

Considerations in the Design of Realistic Agents for Serious Games

Robert A. Sottolare¹, Christopher Ballinger¹ & Christopher McGroarty²

¹Soar Technology, Inc. Orlando, FL, ²US Army CCDC-STTC, Orlando, FL
bob.sottolare@soartech.com, cballinger@soartech.com, christopher.j.mcgroarty.civ@army.mil

Abstract

Artificially intelligent software agents are a staple of the video game and simulation industry, and are used for many purposes ranging from entertainment to training to analysis and decision support. Using data from live training exercises, this paper considers the process of constructing an effective data pipeline to determine the critical features that should be represented in intelligent agents for use in game-based training environments. Intelligent agents are often developed to represent the cognitive processes and behaviors of individuals or aggregate units or teams. The goal of the research described herein is to determine if the construction, training and validation of autonomous agents can be enhanced by modeling their perception, decision, and action processes after real entity behaviors captured during live training exercises.

Introduction

Serious games, computer games used for training or education, have been widely applied to military training. Two meta-analytic reviews of the literature confirm that instructional support in game-based learning environments show significant improvements to learning over non-game conditions. Clark, Tanner-Smith & Killingsworth, (2016) found improved learning ($g = 0.33$, 95% confidence interval $[0.19, 0.48]$, $k = 57$, $n = 209$). Visual and narrative features were factors in the persistence of learners or their engagement in the game. Wouters & Van Oostendorp (2013) found moderately improved learning ($d = .34$, $p < .001$). “Computer games can be considered complex learning environments in which players require instructional support to engage in cognitive processes such as selecting and actively organizing/integrating new information” (Wouters & Van Oostendorp 2013 p. 412).

To promote learning, serious games must enable trainees to interact meaningfully within the system. A lack of realism may cause trainees to discount the validity and training value of the game, and disengage from the training process (Greitzer, Kuchar & Huston 2007). The relationship between trainee engagement and learning is well documented (Rowe et al. 2010; Axelson & Flick 2010), as is the relationship between realism and engagement (Tashiro & Dunlap

2007; Checa & Bustillo 2020). Realism and fidelity are also related. Fidelity refers to how close a simulation or serious game is to reality, influences trainee engagement in three dimensions: physical fidelity, conceptual fidelity, and experiential fidelity (Choi et al. 2017).

Physical fidelity refers to the degree to which the simulation elements approximate visual, tactile, auditory and olfactory reality. Conceptual fidelity refers to the degree to which the game events flow compared to reality. Finally, experiential fidelity refers to the degree to which the game generates the responses (e.g., emotions or stress) that a trainee would expect in a similar real context.

In our study, we are examining the physical fidelity of intelligent agent behaviors and more abstractly their experiential fidelity of decision making processes during the execution of assigned missions/tasks. Intelligent agents perceive their environments (take in information), make decisions, and act autonomously to achieve assigned goals (Russell & Norvig, 1995). Agent performance may be improved though knowledge gained through new experiences. Learning agents often use rewards and penalties to drive decision making and increasingly better outcomes.

AI approaches are widely used in both educational (Woolf 1991) and training systems (McCalla et al. 1997) and have been for decades. Sottolare, Brawner & Ballinger (2021) documented a variety of AI forms including intelligent agents that are used in military simulation-based training applications and serious games. Intelligent agents augment training simulations and serious games to represent real entities as non-player characters (NPCs) or autonomous platforms. Brawner, Ballinger & Sottolare (2022) presented their research on intelligent agents applied in instructional contexts (education and training) and evaluated using a testbed methodology (Hanks, Pollack & Cohen 1993).

Having established the widespread effective use of games in instruction, the use of intelligent agents in serious games, the concept of learning agents, the importance of realistic intelligent agent behaviors and their influence on trainee engagement, and the need for engagement as a pre-requisite for learning. Two driving questions for our research are:

- *What features are transferrable from datasets captured during live training exercises with real people, and equipment on real terrain to agent models?*
- *Do the transferrable features of these datasets suggest that intelligent agents created from these datasets will exhibit more realistic behaviors and thereby be more engaging?*

To begin to answer these questions, next, we examine our approach to managing agent performance.

Realistic or Optimal Agent Performance?

Realistic and optimal intelligent agent behaviors are two different approaches used in the design of non-player characters (NPCs) in serious games. Realistic intelligent agent behaviors are designed to mimic human-like behavior in the game environment. This can include behaviors such as emotional responses, interpersonal interactions, and natural decision-making processes. The goal is to create a believable and engaging game environment that increases player immersion and involvement. On the other hand, optimal intelligent agent behaviors are designed to maximize performance or achieve a specific goal. These behaviors are often optimized for efficiency and can include decision-making processes that are not necessarily realistic or human-like. The goal is to achieve a desired outcome, such as winning the game or achieving a high score. Neither is wrong, but in view of the context of serious games, the level of realism is considered a more important technique for learning and motivation because it affects the learner's sense of immersion and engagement in the game.

A high level of realism can make the game more compelling and enjoyable, which in turn can increase the learner's motivation to play and learn. Additionally, a realistic game environment can help learners to transfer their learning to real-world operational situations, making it more applicable and relevant. When learners are more invested in the game and have a deeper understanding of the content, they are more likely to retain what they have learned and apply it in the future.

Furthermore, a realistic game environment can also support the development of various skills, such as problem-solving, critical thinking, and decision-making. Learners can practice these skills in a safe and controlled environment, allowing them to gain experience and confidence before applying them in real-life situations. Part of a realistic environment is the behaviors of intelligent agents as driven by their perception and action selection. It can be argued that realistic behaviors are preferred over optimal behaviors in serious games.

In games designed for entertainment, it is more prevalent to use optimal intelligent agent behaviors rather than realistic intelligent agent behaviors. This is because the primary goal of games for entertainment is to provide an enjoyable and entertaining experience for the player, rather than support educational or training objectives. Games designed for entertainment often prioritize challenge and gameplay mechanics over realism and believability. As a result, NPC behaviors are often designed to be optimal for achieving specific gameplay objectives, such as winning the game or achieving a high score, and persistence rather than being designed to be realistically human-like.

In contrast, serious games often prioritize realistic intelligent agent behaviors over optimal behaviors. This is because the goal of serious games is to provide a meaningful learning experience for the player, and realism is often seen as a key factor in engaging the trainee and achieving this goal. Since our focus is on training, this paper will address how live simulation datasets can contribute additional realism to intelligent agents used in serious games.

Cultivating Realistic Agent Performance

In this section, we examine the processes and steps required to extract realistic intelligent agent performance from live simulation exercise datasets. Understanding the problem and defining the goals of the intelligent agent to be developed influences all subsequent processes.

This involves providing context for the agent's deployment and in our case the goal is to optimize realism by modeling agent perception and action based on real entities and their behaviors as recorded in live simulation datasets. To facilitate the examination of the available live simulation exercise datasets, we designed a data management pipeline (Figure 1) that enables data preparation, model building and model evaluation processes.

Data Preparation Process

The data preparation process consists of four primary steps: data collection, data cleaning, data enrichment, and feature engineering. Step 1 is data collection which in our case was performed by the National Training Center-Instrumentation System (NTC-IS), an Army digital training capability.

Step 2 is data cleaning, a process of subtraction, used to determine if data in any given live simulation exercise dataset is correct or valid and if it should be included in the final prepared dataset. Step 3 is data enrichment, a process of addition that is used to fill in missing values by either by statistical mean or most probable value techniques.

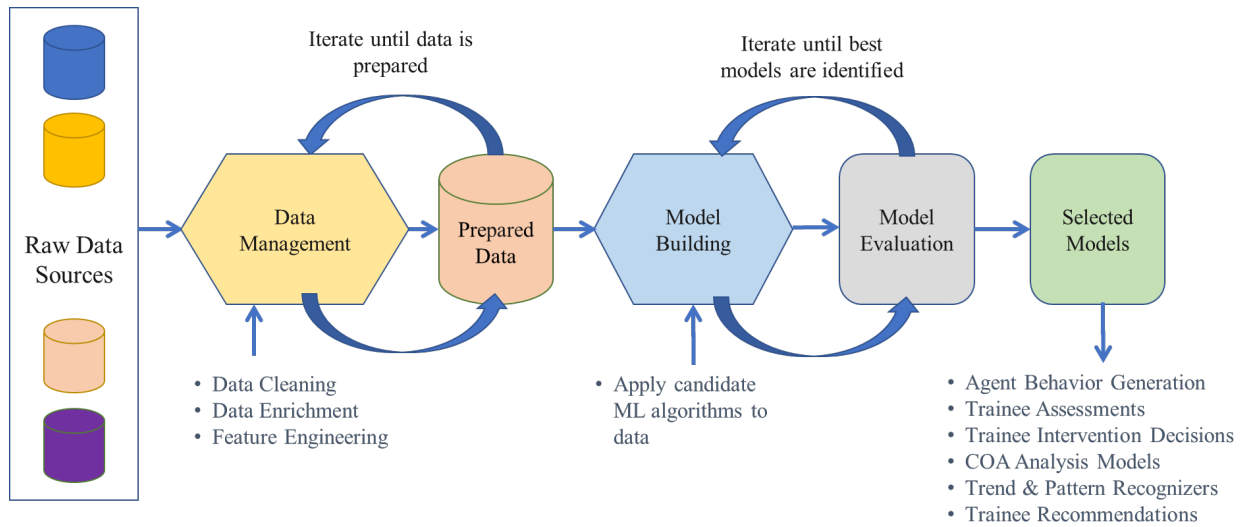


Figure 1. Data management pipeline.

Finally, in Step 4, feature engineering is used to identify critical variables and attributes. Through exploratory data analysis (EDA), we used a clustering algorithm to examine patterns, relationships, and trends to identify features that are relevant and influence individual and aggregate unit performance. The features identified will be used to train the agent. In Figures 2 and 3, we extracted the position and time of attacks origins and impacts of munitions. The Density-Based Spatial Clustering of Applications with Noise (DBSCAN) clustering algorithm was used to identify if the location and timing of the attacks formed identifiable clusters. Figure 2 shows that the source location of the attacks (left) and the target location (right) of the attacks formed strong clusters. Figure 3 also shows positional data points, but categorizes the data points by time using colors to show when the attack occurred.

Figure 2 graphs show that the attack origins and destinations tend to be in close proximity to each other, which could be an indication that the position of the clusters are close combat skirmishes between multiple factions, rather than a single unit using long-range attacks.




Figure 3. Result of EDA – clustering by position and time

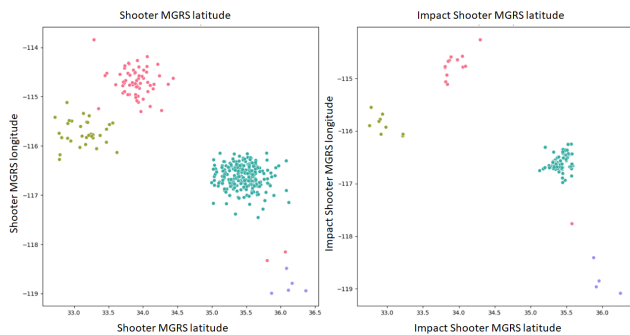


Figure 2. Result of EDA – clustering by position

MGRS stands for "Military Grid Reference System," which is a standardized coordinate system used by military forces around the world to identify locations on the Earth's surface.

Figure 3 graphs show that for the most part, the timing of the attacks is distributed across the clusters, which indicates multiple skirmishes going on in parallel. There may be other combat implications we can learn from these two figures when taken together instead of independently. However, finding subtle combat implications may be difficult or time consuming for an agent to learn directly from the time series data, but alternatively, we can simplify this data to a degree that is useful to the agent. For example, we could reduce the time series to the grouped attack timings that includes a count of how many attacks originated, impacted, hit/missed in each group during that time slice, and other factors. By using a high level of aggregation on this data, we generalize the world state in a way that an agent can more easily ingest

and provides critical information that influences how the agent should respond.

Model Building Process

The model building process takes the resulting prepared dataset and the features identified in the data preparation process and uses them to construct and train candidate models. First, ML algorithms that are most appropriate for the problem and data are chosen. Candidate ML algorithms are available through the Waikato Environment for Knowledge Analysis (WEKA), an open source suite of ML algorithms licensed under the GNU General Public License, and the companion software to Data Mining: Practical Machine Learning Tools and Techniques (Witten et al. 2005). WEKA contains a collection of visualization tools and algorithms for data analysis and predictive modeling together with graphical user interfaces. In Step 2, training the model, the analyst uses the selected features and algorithms to train models on the prepared dataset. Supervised ML models are trained with labeled data sets, which allow the models to learn and grow more accurate over time.

Model Evaluation Process

Model evaluation involves testing the model's predictive performance to determine its accuracy, precision, recall, and F1 score (the harmonic mean of precision and recall). Fine-tuning involves adjusting hyperparameters in the model, to improve its performance. The goal of hyperparameter tuning is to find the combination of hyperparameters that results in the best performance of the model on a specific task. The process of hyperparameter tuning typically involves the steps outlined below.

Step 1 is to define the hyperparameters to be tuned. This involves identifying the parameters that have a significant impact on the model's performance and selecting them for tuning. The learning rate is a hyperparameter in ML models that determines the step size at which the model updates its parameters during training. The learning rate controls how quickly or slowly the model learns from the training data. Regularization strength is a hyperparameter in ML models that controls the degree of regularization applied to the model's parameters. Regularization is a technique used to prevent overfitting, which is a common problem in ML where the model fits the training data too closely and has poor generalization performance on new, unseen data.

In Step 2, the analyst specifies the search space. The search space defines the range of possible values for each hyperparameter. The search space should be large enough to allow for a wide range of values but not too large as to make the tuning process computationally infeasible.

In Step 3, the analyst selects a performance metric which is used to value the model's performance on the validation

set after each iteration. Common metrics include accuracy, precision, recall, F1 score, and area under the curve (AUC).

In Step 4, the analyst chooses a hyperparameter tuning method including grid search, random search, and Bayesian optimization. Grid search involves exhaustively trying out all combinations of hyperparameters within the specified search space, while random search and Bayesian optimization use probabilistic methods to search the space.

In Step 5, once the optimal combination of hyperparameters has been identified, the analyst uses these to train the model using the selected performance metric. This process may be repeated several times, adjusting the hyperparameters after each iteration, until the performance of the model reaches its maximum. Based on the performance metric, the analyst identifies the model with the best performance features to translate into agent behaviors. While we are attempting to optimize performance, our goal in this study is to optimize performance within a range reflected by the live simulation dataset to develop an agent with realistic behaviors.

Finally, in Step 6, testing and deployment, involves the examination of the agent in a controlled environment and then deploying it into its target environment (e.g., serious game). Continuous monitoring is required to improve agent behavior and performance. Next, we discuss the behaviors that we are attempting to model in our intelligent agent.

Features, Cognitive Functions & Performance

The primary design goal for our intelligent agent is to model realistic intelligent agent behaviors based on live simulation datasets to promote enhanced trainee engagement in our serious game. Some of the features identified in the live simulation exercise dataset are relevant inputs to cognitive functions (e.g., decision-making) that influence the behaviors of real entities (soldiers, vehicles, air vehicles) and are also relevant for creating agents with realistic military behaviors:

- **Geolocation data** includes information about the location and movement of military assets, as well as information about the terrain and environment
- **Weapon systems data** includes weapon type, weapon ranges, weapon fire and munition detonation events, and their effectiveness in different contexts
- **Tactical information** includes military units, formations, maneuvers, and communications
- **Resource allocation information** includes troops, vehicles, and supplies that constrain system performance
- **Command and control information** includes structure and flow of information within the military units, and decision-making processes
- **Human factors information** includes the psychological and physiological factors affecting military personnel, such as stress, fatigue, and motivation

- **Environmental information** includes factors affecting performance, such as weather, light, and noise

Additional factors that influence agent cognition in a serious game environment include:

- **Representation of the environment** in the simulation includes its geometry, physical properties, and objects
- **Sensory data** includes the quality, reliability, and accuracy information that affect agent perception
- **Cognitive models** are used by agents to process sensory information and make decisions that influence judgment
- **Prior knowledge and experience** of the intelligent agent can also affect its perception
- **Attention** is the ability to focus awareness
- Rewards structure influences the decision-making of the intelligent agent
- **Constraints** are conditions that must be satisfied

During our explorations, several cognitive functions were identified as critical to developing realistic behaviors during military operations that result in realistic individual and unit performance. The agent should be able to **perceive and process information** from the environment to form an accurate understanding of the situation and make informed decisions. Agent perception during operations in a serious game is influenced by several factors previously defined:

- **Representation of the environment in the simulation:** An inaccurate representation of the environment can lead to incorrect inferences and those inferences can lead to suboptimal decisions by the agent.
- **Sensory data:** If the sensory data is noisy or incomplete, the agent may have difficulty correctly perceiving the environment.
- **Cognitive models:** An agent with a more sophisticated cognitive model may be able to make more accurate inferences about the environment than an agent with a simpler model.
- **Prior knowledge:** An agent with more prior knowledge about a particular environment may be better able to make accurate inferences about it than an agent with limited prior knowledge.
- **Attention:** The attentional focus of the intelligent agent can also impact its perception. An agent that is focused on a particular task or objective may be less aware of its surroundings than an agent with a broader attentional focus.

The agent should be able to make decisions based on the information it perceives using decision trees or reinforcement learning algorithms, to determine the best courses of action. Agent decision-making during operations in a serious game may be influenced by several factors (Wouters & Van Oostendorp 2017):

- **Representation of the environment:** accurate representation of the environment is needed for correct decisions.
- **Perception:** If the agent perceives the environment inaccurately, it may make incorrect decisions

- **Cognitive models:** an agent with a more sophisticated cognitive model may be able to make more accurate decisions than an agent with a simpler model
- **Prior knowledge:** An agent with more knowledge about a particular environment may make better decisions
- **Rewards:** An agent that is rewarded for taking risks may make different decisions than one that is not rewarded
- **Constraints:** Resource limitations or time constraints, can affect agent decision-making.

The agent should be able to **plan and execute actions** to achieve its objectives which involves generating a sequence of actions, considering resources and constraints, and understanding the influence of environment representation, and prior knowledge. A tangible metric for planning is performance in terms of the mission objectives achieved. The agent should be able to **communicate and coordinate** with other agents and entities including exchanging information (e.g., orders, status reports, and sharing situational awareness. Reliable communications can be influenced by the status of communication system nodes or terrain masking.

The agent should be able to **adapt to changing circumstances** and dynamically adjust its actions and strategies. This enables the agent to remain effective in dynamic and uncertain environments. The reward function is a key influencer of adaptivity. The reward function determines the objective that the intelligent agent is trying to achieve, and the agent's actions are guided by the rewards that it receives for its actions. Several factors determine the how the intelligent agent will be rewarded for adapting to changing conditions:

- **Alignment with mission objectives:** The reward function should be aligned with the mission objectives to achieve those objectives. If the mission objective is to complete the mission as quickly and efficiently as possible, then the reward function should be designed achieve that mission.
- **Encouragement of Desired Behaviors:** The reward function can be used to encourage desired behaviors in the intelligent agent. If the desired behavior is to minimize collateral damage, then the reward function should be designed to penalize the agent damage to unassigned targets.
- **Incentivizing exploration:** The reward function can be used to incentivize exploration. If the simulation involves an unknown environment, then the reward function should be designed to incentivize the agent for exploring new areas and gathering information.
- **Balancing tradeoffs:** The reward function can be used to balance trade-offs between different objectives. If the mission involves completing the mission quickly and minimizing collateral damage, then the reward function should be designed to balance these objectives.

Other factors that influence agent adaptivity include:

- **Feedback from the environment:** The feedback from the environment can impact the ability of intelligent agents to adapt during operations. Changes in weather, or lighting conditions can impact the mission.

- **Feedback from other agents:** Feedback from other agents, such as friendly forces or enemy forces, can impact intelligent agent adaptivity. Changes in the behavior of other agents can impact the success of the mission.
- **Feedback from performance metrics:** Feedback from performance metrics, such as accuracy, speed, and efficiency, can also impact intelligent agent adaptivity. If an intelligent agent is not meeting performance goals, it may need to adapt its plan to improve its performance.
- **New information:** Updates to the mission objectives or changes to the operating environment, can impact intelligent agent adaptivity. If new information indicates that the mission objectives have changed, agents may need to adapt their plan to achieve the updated objectives.
- **User input:** User input, such as guidance from a human operator, can also impact intelligent agent adaptivity. If a human operator provides new information or guidance, intelligent agents may need to adapt their plan to incorporate this input.
- **Artificial intelligence algorithms:** The artificial intelligence algorithms used to control the behavior of intelligent agents can also impact their adaptivity. The algorithms used must be capable of processing new information and adapting the plan of action in real-time.

The agent should have a high level of **situational awareness (SA)**, allowing it to understand the current state of the environment and predict future events. SA can be influenced by attention, constraints (e.g., willingness to explore) or other perception limitations. The agent should be able to **learn from experience** (the results of previous actions) and improve its performance over time.

Evaluation Results & Discussion

The current datasets limited the number and type of features that might be transferred from live training data into intelligent agent behaviors. However, we did identify features with a high rate of transfer that can be used to create behavioral or cognitive models. Indirect fire habits were sufficient to create a set of tactical heuristics to guide agent behaviors and decision-making processes. For tactical decision making where thought processes are not observable, we examined conditions and outcomes from fire events. For example, the analysis of outcomes for indirect fire engagements were sufficient to model the conditions under which the attacks were launched (e.g., munition type, firing position, and impact position) and the results which were either friend-on-foe (good outcome) or fratricides (poor outcome).

The features needed to define intelligent agents in a serious game as derived from live simulation datasets can be different based on the level of aggregation that the agent will represent. The choice of the level of aggregation will depend on the goals of the simulation and the intended use of the features derived from the live simulation dataset. The level of aggregation refers to the degree of detail at which the

entities are represented in the simulation. At a higher level of aggregation, a group of entities is represented as a single entity, and the focus is on the overall behavior and performance of the group as a whole. In this case, the features derived from the live simulation dataset may include aggregate metrics such as total number of units, total time taken to complete the mission or total resources expended. At the lowest level of aggregation, the entities are represented in greater detail, and the focus is on the behavior and performance of individuals. In this case, the features derived from the live simulation dataset may include detailed metrics such as individual unit performance, individual decision-making processes, or individual task proficiency.

The next logical step is to examine methods to enhance the explainability of the actions generated by the intelligent agent and its associated models. Explainability refers to the degree to which the behavior and decisions of an AI system can be understood, justified, and communicated to its users. In the context of AI, explainability is important because it helps to build trust and ensure that the system is used appropriately (Miller, 2019). Interpretable representations such as linguistic descriptions or visualizations produced as reports are being considered along with other approaches (e.g., transparent models like decision trees, internal agent state tracking, or human-centered design based on subject matter expert input). The goal of explainability design is to enhance the realism of intelligent agent performance in serious games, and thereby enhance user engagement.

Conclusion

We concluded that the ability to transfer features from live training datasets into usable models of entity behavior and cognition was our best metric of success and for a limited set of features, we demonstrated transfer. Since the live training data contained entity limitations such as situational awareness, performance outcomes, and errors, transferrable features from these datasets suggest that intelligent agents created from these datasets will exhibit more realistic behaviors and thereby be more engaging, but much more comprehensive datasets are needed to demonstrate this assertion for a variety of features under a variety of conditions. Real battlespaces can include phenomena (smoke and weather effects) that are not always present in live training environments, and therefore, would not be captured in the data for transfer directly into agent models.

Acknowledgements: The research described in this paper is sponsored by the U.S. Army (contract W912CG21C0002). Statements and opinions expressed in this paper do not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

- Axelson, R. D., & Flick, A. (2010). Defining student engagement. *Change: The magazine of higher learning*, 43(1), 38-43.
- Brawner, Ballinger & Sottolare (2022, May). Evaluating the Effectiveness of Artificially Intelligent Agents. In Proceedings of the Florida AI Research Society Conference, May 2022. Jensen Beach, Florida.
- Checa, D., & Bustillo, A. (2020). A review of immersive virtual reality serious games to enhance learning and training. *Multimedia Tools and Applications*, 79, 5501-5527.
- Choi, W., Dyens, O., Chan, T., Schijven, M., Lajoie, S., Mancini, M. E., ... & Aggarwal, R. (2017). Engagement and learning in simulation: recommendations of the Simnovate Engaged Learning Domain Group.
- Clark, D. B., Tanner-Smith, E. E., & Killingsworth, S. S. (2016). Digital games, design, and learning: A systematic review and meta-analysis. *Review of educational research*, 86(1), 79-122.
- Greitzer, F. L., Kuchar, O. A., & Huston, K. (2007). Cognitive science implications for enhancing training effectiveness in a serious gaming context. *Journal on Educational Resources in Computing (JERIC)*, 7(3), 2-es.
- Hanks, S., Pollack, M. E., & Cohen, P. R. (1993). Benchmarks, test beds, controlled experimentation, and the design of agent architectures. *AI magazine*, 14(4), 17-17.
- McCalla GI, Greer JE, Kumar VS, Meagher P, Collins JA, Tkatch R, Parkinson B. A peer help system for work-place training. B. d. Boulay, & R. Mizoguchi (Eds.), *AI-ED*. 1997 Aug;97(8):183-90.
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267, 1-38.
- Rowe, J. P., Shores, L. R., Mott, B. W., & Lester, J. C. (2010). Integrating learning and engagement in narrative-centered learning environments. In *Intelligent Tutoring Systems: 10th International Conference, ITS 2010, Pittsburgh, PA, USA, June 14-18, 2010, Proceedings, Part II 10* (pp. 166-177). Springer Berlin Heidelberg.
- Russell, S., & Norvig, P. (1995). A modern, agent-oriented approach to introductory artificial intelligence. *Acm Sigart Bulletin*, 6(2), 24-26.
- Tashiro, J. S., & Dunlap, D. (2007, November). The impact of realism on learning engagement in educational games. In Proceedings of the 2007 conference on Future Play (pp. 113-120).
- Witten, I. H., Frank, E., Hall, M. A., Pal, C. J., & DATA, M. (2005, June). Practical machine learning tools and techniques. In *Data Mining (Vol. 2, No. 4)*.
- Woolf B. *AI in Education*. University of Massachusetts at Amherst, Department of Computer and Information Science; 1991.
- Wouters, P., & Van Oostendorp, H. (2013). A meta-analytic review of the role of instructional support in game-based learning. *Computers & Education*, 60(1), 412-425.
- Wouters, P., & Van Oostendorp, H. (2017). Overview of instructional techniques to facilitate learning and motivation of serious games (pp. 1-16). Springer International Publishing.