early conceptualization before social interactions, while others rely more on peer discussion before formulating responses. This variation may impact how effectively students consolidate their understanding. Finally, the extracted patterns reveal that exploration is a key driver of engagement within the EcoJourneys learning environment. "Exploration and Understanding" behaviors (EU) appear in every major pattern, emphasizing that interaction with the game environment is central to learning. The presence of EU multiple times within single sequences suggests that students may iterate between exploration, receiving feedback, and modifying their approach rather than progressing linearly. This reinforces the importance of designing exploration-rich learning environments where students are encouraged to test hypotheses, seek additional information, and refine their understanding dynamically.

The structure of the game trace causal graph provides additional insights into the relationship between cognitive collaborative problem-solving behaviors. First, the bidirectional associations suggest that students engage in an iterative learning cycle where exploratory patterns often co-occur with monitoring behaviors. This further supports the notion that students often refine ideas through repeated interactions rather than progressing linearly. The discovered causal graph highlights two connections between cognitive behavioral patterns and student learning gains, with "$EU_1 \rightarrow RF_0 \rightarrow EU_1 \rightarrow Social_1 \rightarrow EU_2$" having a positive connection (0.23) to student learning outcomes. A strong negative connection (-0.78) from "$M_0 \rightarrow Social_1 \rightarrow EU_1 \rightarrow Social_1$" to learning gains implies that certain types of monitoring behaviors may fail to support effective learning. Students may be passively consuming feedback rather than actively integrating it into their problem-solving processes. The presence of both positive and negative influences on learning gain suggests that a balance between exploration, representing and formulating ideas, and communication is key. Too much of any one behavior, especially repetitive social engagement without deeper cognitive work, can lead to diminishing returns. These findings underscore the importance of fostering a balanced approach to collaborative problem solving, where students are encouraged to engage in productive cycles of exploration and reflection while avoiding unproductive behavioral loops.

7. LIMITATIONS

This work has two limitations. First, it only uses dialogue and game trace information from 75 students, limiting the generalizability of the current modeling approaches. Future work should evaluate our predictive modeling and causal discovery framework across a larger, diverse range of students. Second, we chose the Peter-Clark algorithm for causal discovery due to its well-established ability to infer causal relationships in scenarios with a limited sample size. We omit comparisons to alternative causal discovery methods; however, it would be important to evaluate how well these causal relationships hold under different algorithmic assumptions to improve the robustness of our derived insights.

8. CONCLUSION

Understanding how student behaviors influence learning outcomes can aid in the development of effective collaborative learning environments. However, this is challenging due to the complex interplay between problem-solving actions and dialogue, as well as the need for interpretable models. Multi-view learning and causal modeling provide potential avenues to address these challenges by integrating diverse behavioral data sources and finding important collaboration dynamics that drive learning outcomes. We have introduced a multi-view predictive student modeling framework that combines constraint-based sequential pattern mining with causal discovery to extract interpretable behavioral features from dialogue data and game-trace logs. By utilizing these methods, we identify meaningful cognitive and social behavioral patterns, improve predictive accuracy, and support the interpretability of learning analytics models.

The results of empirical studies indicate that extracted behavioral patterns from cognitive and social data representations serve as effective predictors of student learning gains, surpassing the predictive performance of a naive baseline and models using raw data alone **(RQ1)**. We also found that game trace log information contains more effective predictors of student learning gain than dialogue data. Moreover, we find that multi-view modeling of student behaviors can achieve improved performance over single-view approaches when using deep learning models to more effectively identify cross-view relationships **(RQ2)**. To better understand the relationship between extracted behavioral artifacts and student learning outcomes, we apply the Peter-Clark causal discovery algorithm to provide a model-agnostic method to find causal relationships from the data-generating process. Extracted causal graphs identify key cognitive and social behavioral patterns that affect predictive modeling decisions **(RQ3)**. The causal graph analysis reveals that maintaining communication, negotiating ideas, and managing off-task behaviors are key predictors of learning outcomes, with off-task behaviors potentially acting as transitional moments in productive collaboration. The findings highlight the importance of fostering learning environments that encourage dynamic engagement and peer interactions. Additionally, a balanced approach to collaborative problem solving, avoiding unproductive behavioral loops, is essential for maximizing learning outcomes.

There are several promising directions for future work. First, because the Peter-Clark algorithm may not fully capture the complex and dynamic relationships between collaborative behaviors and learning outcomes over time, extending the causal discovery framework to incorporate temporal changes in student behavior and contextual factors, such as time spent on task or emotional states, could provide a deeper understanding of how these relationships evolve. Second, while dialogue and game trace data offer valuable insights, integrating multimodal data, such as video data capturing facial expressions, could offer a more holistic view of student behavior and cognitive states. This would improve the granularity of behavioral feature extraction and enhance both predictive and causal modeling. Third, future research should explore methods to improve model generalization, such as using ensemble-based techniques or more effective multi-view representation learning, particularly in the case of multi-view neural networks (MVNN). Additionally, it will be important to investigate how multi-view learning can be used to uncover latent structures underlying collaborative learning processes, offering educational researchers a powerful lens for examining the interplay between cognitive and social dynamics. Finally, cross-view causal inference should also be explored to determine how behaviors in one view, such as gameplay patterns, influence or are influenced by behaviors in another view, such as collaborative dialogue, which can further improve causal discovery methods' capabilities to capture complex student collaborative problem-solving dynamics.

9. ACKNOWLEDGMENTS

This research was supported by funding from the National Science Foundation (NSF) under Grants DRL-2112635, DRL-1561655, DRL-1561486, IIS-1839966, and SES-1840120. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

10. REFERENCES

[1] S. Akaho. A kernel method for canonical correlation analysis. *arXiv preprint cs/0609071*, 2006.

[2] J. Andrews-Todd and C. M. Forsyth. Exploring social and cognitive dimensions of collaborative problem solving in an open online simulation-based task. *Computers in Human Behavior*, 104:105759, 2020.

[3] J. Andrews-Todd, J. Steinberg, M. Flor, and C. M. Forsyth. Exploring automated classification approaches to advance the assessment of collaborative problem solving skills. *Journal of Intelligence*, 10(3):39, 2022.

[4] Y. J. Kim, H. Acosta, W. Min, J. P. Rowe, B. W. Mott, S. Chaturvedi, and J. C. Lester. Dual process masking for dialogue act recognition. In Y. Al-Onaizan, M. Bansal, and Y. Chen, editors, *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15270–15283. ACL, 2024.

[5] R. S. Baker, J. Clarke-Midura, and J. Ocumpaugh. Towards general models of effective science inquiry in virtual performance assessments. *Journal of Computer Assisted Learning*, 32(3):267–280, 2016.

[6] A. F. Botelho, A. H. Closser, A. C. Sales, N. T. Heffernan, and K. P. Vanacore. Causal inference in educational data mining. In *Proceedings of the 17th International Conference on Educational Data Mining*, pages 1034–1036. International Educational Data Mining Society, 2024.

[7] M. Emara, R. Rajendran, G. Biswas, M. Okasha, and A. A. Elbanna. Do students' learning behaviors differ when they collaborate in open-ended learning environments? *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–19, 2018.

[8] C. Huanyou, T. Hu, and L. Wu. Computer-based assessment of collaborative problem solving skills: A systematic review of empirical research. *Educational Research Review*, 43:100591, 12 2023.

[9] M. Huber. An introduction to causal discovery. *Swiss Journal of Economics and Statistics*, 160(1):14, 2024.

[10] J. Kaddour, A. Lynch, Q. Liu, M. J. Kusner, and R. Silva. Causal machine learning: A survey and open problems. *arXiv preprint arXiv:2206.15475*, 2022.

[11] F. S. Kia, S. D. Teasley, M. Hatala, S. A. Karabenick, and M. Kay. How patterns of students' dashboard use are related to their achievement and self-regulatory engagement. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge (LAK)*, pages 340–349, 2020.

[12] J. S. Kinnebrew, K. M. Loretz, and G. Biswas. A contextualized, differential sequence mining method to derive students' learning behavior patterns. *Journal of Educational Data Mining*, 5(1):190–219, 2013.

[13] K. Kitto, B. Hicks, and S. Buckingham Shum. Using causal models to bridge the divide between big data and educational theory. *British Journal of Educational Technology*, 54(5):1095–1124, 2023.

[14] G. Kostopoulos, S. Karlos, and S. Kotsiantis. Multiview learning for early prognosis of academic performance: A case study. *IEEE Transactions on Learning Technologies*, 12(2):212–224, 2019.

[15] L. Liu, J. Hao, A. A. von Davier, P. Kyllonen, and J.-D. Zapata-Rivera. A tough nut to crack: Measuring collaborative problem solving. In *Handbook of Research on Technology Tools for Real-World Skill Development*, pages 344–359. IGI Global, 2016.

[16] R. K. Lu, J. W. Liu, S. M. Lian, and X. Zuo. Multiview representation learning in multi-task scene. *Neural Computing and Applications*, 32:10403–10422, 2020.

[17] A. R. Nogueira, A. Pugnana, S. Ruggieri, D. Pedreschi, and J. Gama. Methods and tools for causal discovery and causal inference. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(2):e1449, 2022.

[18] L. Paquette and N. Bosch. The invisible breadcrumbs of digital learning: How learner actions inform us of their experience. In *Handbook of Research on Digital Learning*, pages 302–316. IGI Global, 2020.

[19] J. L. Plass, B. D. Homer, and C. K. Kinzer. Foundations of game-based learning. *Educational Psychologist*, 50(4):258–283, 2015.

[20] L. K. Poon, S. C. Kong, M. Y. Wong, and T. S. Yau. Mining sequential patterns of students' access on learning management system. In *Data Mining and Big Data: Second International Conference, DMBD 2017, Fukuoka, Japan, July 27–August 1, 2017, Proceedings 2*, pages 191–198. Springer, 2017.

[21] D. Rai, J. Beck, and I. Arroyo. Causal modeling to understand the relationship between student attitudes, affect and outcomes. In *Educational Data Mining 2013*, 2013.

[22] A. Saleh, C. E. Hmelo-Silver, K. D. Glazewski, B. Mott, Y. Chen, J. P. Rowe, and J. C. Lester. Collaborative inquiry play: A design case to frame integration of collaborative problem solving with story-centric games. *Information and Learning Sciences*, 120(9/10):547–566, 2019.

[23] B. I. Smith. Causal discovery and counterfactual explanations for personalized student learning. *arXiv preprint arXiv:2309.13066*, 2023.

[24] W. Song, W. Ye, and P. Fournier-Viger. Mining sequential patterns with flexible constraints from mooc data. *Applied Intelligence*, 52(14):16458–16474, 2022.

[25] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2001.

[26] K. Sridharan and S. M. Kakade. An information theoretic framework for multi-view learning. In *COLT*, number 114, pages 403–414, 2008.

[27] S. Sun. A survey of multi-view machine learning. *Neural Computing and Applications*, 23:2031–2038, 2013.

[28] Z. Swiecki, Z. Lian, A. Ruis, and D. W. Shaffer. Does order matter? investigating sequential and cotemporal models of collaboration. *A Wide Lens: Combining Embodied, Enactive, Extended, and Embedded Learning in Collaborative Settings*, 1:112–119, 2019.

[29] E. Tartaglia and P. Rankin. Applying causal inference to inform early-childhood policy from administrative data. *arXiv preprint arXiv:2301.06219*, 2023.

[30] B. Venkatachalam and K. Sivanraju. Enhanced student performance prediction using data augmentation with ensemble generative adversarial network. *Indian Journal of Science and Technology*, 17(44):4619–4632, 2024.

[31] S. Verstege, Y. Zhang, P. Wierenga, L. Paquette, and J. Diederen. Using sequential pattern mining to understand how students use guidance while doing scientific calculations. *Technology, Knowledge and Learning*, 29(2):897–920, 2024.

[32] M. Videnovik, T. Vold, L. Kiønig, A. Madevska Bogdanova, and V. Trajkovik. Game-based learning in computer science education: A scoping literature review. *International Journal of STEM Education*, 10(1):54, 2023.

[33] A. A. von Davier. Computational psychometrics in support of collaborative educational assessments. *Journal of Educational Measurement*, 54(1):3–11, 2017.

[34] S. Wadhwa and R. Dong. On the sample complexity of causal discovery and the value of domain expertise. *arXiv preprint arXiv:2102.03274*, 2021.

[35] H. Wang, H. Lu, J. Sun, and S. E. Safo. Interpretable deep learning methods for multi-view learning. *BMC Bioinformatics*, 25(1):69, 2024.

[36] X. Wang, A. Hosseininasab, P. Colunga, S. Kadıoğlu, and W.-J. van Hoeve. Seq2pat: Sequence-to-pattern generation for constraint-based sequential pattern mining. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 12665–12671, 2022.

[37] P. H. Winne. Construct and consequential validity for learning analytics based on trace data. *Computers in Human Behavior*, 112:106457, 2020.

[38] Z. Yu, Z. Dong, C. Yu, K. Yang, Z. Fan, and C. P. Chen. A review on multi-view learning. *Frontiers of Computer Science*, 19(7):197334, 2025.

[39] Y. Zhang and L. Paquette. Sequential pattern mining in educational data: The application context, potential, strengths, and limitations. In *Educational Data Science: Essentials, Approaches, and Tendencies: Proactive Education Based on Empirical Big Data Evidence*, pages 219–254. Springer, 2023.

[40] J. Zheng, W. Xing, and G. Zhu. Examining sequential patterns of self-and socially shared regulation of stem learning in a cscl environment. *Computers & Education*, 136:34–48, 2019.

[41] G. Zhu, W. Xing, and V. Popov. Uncovering the sequential patterns in transformative and non-transformative discourse during collaborative inquiry learning. *The Internet and Higher Education*, 41:51–61, 2019.