Introducing RIDE: Lowering the Barrier of Entry to Simulation and Training through the Rapid Integration & Development Environment

Arno Hartholt, Kyle McCullough, Ed Fast, Adam Reilly, Andrew Leeds, Sharon Mozgai, Volkan Ustun & Andrew S. Gordon
USC Institute for Creative Technologies
12015 Waterfront Dr.
Playa Vista, CA 90094
{hartholt, mccullough, fast, reilly, leeds, mozgai, ustun, gordon}
@ict.usc.edu

Keywords: simulation, training, rapid iteration, system integration, AI, ML

ABSTRACT: *This paper describes the design, development, and philosophy of the Rapid Integration & Development Environment (RIDE). RIDE is a simulation platform that unites many Department of Defense (DoD) and Army simulation efforts to provide an accelerated development foundation and prototyping sandbox that provides direct benefit to the U.S. Army's Synthetic Training Environment (STE) as well as the larger DoD and Army simulation communities. RIDE integrates a range of capabilities, including One World Terrain, Non-Player Character AI behaviors, xAPI logging, multiplayer networking, scenario creation, destructibility, machine learning approaches, and multi-platform support. The goal of RIDE is to create a simple, drag-and-drop development environment usable by people across all technical levels. RIDE leverages robust game engine technology while designed to be agnostic to any specific game or simulation engine. It provides decision makers with the tools needed to better define requirements and identify potential solutions in much less time and at much reduced costs. RIDE is available through Government Purpose Rights. We aim for RIDE to lower the barrier of entry to research and development efforts within the simulation community in order to reduce required time and effort for simulation and training prototyping. This paper provides an overview of our objective, overall approach, and next steps, in pursuit of these goals.*

1. Overview

Military simulations are complex systems, requiring models of allies, enemies, neutral agents, weapons and vehicles to work seamlessly together, in arbitrary environments, reactive to user input, and per established doctrine. For training in particular, these systems need to be real-time, multi-user, and effective in meeting learning objectives. Creating these systems is a challenging undertaking. It requires a multidisciplinary team that includes expertise in integrated systems design, computer science, artificial intelligence (AI), machine learning (ML), learning sciences, real-time graphics, and military doctrine. Few teams possess all these skills, let alone the time and resources to design and develop these systems from scratch.

To lower the barrier of entry into this field, we introduce the Rapid Integration & Development Environment (RIDE). RIDE is a research and development platform that offers a range of common simulation capabilities to facilitate rapid prototyping. It has been designed and developed from the ground up to 1) integrate many current and upcoming simulation-related capabilities into a single framework, 2) support decision making and scenario development for simulation and training, 3) target novel and advanced technologies, in particular within the field of AI and ML, 4) be in principle agnostic to any particular simulation or game engine, and 5) vastly reduce the required time to develop new capabilities. RIDE combines a range of features in a drag-and-drop development environment, including One World Terrain (OWT), AI and ML, battle drill performance assessors, scenario creation tools, Non-Player Character (NPC) behaviors, xAPI logging, multiplayer networking, and multi-platform support. RIDE leverages robust game engine technology while designed to be agnostic to any specific game or simulation engine. By facilitating rapid prototyping specifically aimed at the simulation domain, RIDE enables decision makers to better define requirements and identify potential solutions in much less time and at much reduced costs.

In this paper we provide an overview of RIDE, with an emphasis on its architecture and key capabilities. We discuss three use cases that leverage RIDE and end with a discussion on its current status and next steps.

2. Background

Existing Army training systems are typically either not interoperable or are combined in a federated system of systems. For example, Bradley simulators cannot synchronize with UH-60 simulators, which cannot natively interface to Stryker simulators, and so on. As a result, combined training is either impossible or difficult to execute, data cannot easily be reused or updated, fidelity and features differ per system, and cross-system tracking and analysis of overall training effectiveness is near impossible. Individual technical solutions and the vendors that provide them result in a fractured landscape in efficiency and effectiveness are difficult to assess in a unified and objective fashion.

In order to address this, the Army Modernization priorities have directed the need for a Synthetic Training Environment (STE) that is a holistic training system which leverages the latest in real-time simulation capabilities, allows domain experts to create and modify content, and allows Warfighters to train on demand at the Point of Need. While many of the individual capabilities that STE requires exist in one form or another, they lack a common environment in which to integrate them. RIDE provides that environment for research and development purposes.

RIDE is being developed at the University of Southern California's Institute for Creative Technologies (USC ICT). USC ICT is a University Affiliated Research Center (UARC), working in collaboration with the Army Research Laboratory (ARL). USC ICT has a rich history of performing research and creating systems that teach, train, assess and heal. For example, UrbanSim is a PC-based virtual training application for practicing the art of mission command in complex counterinsurgency and stability operations, consisting of a game-based practice environment, a web-based multimedia primer on doctrinal concepts of counterinsurgency and a suite of scenario authoring tools [1]. ELITE/INOTS is a virtual human and intelligent tutoring system for instruction, practice and assessment of interpersonal communication skills for Army and Navy leadership and counseling [2]. BRAVEMIND is an evidence-based Virtual Reality Exposure Therapy (VRET) system aimed at providing relief from post-traumatic stress, allowing clinicians to gradually immerse patients into virtual environments representative of their traumatic experiences in a controlled, stepwise fashion [3,4]. SimSensei is a basic research clinical decision support tools and interactive virtual agent-based healthcare dissemination system that is able to recognize and identify psychological distress from multimodal signals, including signs that may be indicative of post-traumatic stress [5]. It uses MultiSense, a framework for audio-visual sensing modules, to gather rich data of a user interacting with a virtual human [6]. Many of the virtual human capabilities are available for the R&D community through the Virtual Human Toolkit [7].

As a UARC, USC ICT is a neutral and trusted advisor in support of validating and informing Army program requirements and acquisitions. While RIDE originated to target the STE specifically, it has been designed and developed to support a wide range of research and development needs, and is used for non-Army and non-training purposes by a range of organizations, including the Office of Naval Research (ONR).

3. Architecture

The RIDE ecosystem consists of three main layers: the Engine Layer, the Middleware Layer, and the Project Layer (see Figure 1).

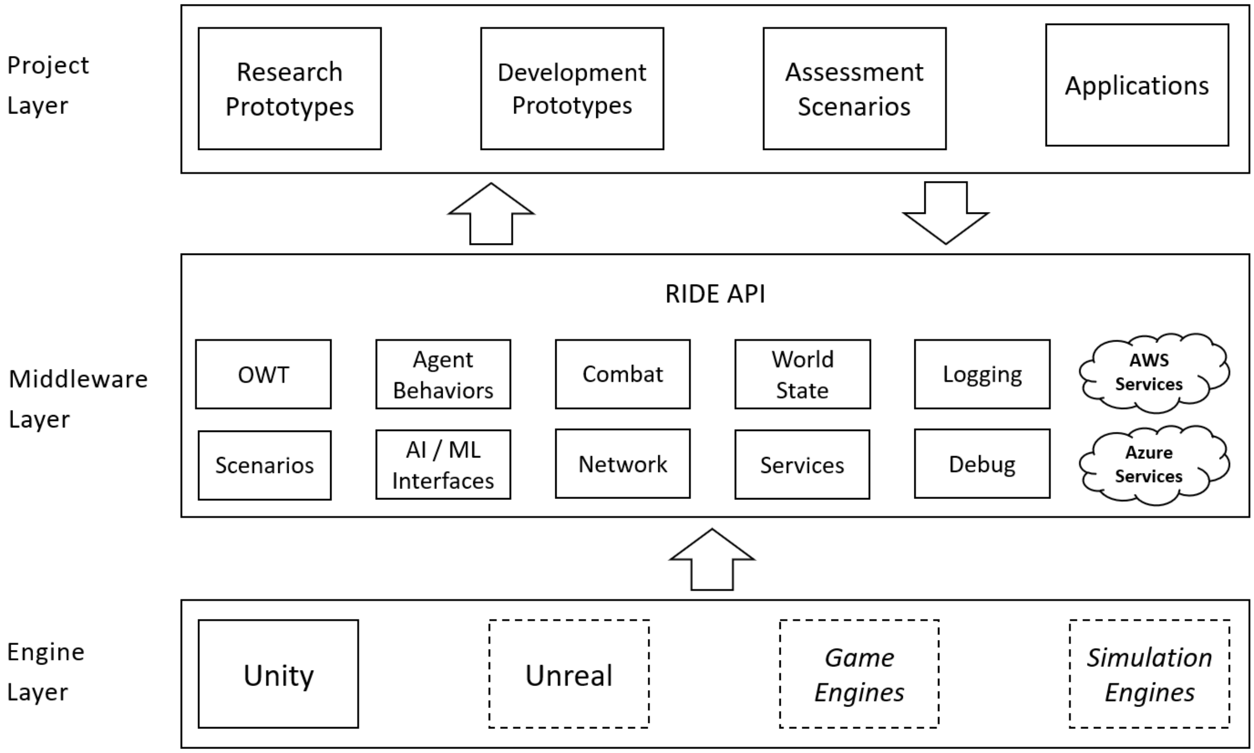

Figure 1: The RIDE architecture. RIDE leverages robust commercial engine capabilities while being agnostic to any specific game or simulation engine. It provides simulation-specific capabilities to enable a range of R&D projects.

The *Engine Layer* allows RIDE to leverage robust gaming technologies that provide common capabilities, including rendering, physics, animation, pathfinding, UI, audio, and network protocols. RIDE has been designed and developed to be agnostic to any specific game or simulation engine by abstracting core engine capabilities into its own API. This allows users of RIDE to develop simulation scenarios without the need to have experience with a specific engine. Currently, the main engine targeted is Unity, a very popular game engine for R&D purposes within the DoD. We are in the process of porting key RIDE functionality to Unreal.

The *Middleware Layer* abstracts and augments the Engine Layer with simulation-specific capabilities, including One World Terrain (OWT), AI agent behaviors, combat system, scenario system, machine learning (ML) interfaces, and networked multi-user capabilities. RIDE is designed to provide architectural flexibility in order to facilitate R&D in support of systems design, technical performance assessment, and scalability. It contains and is extendable with 3rd party assets and libraries, has native support for AWS and Azure cloud web services (e.g., storage, AI services), and can itself act as a web service for any of its capabilities. The RIDE API encapsulates these available capabilities in a well-designed suite of systems and services. The API follows interface-based design, which allows for principled implementations of new technologies; new concrete implementations simply need to implement the appropriate interface in order to be plugged into the RIDE and put to use. This not only provides the ability to extend RIDE with relatively little effort, it also enables multiple technical implementations of a single feature under unified interfaces, making it possible to contrast and compare, simultaneously or independently, separate approaches within a single platform, which is crucial for R&D purposes.

The *Project Layer* allows researchers and developers to leverage RIDE as a foundation for their own projects. RIDE provides common functionality through a drag-and-drop interface in combination with the API. This enables researchers and developers to rapidly create new scenarios as a starting point for their specific needs. The Project Layer acts as an incubation area, where new technologies and approaches can be explored safely, with mature results moving back to the Middleware Layer in order to advance RIDE and benefit all of its users. See Section 5 for several use cases.

4. Capabilities

4.1 One World Terrain (OWT)

Core to RIDE's foundational capabilities is high-resolution geo-specific terrain data. Typically the pathway to usage for high quality datasets requires a substantial effort covering sourcing the data, getting the coordinate systems right, optimizing the meshes for runtime, and building out custom terrain loaders in the runtime engine that integrate simulation necessities such as colliders and navigational meshes. For RIDE however, we've created a straightforward path that makes utilizing simulation ready real-world terrain meshes much simpler to the end user, while supplying all of the tools to access the inherent data for advanced simulation utilization. The integration of the pipeline and toolset within RIDE allows users to simulate, experiment, and leverage highly accurate real-world datasets for numerous use-cases with a drag-and-drop capability. Researchers at ICT have been working on the US Army's One World Terrain (OWT) project for a number of years, and have made significant contributions to the program's body of knowledge. OWT seeks to provide a set of 3D global terrain capabilities and services that can replicate the coverage and complexities of the operational environment and covers a five step pipeline: collection, processing, storage, distribution, and runtime. This data is currently used for a number of applications covering simulation, training, mission command, operations, and intelligence. RIDE supports full evaluation and research of all of the key areas of the 5-step process. Evaluating multiple sources of collected data, utilizing rich feature attribution and semantic classification from the processing phase, exploring optimal storage formats for the 3D data, distributing the data securely via the cloud, and then providing the runtime simulation environment. This aspect of RIDE is key not only to the OWT program, but a number of research objectives that are able to test their research holistically with RIDE acting as a one-stop shop for large-scale research initiatives.

A key area of ICT research for OWT is the production of highly accurate geo-spatial datasets with rich feature attribution and semantic classification processed through a fully automated pipeline [8]. These datasets are critical to a number of use cases and their interoperability means that RIDE's cross-application integrations are simpler to create and evaluate. The same terrain dataset that may be used in a mission planning toolkit, mission command system, or operational program of record can be leveraged directly in RIDE to supplement the external application, or simply provide an evaluation of the terrain data directly. An excellent example of this would be a situation where RIDE is being utilized for simulation, path planning & line of sight assessment and the exact same terrain dataset is being used simultaneously within a situational awareness application on a user's tactical device. This agnostic and interoperable foundation is core to the OWT program as well as RIDE.

RIDE currently supports OWT derivative datasets and researchers are working to fully integrate the OWT Well-Formed Format (WFF) as it's being developed by the STE's chosen vendor. Ensuring that RIDE supports the foundational data and ultimately the complete terrain packs as distributed will help to ensure that research does not need to be explored anew when the OWT WFF is finalized, and researchers are not delayed in meeting their goals to help inform requirements of various applied efforts.

4.2 Artificial Intelligence (AI) & Machine Learning (ML)

RIDE provides researchers and developers with a range of AI and ML related capabilities that can be leveraged to either directly create new entity behaviors or to explore new approaches to generate behaviors. This AI Behavior Framework provides a common interface for implementing behaviors in a simulation-agnostic manner. This framework allows developers to:

- Author behaviors
- Test / validate behaviors
- Train behaviors
- Execute behaviors
- Integrate externally created behaviors

Figure 2 provides an overview of the RIDE AI Architecture. It shows how the AI Behavior System interfaces with the core RIDE systems, the specific behavior types and how they are implemented via the generic interface provided within the AI Behavior System, and the tools used in conjunction with the behavior types. RIDE's AI Behavior Framework is initially written in Unity. However, it is designed in such a way that behaviors using this framework could be used in different simulation engines. This gives the ability to take advantage of certain Unity specific features, but not solely depend on them.

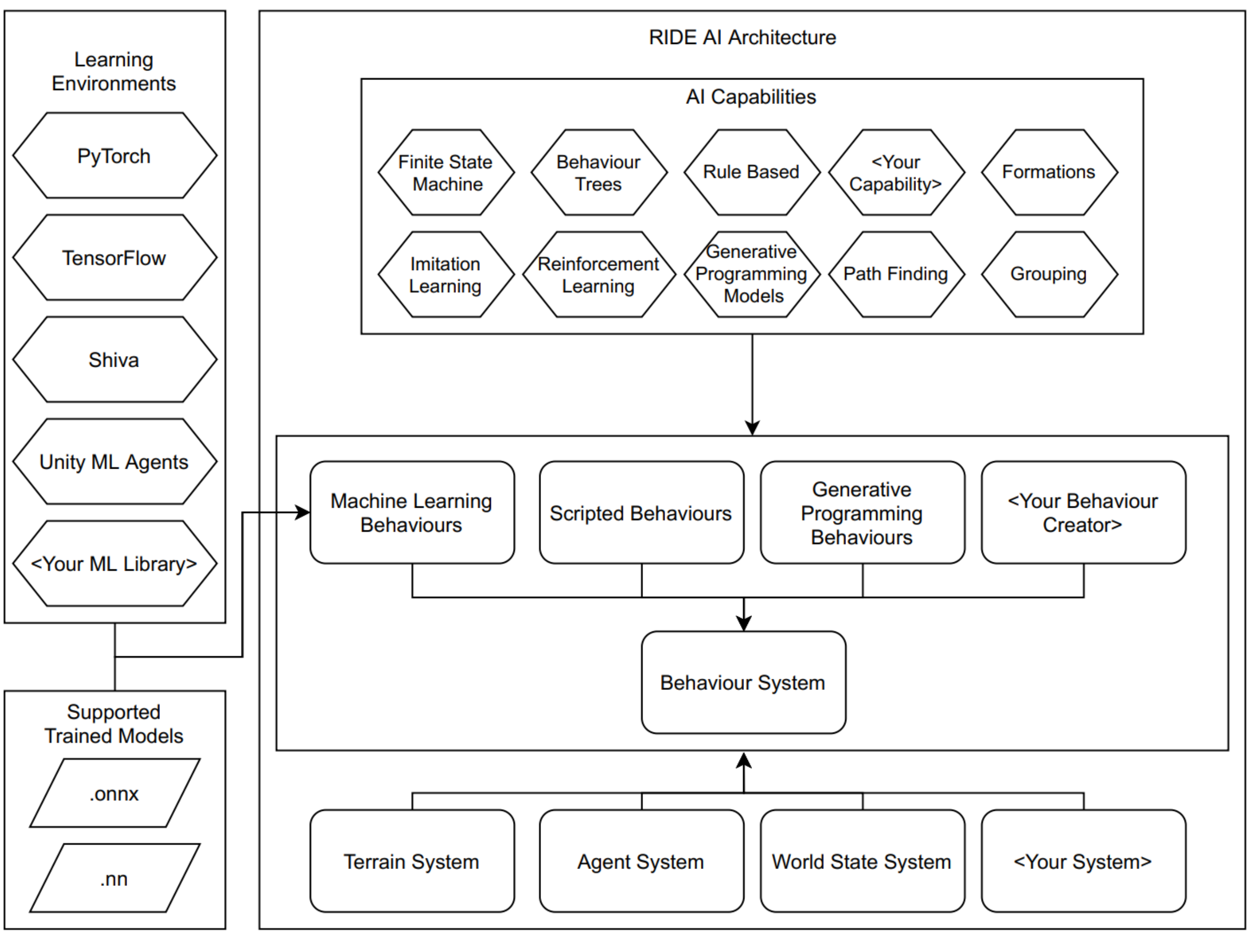

Figure 2. An overview of RIDE's AI Architecture, including its behavior system and interfaces.

A common interface provides the main agent behaviors, including movement, attack, grouping, and formations. These behaviors can either be scripted through traditional means or learned through ML. Scripted behaviors use either state machines or behavior trees, both native to RIDE. Learned behaviors are typically generated in the Project Layer (see Section 3), using TensorFlow, PyTorch, or custom systems. See Section 5 for example use cases. Inference models can be used directly in RIDE, either as .onnx or .nn models, or through Unity's ML-Agents.

4.3 Logging, Messaging, and Interoperability

The architectural flexibility discussed in Section 3 extends to the various ways in which RIDE can interface with existing and upcoming systems. RIDE supports the Distributed Interactive Simulation (DIS) messaging protocol used in many military simulations [9] as well as the xAPI data format that is core to many learning and training systems [10]. In addition, it supports the ActiveMQ messaging protocol [11], in particular through VHMsg, which allows communication with common virtual human and embodied conversational agent systems [12]. RIDE provides a common interface for data logging and storage. Multiple implementations allow for straightforward storage of data locally, in the cloud using AWS or Azure, or through custom means.

RESTful web services are a core pillar of the modern web, allowing vendors, producers, and consumers to interact in a flexible, secure, and real-time manner. The ability of RIDE to act as a web service provides interoperability between RIDE and other applications without the need for those applications to know about RIDE. This enables other applications (including those created years ago) to take advantage of RIDE's deep suite of capabilities with minimal work on their end. This makes RIDE highly interoperable with other systems, increasing its architectural flexibility, while maintaining efficiency and security. For example, any web browser can send and receive web requests without having to know the details of RIDE itself. It enables many simulation capabilities, especially useful for data and analysis without the need to host the RIDE platform or data locally.

4.4 Multiplatform Support

Modern training and simulation systems require flexibility in terms of delivery mechanisms in order to allow access by end users at the point of need. RIDE has been designed and developed from the ground up to account for this need through the support of multiple target platforms. This is made possible by the multi-platform support that current game engines offer. RIDE supports development on Windows, Mac, and Linus machines, and can target applications for mobile (Android, iOS), web (WebGL), AR/VR (e.g., SteamVR, OpenXR, HoloLens), and streaming solutions (e.g., Google Stadia, Microsoft xCloud, nVidia Geforce Now).

5. Use Cases

5.1 Battle Drills

It is critical to have useful tool sets for evaluation and testing of simulation capabilities, especially when there are a number of disparate research initiatives, use cases, and external collaborations involved in a field of research as great as that required of the STE. Recently, researchers at ICT have begun applying military doctrine, using dynamic simulations of the Army's infantry battle drills in an effort to provide not only a useful tool for predictive analysis, but to also provide a common integration point to evaluate the multiple research lines. This has been incredibly helpful at information gathering and informing the foundational capabilities. New research can be directly applied to the battle drills in order to provide immediate feedback on its usefulness and applicability to the STE, and the RIDE platform more generally. Of note, researchers have built a number of foundational capabilities for the "React to Ambush" battle drill. This drill features two blufor squads who enter the killzone of an enemy ambush. In response to the immediate ambush, the blufor units immediately deploy smoke and assault through the killzone, and units outside the killzone take defensive positions and apply suppressive fire. By creating this cohesive scenario, multiple lines of research can be applied to it, and the result can be directly measured. For instance, researchers at ICT are now looking at how to apply material attribution information from the OWT Data Model to cover and concealment behaviors so the autonomous agents within the battle drill simulation are able to make effective decisions as to what nearby objects and terrain geometry would offer sufficient cover. The agents are able to pull information from the data model, being imbued with knowledge that of the two cover positions ahead of them, one is concrete and one is plywood. Not only does this increase the realism and effectiveness of the battle drill simulation itself, but it directly informs the OWT Data Model by providing relevant information that can be used for optimization and evaluation. This creates a feedback loop in which updates to OWT research are then re-applied to the battle drill simulation and effectiveness can be measured against the previous baseline. This same drill can be used to apply additional research exploring dynamic real-time destructibility and deformation of meshes based on material classification [13], which is useful for visualization, user-feedback, and simulation fidelity.

Currently, RIDE and the battle drill are being explored as an integration point for a number of STE research objectives from external collaborators under the Army's Simulation & Training Technology Center (STTC) umbrella. This is already proving to be useful even as a way to help define our individual research lines and application use-cases. With RIDE offering a simple entry and jump-start into demonstrating integrated capabilities, it's much easier for even non-developer collaborators to begin to scope out their work as it will be applied to military training doctrines such as the infantry battle drills. As ICT researchers have been able to build a baseline for the necessary components, such as terrain, behavior trees, combat systems, unit states, and animations, the RIDE core already supports the necessary requirements. For instance, if an additional collaborator is working on a new Machine Learning method for road detection, and that information is passed to RIDE through the OWT Data Model, the AI units within a Battle Drill scenario will be able to take advantage of the new information for terrain traversal. If the road information is disabled, the agents lose that knowledge and the outcome of the Battle Drill could be significantly changed providing validation to external development and informing requirements across disparate research objectives.

Continuing to expand external collaboration capabilities, and internal foundational development efforts is key to the ongoing success of the platform and the quality of the simulation fidelity. As RIDE isn't a one-off development process for a singular prototype, capabilities can be continually layered. By focusing on doctrine oriented scenarios such as the Battle Drills, which have a relatively limited scope of function, this layering begins to create incredibly useful data and a toolset for continued extension. By expanding foundational capabilities, systems become stronger and more realistic. For instance, if a unit's health is calculated relative to damage from bullet hits, and those hits only take into account the hit itself, the simulation will still function but the analog to reality is hampered. If one research group is working to effectively model the damage from specific weapon and ammunition combinations at various ranges, and another group is working to create an effective system for damage based on the target location of a bullet hit, armored vs. bare or head vs. leg, for instance, these systems can be built independently in the same foundational code base. This lets both projects work at their own pace, demonstrating together or separately using built-in Battle Drill scenarios as a doctrine-based method of research validation.

5.2 Interpretation and Narrative Summarization of Simulated Battles

The RIDE platform serves as the testbed environment for ongoing applied AI research on automated interpretation and narrative summarization. In previous basic research funded by the Office of Naval Research (ONR), USC ICT pursued new technologies for real-time interpretation of behavior, combining ML methods for action perception with symbolic reasoning methods to search for assumptions about intentions and plans that best explain time-series observations [14]. To identify the best explanations, this work used probability-ordered logical abduction [15], which searches through the antecedents of definite clauses expressed in first-order logic to find the most probable set of assumptions that logically entail the observations. The resulting proof graphs can then be translated into English-language narrative summaries by combining template-based techniques with statistical ranking algorithms [16]. When pursued as a basic research effort, these technologies were explored in consideration of a single 90-second animated film crafted in 1944 by social psychologist Fritz Heider and his student Marianne Simmel, which subsequently became influential in basic research on social perception [17]. To advance these technologies toward more practical application, USC ICT researchers adopted RIDE as a simulation testbed for the real-time interpretation of simulated battles between virtual forces, funded by the US Army in support of the STE.

As testbed for applied research, the RIDE platform provided several key supporting capabilities. Chief among them is the ability to conduct networked multiplayer military exercises on realistic terrain, allowing researchers to collect test data from volunteers for the evaluation of the reasoning pipeline. Specifically, squad-sized groups of volunteers executed infantry squad maneuvers in an urban environment, providing a gold-standard dataset for the recognition and interpretation of these maneuvers. To accurately recognize squad maneuvers, ML models were trained using synthetic training data, i.e., where the annotated examples of squad maneuvers were generated automatically by assigning AI behaviors to groups of autonomous agents [18,19]. These AI behaviors were authored as scripted finite-state machines in RIDE, and used to collect unlimited examples of simulated squad maneuvers on varied terrain.

In exchange for the functionality provided by RIDE, several new capabilities were added to RIDE by the researchers pursuing this work, and made available for other RIDE users. These capabilities include a means of monitoring the movement of individuals and groups in and out of defined regions, a mechanism for representing all RIDE system events into first-order logical literals, and various mechanisms for triggering scenario events in the Master Scenario Events List (MSEL) of a training scenario. These technologies are in addition to the primary capabilities for real-time interpretation and narrative summarization, which have been transitioned in the distributed RIDE codebase for use in future research and development applications.

5.3 Generating Challenging Opponents for Military Training Simulations

Another ongoing applied AI research project at ICT, A Hybrid Architectural Approach to Adaptive Training (HA^{3}T), focuses on observation-based behavior model adaptation for synthetic characters in military training simulations. HA^{3}T leverages multi-agent reinforcement learning and synthetic entities' experience in combination with appropriate prior knowledge to generate adaptive and intelligent synthetic characters in the military domain.

Multi-agent Reinforcement Learning (MARL) models multiple agents that learn by dynamically interacting with an environment and each other, providing a framework for evaluating competitive and collaborative dynamics between these agents [20]. Current state-of-the-art blends advances in Artificial Neural Networks with algorithms from MARL research yielding promising models to generate adaptive opponent behaviors [21]. Consequently, MARL presents opportunities to train simulated enemies to become challenging opponents in military training simulations. However, these simulations take place in complex, continuous, stochastic, partially-observable, non-stationary, and doctrine based environments with multiple players, either collaborating or competing against each other. The combination of these challenges makes the computational generation of intelligent behavior a very challenging task.

Furthermore, these challenges require MARL models to interact with realistic representations of these environments to learn realistic behavior policies. MARL models are complex, and they need a lot of fine-tuning to yield robust behavioral models, which demand a significant amount of computation. Therefore, for timely and efficient training of behavior policies, MARL requires simulation environments capable of running very quickly (much faster than real-time) while still being realistic.

HA^{3}T project brings together Shiva, a novel and adaptable multi-agent reinforcement and imitation learning framework, with RIDE [22] to train challenging synthetic characters as opponents for military training simulations. RIDE offers the desired capabilities to create complex military training simulations on realistic terrain that seamlessly integrate with Machine Learning (ML) frameworks like Shiva through Unity's ML-Agents Toolkit [23]. Additionally, RIDE provides a fast simulation environment that supports the demanding computational requirements for MARL experiments. The in-house comparisons performed have shown that RIDE can run as fast as environments like Neural MMO [24], which is a massively multi-agent game engine for reinforcement learning agents. Utilizing RIDE, we have run several MARL experiments in which agents in a red team train to learn behavior policies against a blue team whose behavior is dictated by doctrine-inspired computational scripts. Our experiments yielded successful proof-of-concept behavior models for the red team, which is an essential step towards bringing ML into military training simulations leveraging RIDE as a simulation environment to support large scale MARL experiments.

6. Conclusion

We have provided an overview of RIDE, the Rapid Integration & Development Environment. RIDE's architecture leverages game engine technology, industry cloud services, ML frameworks, and common simulation standards to provide a solid foundation for simulation related research and development. Dedicated capabilities, including One World Terrain, character AI behaviors, logging, and multiplatform support, all encapsulated within a principled API and drag-and-drop interface, enable rapid prototyping for training, planning, operations, and analysis systems. We discussed three use cases that highlight how RIDE has facilitated research and development. While it is early days for RIDE, it has already shown its value to the military R&D community. Ongoing work includes improving scenario creation capabilities, extending the AI/ML frameworks, and enhancing combat simulation models. For more information, see https://ride.ict.usc.edu.

Acknowledgements

We would like to thank all of our partners and collaborators at USC ICT, STTC, ONR, STE CFT and others, who have been instrumental in making RIDE a success. Part of the efforts depicted were sponsored by the US Army Research Laboratory (ARL) under contract number W911NF-14-D-0005. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

[1] McAlinden, R., Durlach, P., Lane, H., Gordon, A., and Hart, J. (2008) "UrbanSim: A game-based instructional package for conducting counterinsurgency operations". Proceedings of the 26th Army Science Conference, Orlando, FL, December 1-4, 2008.

[2] Hays, M. J., Campbell, J. C., Trimmer, M. A., Poore, J. C., Webb, A. K., & King, T. K. (2012). "Can role-play with virtual humans teach interpersonal skills?". University of Southern California, Los Angeles, Institute for Creative Technologie.

[3] Rizzo, A., Hartholt, A., Grimani, M., Leeds, A., & Liewer, M. (2014). "Virtual reality exposure therapy for combat-related posttraumatic stress disorder. Computer", 47(7), 31-37.

[4] Mozgai, S., Hartholt, A., Leeds, A., & Rizzo, A. S. (2020, April). "Iterative Participatory Design for VRET Domain Transfer: From Combat Exposure to Military Sexual Trauma". In Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (pp. 1-8).

[5] DeVault, D., Artstein, R., Benn, G., Dey, T., Fast, E., Gainer, A., Georgila, K., Gratch, J., Hartholt, A., Lhommet, M. and Lucas, G., (2014, May). "SimSensei Kiosk: A virtual human interviewer for healthcare decision support". In Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems (pp. 1061-1068).

[6] Scherer, S., Marsella, S., Stratou, G., Xu, Y., Morbini, F., Egan, A. and Morency, L.P., 2012, September. "Perception markup language: Towards a standardized representation of perceived nonverbal behaviors". In International Conference on Intelligent Virtual Agents (pp. 455-463). Springer, Berlin, Heidelberg.

[7] Hartholt, Arno, David Traum, Stacy C. Marsella, Ari Shapiro, Giota Stratou, Anton Leuski, Louis-Philippe Morency, and Jonathan Gratch. (2013) "All together now: Introducing the Virtual Human Toolkit." In International Workshop on Intelligent Virtual Agents, pp. 368-381. Springer, Berlin, Heidelberg.

[8] Chen, M., Feng, A., McCullough, K., Prasad, P. B., McAlinden, R., Soibelman, L., and Enloe, M. (2019). “Fully Automated Photogrammetric Data Segmentation and Object Information Extraction Approach for Creating Simulation Terrain”. Interservice/Industry Training, Simulation and Education Conference, November 29-December 3, 2019, in Orlando, FL.

[9] Hofer, Ronald C., and Margaret L. Loper. "DIS today [Distributed interactive simulation]." Proceedings of the IEEE 83.8 (1995): 1124-1137.

[10] https://xapi.com, accessed January 15 2020.

[11] https://activemq.apache.org, accessed January 15 2020.

[12] Hartholt, Arno, Ed Fast, Adam Reilly, Wendy Whitcup, Matt Liewer, and Sharon Mozgai. "Multi-Platform Expansion of the Virtual Human Toolkit: Ubiquitous Conversational Agents." International Journal of Semantic Computing 14, no. 03 (2020): 315-332.

[13] Nam, N., McCullough, K., New, R., McAlinden, R.. (2019). “Methodology to Utilize Pre-Computed Voronoi Diagrams to Enable Dynamic Deformation and Destructibility of Environmental Meshes Within A Simulation Environment”. Interservice/Industry Training, Simulation and Education Conference, November 30-December 4, 2020, in Orlando, FL.

[14] Gordon, Andrew S. (2018) “Interpretation of the Heider-Simmel Film Using Incremental Etcetera Abduction”. Advances in Cognitive Systems 7 (2018), pp. 23-38.

[15] Gordon, A. (2016) “Commonsense Interpretation of Triangle Behavior”. In Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16), February 12-17, 2016, Phoenix, Arizona.

[16] Ahn, E., Morbini, F. and Gordon, A. (2016) “Improving Fluency in Narrative Text Generation With Grammatical Transformations and Probabilistic Parsing”. Proceedings of the 9th International Natural Language Generation Conference (INLG-2016), September 5-8, 2016, Edinburgh, Scotland.

[17] Heider, F., and Simmel, M. 1944. “An experimental study of apparent behavior”. The American Journal of Psychology 57(2):243–259.

[18] Feng, A., and Gordon, A. (2020) “A Framework for Action Detection in Virtual Training Simulations using Synthetic Training Data”. Interservice/Industry Training, Simulation and Education Conference, November 30-December 4, 2020, in Orlando, FL.

[19] Feng, A., and Gordon, A., (2020) “Recognizing Multiplayer Behaviors Using Synthetic Training Data”. 2020 IEEE Conference on Games, Online, August 24-27, 2020.

[20] Buşoniu, L., Babuška, R., & De Schutter, B. (2010). “Multi-agent reinforcement learning: An overview”. In Innovations in multi-agent systems and applications-1 (pp. 183-221). Springer, Berlin, Heidelberg.

[21] Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., & Mordatch, I. (2017). “Multi-agent actor-critic for mixed cooperative-competitive environments”. In Advances in Neural Information Processing Systems (pp. 6379-6390).

[22] Ustun, V., Kumar R., Reilly A., Sajjadi S., & Miller A. (2020). “Adaptive Synthetic Characters for Military Training Simulations”. In virtual Interservice/Industry Training, Simulation, and Education Conference.

[23] Juliani, A., Berges, V. P., Vckay, E., Gao, Y., Henry, H., Mattar, M., & Lange, D. (2018). “Unity: A general platform for intelligent agents”. arXiv preprint arXiv:1809.02627.

[24] Suarez, J., Du, Y., Isola, P., & Mordatch, I. (2019). “Neural MMO: A Massively Multiagent Game Environment for Training and Evaluating Intelligent Agents”. arXiv preprint arXiv:1903.00784.

Author Biographies

ARNO HARTHOLT is the Director of R&D Integration at USC ICT where he leads the central tech integration group. He is responsible for much of the technology, art, and processes related to integrated systems, with a particular focus on the interchange of research and industry capabilities. He has a leading role on a wide variety of research prototypes and applications, ranging from medical education to military training and treatment. Hartholt studied computer science at the University of Twente in the Netherlands where he got his Master’s degree. He worked at several IT companies, from large multinationals to early start-ups, before accepting a position at ICT in 2005. He has over a decade's worth of experience in leading multidisciplinary research and commercial projects, with an emphasis on virtual humans, virtual reality, augmented reality and serious games. He is an experienced speaker with appearances at SXSW, SIGGRAPH, XDC, LeapCon, IEEE VR, CNN, Unite, UCLA, and VR/AR Vision Summit.

KYLE MCCULLOUGH is the Director of Modeling & Simulation at USC-ICT. His research involves geospatial initiatives in support of the Army's One World Terrain project, as well as advanced prototype systems development. His work includes utilizing AI and 3D visualization to increase fidelity and realism in large-scale dynamic simulation environments, and automating typically human-in-the-loop processes for Geo-specific 3D terrain data generation. Kyle received awards from I/ITSEC and the Raindance festival, winning "Best Interactive Narrative VR Experience" in 2018. He has a B.F.A. from New York University.

ED FAST is a Computer Scientist at the University of Southern California Institute for Creative Technologies where he has been the Technical Lead on a number of projects for over a decade. While his primary projects have been virtual humans related, the main goal has been to build reusable systems for a variety of research prototypes and applications. He previously worked in the video game industry and has enjoyed rewarding success in applying that expertise to simulations. Ed holds a BS in Computer Science from California State Polytechnic University, Pomona.

ADAM REILLY is an ICT Computer Scientist developing the Rapid Integration and Development Environment (RIDE) software platform on the Training Simulation Software (TSS) team for the Synthetic Training Environment (STE). Since joining ICT in 2010, his work has focused on building large, scalable, multi-domain systems that facilitate research and production software, much of which focuses on the creation and authoring of realistic virtual agents across desktop, web, and VR/AR platforms. His current work on RIDE focuses on the development, integration, distribution, and scalability of research and production level capabilities, used across the DOD and research communities. Prior to working at ICT, Adam worked in the video games industry for several years. Apart from his career as a Computer Scientist, he has taught software development at the New York Film Academy since 2015.

ANDREW LEEDS is a technical project specialist at the USC Institute for Creative Technologies. His professional work focuses on medical and training application development and the integration of novel, multisensory VR systems. Leeds received a B.A. in Film and Video from Columbia College Chicago.

SHARON MOZGAI is the Associate Director of the Medical Virtual Reality (VR) lab at USC ICT where she heads up the research and development efforts for clinical applications of VR, including VR Exposure Therapy, in addition to Virtual Human training and education. Her background is in psychology with a Master's degree from Harvard University. She has worked at a range of organizations, focusing on computational linguistics at MIT, organizational behavior at Harvard Business School, cognitive behavioral therapy and motivational interviewing at the San Francisco VA Medical Center, and VR, Virtual Humans, and AI at USC ICT and at several tech startups.

VOLKAN USTUN is the Associate Director of the Cognitive Architecture Group at USC ICT. His general research interests are cognitive architectures, computational cognitive models, multi-agent systems, and simulation. He has a Ph.D. degree in Industrial and Systems Engineering and held postdoc positions at Georgia Tech and Rice University before joining ICT.

ANDREW S. GORDON is the Director of Interactive Narrative Technologies at USC ICT, and Research Associate Professor in the USC Department of Computer Science. He is the author of the 2004 book, *Strategy Representation: An Analysis of Planning Knowledge*, and the 2017 book, *A Formal Theory of Commonsense Psychology: How People Think People Think* (with Jerry R. Hobbs). He received his Ph.D. in 1999 from Northwestern University.