Developing Adaptive Team Coaching in GIFT: A Data-Driven Approach

Randall D. Spain[1], Jonathan P. Rowe[1], Benjamin Goldberg[2], Robert Pokorny[3], Michael Hoffman[4], Steven Harrison[4], and James Lester[1]

[1] *North Carolina State University, Raleigh, NC 27695, USA*
[2] *US Army DEVCOM Soldier Center, Orlando, FL 32826, USA*
[3] *Intelligent Automation Inc., Rockville, MD 20855, USA*
[4] *Dignitas Technologies, Orlando, FL 32817, USA*

Abstract

A critical step towards leveraging machine learning to create models for adaptive team-based coaching is collecting a large corpus of training data that can serve as a source for inducing data-driven coaching policies. In this paper we describe ongoing work using the Generalized Intelligent Framework for Tutoring (GIFT) to build reinforcement learning (RL) based coaching policies for promoting team performance in the domain of crew gunnery training. We describe our approach towards collecting a corpus of multimodal training data including video, communication, and simulation-trace data from U.S. Army gunnery crews who are completing simulation-based training exercises to prepare for crew gunnery qualification and sustainment. The dataset will be used to induce data-driven coaching policies for promoting individual and crew gunnery performance. We are using a component of GIFT called GIFT Data Collector to collect multimodal training data from simulation-based training stations as U.S. Army gunnery crews complete their assigned training exercises using the Virtual Battlespace 3 (VBS3) simulation platform. We discuss the data analysis pipeline that the team is developing to support data formatting, cleaning, filtering, and feature extraction processes that will be used to induce coaching policies. We also discuss how we plan to utilize multimodal training data, including crew communication logs, to iteratively refine the domain assessment model to support individual and team performance analysis and state classification. We conclude with a discussion of our upcoming research activities that aim to evaluate the acceptance and effectiveness of the coaching policies in an empirical study.

Keywords

Tutoring, Tutoring, Feedback, Reinforcement Learning

1. Introduction

Military training is designed to develop teams and units, ensure they can perform their required tasks and actions to specific standards, build cohesion, and provide opportunities to demonstrate team adaptability during novel combat situations [1]. Simulation-based training systems have long been used to achieve the benefits of live military training events. Simulation offers unique benefits compared to live training, including the ability to record and replay training exercises as well as the ability to automatically measure performance outcomes. Adaptive instructional systems (AISs) that integrate with simulation-based training environments can significantly improve training outcomes by presenting trainees with simulated scenarios that are high in psychological and physical fidelity, and by automatically assessing individual and team performance, diagnosing learners' strengths and weaknesses, and providing coaching and feedback.

Proceedings of the Challenges and Advances in Team Tutoring Workshop during the Artificial Intelligence in Education Conference, June 15th, 2021, Virtual from Utrecht, The Netherlands
EMAIL: rdspain@ncsu.edu (A. 1); jprowe@ncsu.edu (A. 2); benjamin.s.golberg.civ@army.mil (A. 3); bapokorny@gmail.com (A. 4); mhoffman@dignitastechnologies.com (A. 5); sharrison@dignitastechnologies.com (A. 6); lester@ncsu.edu (A. 7)

Research in the learning sciences and science of training consistently shows that feedback and coaching are critical for learning and skill development [2, 3]. Feedback can lead to increased motivation and reduce uncertainty by showing individuals how to correctly perform a task [4]. Specific forms of feedback such as corrective feedback can function to highlight errors, replace incorrect actions with appropriate responses, and reinforce the response so individuals will be more likely to act appropriately in the future. It can also encourage and improve trainee motivation [5]. Determining how feedback should be delivered, the content of the feedback message, and the level of cognitive processing required of the trainee is critical for optimizing learning and promoting engagement. These decisions become more complex when feedback and coaching are intended to improve team performance in AISs.

There is growing evidence that machine learning techniques, such as reinforcement learning (RL), provide an effective data-driven approach for modeling pedagogical coaching and feedback decisions in AISs [6, 7, 8, 9, 10]. RL has shown promise for automatically inducing tutorial policies that optimize student learning outcomes without requiring pedagogical policies to be manually programmed or demonstrated by expert tutors [8]. To illustrate the usage of RL for team-based pedagogical modeling, consider the task of delivering feedback after a team member failed to provide critical backup behavior during a collective training exercise, which negatively impacted the team's performance. The pedagogical model could select among several possible instructional strategies: 1) provide feedback to the individual directly and immediately; 2) provide delayed feedback to the individual directly (e.g., after the exercise is completed); 3) provide immediate feedback to the individual but include the entire team as secondary observers; 4) halt the task and provide immediate feedback to the individual and team. Several factors could influence the choice of instructional strategy, including the estimated competency of the individual team member, the competency of the collective team, the recent feedback history to date, the severity of the error, and the prospective impact of the team member's decision. Over time, pedagogical strategy selection can impact the performance of the individual and team, and their respective learning outcomes and motivation, which points toward the opportunity to systematically improve feedback strategies to optimize team training outcomes.

A critical step towards leveraging RL for team-based coaching is collecting a large corpus of training data that can serve as a source for inducing RL-based coaching policies. In this paper, we describe ongoing work using the Generalized Intelligent Framework for Tutoring (GIFT), an open-source service-oriented framework for designing, developing, and evaluating AISs to build RL-based coaching policies for promoting team performance in the domain of crew gunnery training. Specifically, we discuss our data collection approach which utilizes a component of GIFT called GIFT Data Collector that permits the team to collect multimodal training data from simulation-based training stations as U.S. Army gunnery crews complete training exercises using the Virtual Battlespace 3 (VBS3) simulation platform. We discuss the data analysis pipeline that the team is developing to support data formatting, cleaning, filtering, and feature extraction processes that will be used to induce coaching policies. We will also discuss how we plan to utilize these data, as well as crew communication logs to iteratively refine the domain assessment model to support individual and team performance analysis and state classification. We conclude with a discussion of our upcoming research activities that aim to evaluate the acceptance and effectiveness of the coaching policies in an empirical study.

2. Data-Driven Team Feedback and Coaching in Synthetic Training Environments with GIFT

Devising computational models that provide instructional support effectively—determining when to present feedback, what type of feedback to deliver, and how it should be realized—is a critical challenge facing AIS developers. In recent years, GIFT has become a key exemplar of how these challenges can be addressed [11]. GIFT provides a suite of software tools and standards that can be used to create models of adaptive team feedback in adaptive instructional systems for teams. By providing a reusable framework for adaptive training, GIFT can support the design, development, and investigation of Artificial Intelligence (AI)-based feedback models that operate at both individual and team levels across a range of domains and simulation-based training environments.

Recent enhancements to GIFT afford course authors with new opportunities to support team tutoring. These include the ability to model team structures in GIFT's assessment architecture, the addition of new condition classes and scenario adaptations to support automated assessment of team behaviors and team performance in simulation-based training environments, and the ability to deliver adaptive feedback and coaching messages at the individual and team level. For example, GIFT's Course Creator enables AIS developers to customize the content and timing of feedback messages, as well as select the intended audience for the feedback. GIFT also offers the ability to present feedback statements directly in a simulation-based training environment or in a separate tutoring interface, referred to as GIFT's Tutor User Interface. These feedback statements can be delivered visually as instructional text, verbally as spoken directives, or they can be delivered by a pedagogical agent.

In addition to these refinements, there have been considerable enhancements to GIFT's instructor tools, referred to as the GIFT Game Master. This interface facilitates a "human in the loop" AIS interaction model for assessing performance and injecting scenario adaptations during collective simulation-based training events. GIFT Game Master enables observer controllers to visualize unit progress through a live map view; monitor completed tasks, active tasks, and upcoming tasks; provide manual performance assessments; and inject scenario adaptations to shape unfolding simulation-based training scenarios at run-time. Game Master also enables instructors to provide pre-configured feedback messages and coaching, or tailor their own message, and select to whom feedback and coaching is delivered (e.g., individuals, subset of team, entire team). Furthermore, Game Master offers instructors the opportunity to tag and place bookmarks at specific timepoints during, or after, a training event and dictate notes that can be used to guide coaching and after-action review (AAR) discussions.

The aforementioned enhancements to GIFT provide instructors with a myriad of ways to deliver adaptive feedback and coaching in simulation-based training. Additionally, GIFT Game Master's "human in the loop" interaction model supports expert annotation of data from collective simulation-based training events, which is highly beneficial for the creation of data-driven models of adaptive coaching. Leveraging RL to devise data-driven models for delivering feedback and coaching offers a promising approach for supporting adaptive team training in GIFT. In the following sections we describe our approach towards collecting a corpus of multimodal training data including video, communication, and simulation-trace data from U.S. Army gunnery crews who are completing simulation-based training exercises to prepare for crew gunnery qualification and sustainment to support the development of team coaching policies in GIFT.

2.1. Crew Gunnery Testbed

A critical step towards realizing the vision of automated team feedback and coaching in synthetic training environments is establishing a corpus of team interaction data that can be utilized to devise RL-based adaptive feedback policies. To support the collection of a corpus of multimodal data that can be used to design and develop adaptive team feedback models with GIFT, we are using crew gunnery training in VBS3 as a testbed. VBS3 is an immersive virtual training environment used by the U.S. Army to support individual and collective training. Specifically, we are using a set of VBS3 missions developed at the Warrior Skills Training Center (WSTC) in Fort Hood, Texas that emulate real-world courses that gunnery crews use for live fire training and qualification. The virtual scenarios offer crews the opportunity to practice and rehearse crew coordination activities prior to engaging crew qualification on a live training range. Specifically, the VBS3 crew gunnery exercises involve a progression of training tables in which gunnery crews practice skills and procedures for engaging moving and stationary targets from a mounted-weapons system utilizing the direct fire engagement process. A key component of the direct fire engagement process is the coordination of actions and behaviors among the vehicle commander, gunner, and driver. During an engagement, crew members coordinate actions and exchange information pertaining to potential threats in the environment. Once a threat has been identified, crews engage in a fire command sequence, which is a well-defined protocol for communicating information and actions to facilitate a coordinated response to a threat. Each engagement is assigned a point-based score that ranges between 0 and 100 points. The points are calculated using a rubric that incorporates a combination of the firing vehicle's position, target type, target movement, target range, and target neutralization time to determine the engagement score. Crews can receive point deductions for violating the actions outlined in the fire command protocol and for committing safety violations.

Upon conclusion of the training session, crews conduct an AAR to discuss areas for improvement. The AAR is typically led by a certified Vehicle Crew Evaluator (VCE) who uses their domain expertise to guide the discussion and offer coaching and feedback. Devising computational models of adaptive feedback that provide timely coaching and feedback during training could complement existing crew gunnery training workflows by enhancing the feedback that can be delivered during simulation-based crew gunnery practice even when VCEs are not necessarily available.

Crew gunnery training in VBS3 is an attractive testbed for developing initial models of adaptive team feedback for several reasons. First, the VBS missions developed at WSTC are used for training by actual Army units, which is important to their relevance and alignment with Army training needs. Second, the VBS3 platform is well integrated with GIFT, which lends itself to instrumenting crew gunnery training exercises and delivering adaptive feedback via GIFT interfaces. Third, there are well-defined assessment and collective training procedures for crew gunnery, which provide a wealth of resources for developing domain models that can be translated into the representations and standards utilized by GIFT. Fourth, communication and behavior in crew gunnery is highly structured and procedural, constraining the space of assessment and feedback, which is a useful property of an initial test case. By establishing an adaptive team feedback capability in the domain of crew gunnery, this work aims to serve as a launch-point for developing adaptive feedback functionalities in other domains, learner populations, and training environments. Finally, the scenarios are utilized by hundreds of gunnery crews for sustainment and rehearsal training purposes which provides the research team with a unique opportunity to gather a rich corpus of data from crews that vary widely in skill and experience.

2.2. GIFT Data Collector System

Collecting and analyzing crews' verbal and non-verbal behaviors is critical for diagnosing crew coordination breakdowns and prescribing feedback as they complete gunnery engagements. To facilitate the collection of these data sources, we are utilizing GIFT Data Collector which is a light-weight version of GIFT that is used to facilitate data collection from simulation-based training events. GIFT Data Collector provides a subset of GIFT functionality focused on instrumenting a simulation-based training exercise to collect multiple data streams. It can be thought of as a "stealth" data collector that is able to passively collect multimodal data from crews as they complete the VBS3 training scenario. Figure 1 shows how GIFT Data Collector is being used in the WSTC training environment to collect multimodal training interaction data.

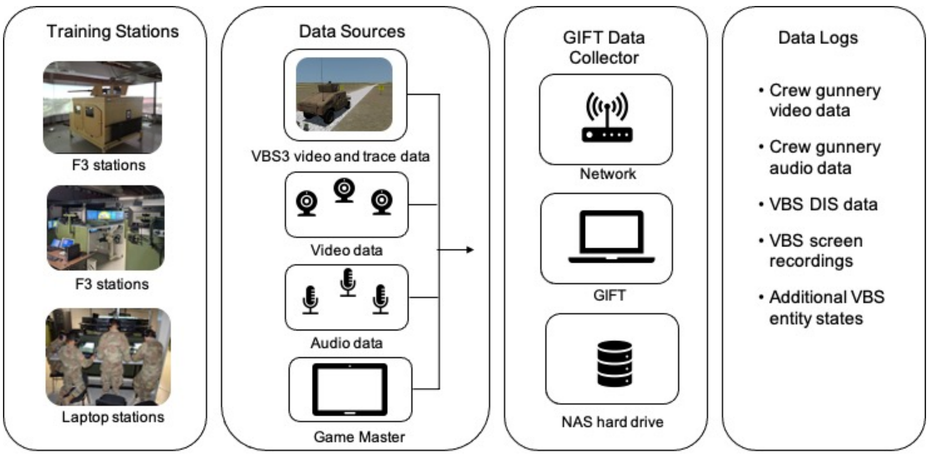

Figure 1: GIFT Data Collector setup at WSTC.

For our testbed, GIFT Data Collector collects multiple sources of data, including video, audio, and VBS3 simulation data. The training space at WSTC includes multiple form-fit-function (FFF) VBS3 training stations that provide crews with a physical mock-up of the vehicle system they are training on as well as VBS3 laptop training stations. Each FFF station can be instrumented with video and audio recording capabilities to capture crew communication and coordination behaviors. The purpose of capturing video and audio data is to facilitate analysis of trainee communication (verbal and non-verbal) during the training engagement. The GIFT Data Collector also records Distributed Interactive Simulation (DIS) data from VBS3, which represents a form of interaction trace log data for teams' simulation-based training interactions. The VBS3 DIS data consists of timestamped sequences of simulation states and actions related to specific entities in a VBS mission. A fourth data stream in our testbed includes instructor notes and bookmarks collected from the GIFT Game Master. We plan to provide instructors with the Game Master interface so they can place bookmarks, notes, and coaching recommendations during key training events that can be used to guide subsequent AAR sessions.

The goal of utilizing GIFT Data Collector is to establish a leave-behind data collection capability that can passively record data as units rotate through the training stations and perform crew gunnery training tasks. As training is being performed on the WSTC training stations, GIFT Data Collector silently collects data for each practice session. Once training is completed, GIFT Data Collection system can be removed from the training site to facilitate analysis of the captured data.

2.3. Data Analysis Pipeline

GIFT Data Collector consumes VBS3 DIS traffic and produce JSON-formatted log files containing timestamped records of the team's behavior in the VBS3 crew gunnery training scenario. A critical step towards using these data sources to devise data-driven models of team coaching and feedback is developing an automated (or semi-automated) software pipeline that consumes data from GIFT Data Collector and perform data cleaning and filtration, multimodal data alignment, extract features, and generate formatted output that is conducive to analysis by statistical software packages and machine learning toolkits. Our team is currently developing a series of Python scripts to facilitate this process. The formatted output will consist of vector-based feature representations, sampled coaching actions, and reward values for implementing reinforcement learning-based models of adaptive team coaching for crew gunnery training.

A critical step towards using the multimodal data sources collected from GIFT Data Collector is synchronizing the audio and video logs collected from crews to analyze crew communication and coordination behaviors. GIFT provides a robust set of tools for capturing and aligning learner performance data during simulation-based training on a single machine, but in a team context with multiple networked GIFT clients, effective workflows and tools for multimodal data capture and management become critical. We are currently working towards extending GIFT's capabilities to efficiently and accurately align and synchronize many different data streams from multiple GIFT clients at a high level of precision. Meeting this need is especially important for reducing the significant burden of manual data synchronization and alignment.

2.4. Refinements to the Crew Gunnery Domain Assessment Model

An important outcome of our initial data collection efforts at WSTC have included the ability to review the team communication data (video and audio data) and log data collected from GIFT Data Collector and refine our crew gunnery assessment model. The baseline Domain Knowledge File (DKF) implemented in GIFT models each gunnery engagement as a unique task in GIFT, with supporting subtasks aligning to the Observe, Detect, Engage, and Report phases of the direct engagement process. These subtasks are further classified into engagement specific activities that crew members must complete. The team developed this initial domain and assessment model by reviewing the crew gunnery training and qualification curriculum (TC 3-20.31), reviewing relevant field manuals, interviewing former VCEs and crew gunnery Subject Matter Experts (SMEs), and devising an abbreviated task analysis of the crew gunnery domain.

Using data and observations gathered from WSTC, the team has been iteratively refining our assessment model, using input guided by SMEs, to ensure key concepts and tasks are being accurately assessed and to identify strategies for introducing feedback and coaching. We plan to leverage notes and data provided by instructors through GIFT Game Master to gather coaching and feedback exemplars that can be used to refine coaching statements and assessment concepts utilized in GIFT. Decisions about what type of feedback to present, the sequencing and timing for presenting feedback, and when to interject coaching statements will be critical to model to promote trainee engagement and facilitate learning. By reviewing crew interaction data as well as coaching and feedback statements from instructors that are collected through Game Master, we will establish baseline parameters for providing adaptive team coaching to crews.

3. Conclusions and Future Directions

Developing computational models that automatically determine when and how feedback should be delivered to team members is critical for realizing the potential of AISs for team. We are investigating the creation of data-driven adaptive team feedback models in the domain of crew gunnery training using GIFT. Towards this goal, we are working with stakeholders at WSTC in Fort Hood, Texas to establish a testbed centered on crew gunnery training in VBS3. To facilitate the collection of multimodal training data to including video, audio, and simulation trace data, we are utilizing GIFT Data Collector which serves as a lightweight field-based data collection architecture. GIFT Data Collector supports the instrumentation of simulation-based training stations to passively collect and log interaction and multimodal training data. The recorded data can then be analyzed (offline) to examine trends and develop models using machine learning techniques. We plan to use the coordination and interaction data as well as coaching and feedback statements collected from instructors to devise an adaptive team feedback and coaching model in GIFT.

In the upcoming months, our team will be actively working towards collecting training interaction data and investigating the effectiveness of different forms of adaptive team-based feedback and coaching. As an initial starting point, we plan to investigate three types of feedback strategies for modeling and delivery by RL-based adaptive team feedback models: 1) Feedback timing, 2) Feedback content, and 3) Direct targets of feedback. Integrating data-driven models of team feedback and coaching with GIFT shows significant promise for enhancing the effectiveness of adaptive training tools for a range of military tasks.

4. References

[1] R.A. Sottilare, C.S. Burke, E. Salas, A.M. Sinatra, J.H. Johnston, S.B. Gilbert, Designing adaptive instruction for teams: A meta-analysis, International Journal of Artificial Intelligence in Education, 28(2) (2018): 225-264.

[2] V.J. Shute, Focus on Formative Feedback. Review of Educational Research, 78(1) (2008): 153–189.

[3] D.R. Sadler, Formative assessment and the design of instructional systems. *Instructional Science, 18* (1989): 119-144.

[4] D.R. Billings, Efficacy of adaptive feedback strategies in simulation-based training. Military Psychology, 24, (2012): 114–133.

[5] R. Badami, M. VaezMousavi, G. Wulf, M. Namazizadeh, Feedback after good versus poor trials affects intrinsic motivation. Research Quarterly for Exercise and Sport, 82(2) (2011): 360-364.

[6] M. Chi, K. VanLehn, D. Litman, P. Jordan, Empirically evaluating the application of reinforcement learning to the induction of effective and adaptive pedagogical strategies. User Modeling and User-Adapted Interaction, 21(1–2) (2011): 137–180.

[7] J. Rowe, J. Lester, Improving student problem solving in narrative-centered learning environments: A modular reinforcement learning framework. In Proceedings of the 17th International Conference on Artificial Intelligence in Education (2015): 419-428.

[8] S. Doroudi, V. Aleven, E. Brunskill, Where's the Reward? International Journal of Artificial Intelligence in Education, 29(4) (2019): 568–620.

[9] S. Shen, B. Mostafavi, T. Barnes, M. Chi, Exploring induced pedagogical strategies through a Markov decision process framework: Lessons learned. Journal of Educational Data Mining, 10(3) (2018): 27-68.

[10] F. Fahid, J. Rowe, R. Spain, B. Goldberg, B. Pokorny, J. Lester, Adaptively scaffolding cognitive engagement with batch constrained deep Q-networks. In Proceedings of the Twenty-Second International Conference on Artificial Intelligence in Education (2021): 113-124.

[11] R.A. Sottilare, K.W. Brawner, B.S. Goldberg, H.K. Holden, The generalized intelligent framework for tutoring (GIFT). In Fundamental issues in defense training and simulation (2017): 223-233.