A Neuro-Symbolic Approach to AI Agents within Simulation

Samuel F. Griffith, Todd W. Griffith, Ph.D., Isaiah B. Mallery, Brady D. Kirkpatrick, Hudson Pepper
Discovery Machine®, Inc.
Williamsport, PA
tgriffith@discoverymachine.com, samuel@griffith.llc, imallery@discoverymachine.com,
bradydkirkpatrick@gmail.com, hap5183@psu.edu

ABSTRACT

Military training systems have reached an impasse. In any complex realistic training situation, trainees must communicate with other participants who are supporting the training. For example, a Joint Terminal Attack Controller (JTAC) talking to an A10 Pilot or to the Air Support Operations Center (ASOC). Typically, these other participants are human role-players, and often, they have outdated or limited expertise in the role they are playing. We show how a neuro-symbolic approach to Artificial Intelligence (AI) can support conversational agents that effectively replace the human role-player for specific missions. As machine learning (ML) attempts to move into highly specific domains, its voracious need for more data increases and the program's ability to explain itself wanes. To combat these problems, we developed an innovative technique for incorporating knowledge representation and machine learning via an augmented ensemble approach we call a Parallel Processing Neural Net (PPNN). We outline the technique in its current form and highlight areas for potential improvement and applications. Initial experiments involve training over commentary data from a popular strategy game called League of Legends. We then describe how our neuro-symbolic approach has leveraged TensorFlow and transformer technology to create a conversational A10 agent for Combat Air Support (CAS) missions. We also detail our exploration of quantum theory's applications to the field of AI, particularly in the area of projective simulation, which has benefited in recent years by utilization of quantum-like algorithms.

There are two problems that prevent agents from being widely used in training systems. The first is that agents within simulations are unable to communicate, and the second is that agents are unable to make predictions of the future based on current scenarios. For this reason, our work pertains to creating agents which can extrapolate future events from the environment and communicate accurately and efficiently with the user.

ABOUT THE AUTHORS

Samuel F. Griffith is a recent graduate of The Pennsylvania State University with a Bachelor of Science in Physics in the computational option. He has done several internships for Discovery Machine®, Inc. and has an interest in quantum information science. He will be attending graduate school in the Fall.

Todd W. Griffith, Ph.D. has been working in the area of intelligent systems research for 25 years and has published papers in the areas of cognitive science, HCI, and intelligent systems. He is frequently invited to present at workshops, conferences, and panels. Prior to founding Discovery Machine®, Inc., Dr. Griffith taught computer science at Georgia Tech and Bucknell University. Dr. Griffith has focused his research on building knowledge acquisition and deployment tools that enable SMEs to encode their own mental models and problem-solving strategies on the computer. He has obtained grant funding through DARPA, NASA, NAVAIR, ONR, ERDC, MDA, USAF, and NSF. Dr. Griffith holds four U.S. Patents, and the RESITE® Suite and Knowledge Service Engine are commercial realizations of his research.

Isaiah B. Mallery has recently graduated from The Pennsylvania State University with a Bachelors of Engineering in Computer Science and a minor in mathematics. He has begun working for Discovery Machine®, Inc., with hopes of focusing his future efforts on intelligent systems and virtual reality systems design.

Brady D. Kirkpatrick is a Graphics Software Engineer at Leidos. He received his Bachelor of Science degree in Physics from The Pennsylvania State University in 2021 along with minors in Computer Science and Mathematics.

Hudson Pepper is an undergraduate student at Penn State University majoring in Physics and Mathematics. He specializes in representation theory with a focus in information theory.

A Neuro-Symbolic Approach to AI Agents within Simulation

Samuel F. Griffith, Todd W. Griffith, Ph.D., Isaiah B. Mallery, Brady D. Kirkpatrick, Hudson Pepper
Discovery Machine®, Inc.
Williamsport, PA
tgriffith@discoverymachine.com, samuel@griffith.llc, imallery@discoverymachine.com,
bradydkirkpatrick@gmail.com , hap5183@psu.edu

1 INTRODUCTION

Machine learning, especially deep learning techniques, have proven themselves to be powerful for general purpose pattern recognition applications. This is not a small feat, and for this reason, they have become a popular tool for data scientists and computer scientists alike, especially when tackling correlation problems. However, these techniques frequently possess two key limitations that only increase as the complexity scales: data requirements and program 'black-boxing'. In this paper, we explore how to integrate ML with other AI techniques to increase applicability. There are two areas applicable to modeling and simulation that we have explored. The first is natural language processing and the second is projective simulation.

Deep Learning is at a crossroads. It can find correlations in data that are often overlooked by human analysts, but it says very little about causation. Deep Learning seeks patterns in sensors, sounds, images, videos, documents and databases. In its best incarnation it can selectively answer questions by applying categorization and analysis to the data to answer well-formed questions. It is not, however, a solution to all the problems of cognition. It does not offer explanations of "why" but rather presents statistical beliefs based on patterns found in the data.

Cognition is situated and specific. It is not buried deep in some algorithm or rooted in big data analysis. It is the relationship between people's understanding of their circumstance and the circumstance itself. It is tied to what they can share and how they interact with the world and each other; the language that they use and the objects with which they interact; the people with whom they discourse and the topics that they discuss.

Intelligent agents have been used in modeling and simulation applications since the first semi-automated forces (SAF) systems were developed in the early 1980s. These agents, however, have always been limited in real world military training because while they can make simple decisions, they cannot communicate like a human and they cannot think ahead to see what is coming in the near future. In order for the trainee to experience a realistic scenario, generally requires a "white cell" that can role play the communications of the agent; this limits the number of agents that can be added to the simulation and increases the workload for instructors (i.e., the white cell) conducting the training exercise. In the same respect, simple agents within the SAF must be continuously monitored by the white cell to make sure that they are doing what is expected. The agents are not intelligent enough to make their own decisions, which often requires looking out into the future and reacting to the potential changes and decisions of both friendly and adversary agents within the simulation environment.

Our efforts in this research are focused on addressing these two challenges to the M&S community: communication via natural language and projective simulation. We have demonstrated voice-enabled agents in our work on the Joint Theater Air Ground Simulation System (JTAGSS) which is part of the Joint Terminal Control Training and Rehearsal System (JTC-TRS) program of record being fielded by the Air Force. The work described in this paper is also being leveraged to advance capabilities for larger Air Combat Command distributed training within the Distributed Missions Operations Network (DMON). We envision this early work in projective simulation as a way to provide agents with more independence, thereby enabling instructors in the white cell to focus more attention on the trainee and less attention on directing agents.

In a neural net with the goal of interpretating speech from a user within a simulation, there is a key difference from a human interpreter. Whether or not we realize it, our minds are able to interpret speech patterns within the context of the current situation, and this significantly reduces the potential interpretations of an utterance. To combat this inherent flaw, we have endeavored to create a "guided" neural net which is able to filter possible interpretations based on the factors present within a simulation. A neural net may be capable of recognizing overall patterns in semantics and form an accurate representation of meaning much of the time, however this is much more difficult in domains where data is sparse and language is specific to the domain, e.g., commentary for competitive on-line games or military conversations during combat air support missions. By creating a PPNN we create several neural nets which are independently trained for specific situations in the simulation.

We also detail our exploration of quantum AI particularly how a neuro-symbolic approach might create superior quantum agents. We explore several areas; however, we currently believe that this approach could provide the greatest improvement in projective simulation algorithms which use random walks to explore a quantum memory space. We are aware that this work dabbles in several different subtopics. We believe it necessary to take a moment before continuing to explain the interplay of these subdisciplines in our work, and the purpose for doing so. All of this work pertains to two goals: to advance agents' abilities to predict the future as well as their ability to communicate. Both of these goals could benefit from the use of neuro-symbolic approaches as well as quantum theory. In the speech recognition realm, neuro-symbolic approaches can potentially reduce data costs and improve accuracy. In projective simulation, a neuro-symbolic approach provides the potential for more intricate situational analysis. Additionally, both areas could potentially be improved by quantum algorithms, which have recently been shown to have great potential in the world of AI.

1.1 KNOWLEDGE REPRESENTATION

Our research at Discovery Machine®, Inc. (DMI) is focused on AI training programs for both military and civilian domains specifically knowledge representation. Over the past 20 years, DMI has demonstrated the power of knowledge capture and deployment and has successfully developed agents to aid in training such as virtual instructors for Air Force Pilot Training, well-site technicians, hospital staff, and many more both currently in the field and in development. Knowledge representation is a way of capturing the knowledge of subject matter experts such that an intelligent agent shares the mental models of the expert and can apply those models for problem solving, explanation and instruction. This is useful for training programs for obvious reasons: a trainee needs to know when they are wrong, why they are wrong, and what the correct action is and why. Additionally, for debugging purposes, whenever an agent preforms a task in a seemingly unusual way, it is easy to find out whether the agent was acting correctly or incorrectly depending upon its explanation. However, there is one area where machine learning shines brighter than knowledge representation and traditional simulation programs: recognizing patterns and forming correlations quickly. Using Lanczos and other highly efficient tensor analysis techniques allow rapid data processing which is far beyond traditional programs. It was our goal to extract the advantages of both of these approaches.

1.2 ON DATA AND NATURAL LANGUAGE PROCESSING

Statistical Machine Learning (SML) relies on having large quantities of data in a consistent format. This data is used to train classifiers. Large companies such as Google, Meta, Apple & Amazon, data is formatted by design and leveraged daily. It is also many orders of magnitude larger than any data used in military simulations. In other cases, such as AlphaGo Zero which was trained to play the game Go at a world class level, the data could be collected and generated from online sources including games played against itself. Deep Learning was also used effectively to play the online game StarCraft at a professional level. Once again this was achieved by having enormous amounts of well formatted data of prior games. Finally, there are situations in which statistical machine learning classifiers can be created with smaller amounts of data. These include simple pattern recognition examples such as recognizing handwritten characters or understanding simple speech commands. The improvements described in this paper pertain to complex situations where data is sparse (by comparison) and often inconsistently formatted.

Natural Language Processing (NLP) involves text analysis of a traditional spoken or written language (English, French, Spanish, etc.). The most common way of analyzing this data at present is to semantically represent words as vectors and create SML algorithms to analyze this data. Prior to the pervasive use of SML algorithms, NLP was done by directly programming rules by which the computer could understand language. Languages, by their nature are incredibly complex and contain far more data and rules than could ever be directly coded in a practical manner. Indeed, it turns out that even with the utilization of SML, the data required to properly analyze speech from human beings is incredibly demanding. We attempt to make the data requirement less demanding by breaking a problem down into several neural networks which each have a lesser data requirement for performance and sort them based on interpretation from a knowledge representation (KR).

1.3 HYBRID AI

There are many areas of AI that can be leveraged for modeling and simulation (machine learning, knowledge representation, reflection, etc.). For the purposes of this paper, hybrid AI is a combination of AI approaches forming a multi-strategy application. We explore the applications of hybrid AI to language processing as well as the potential of superior quantum AI, especially in the context of projection simulation (PS).

Aside from the possible practical benefit of combining multiple AI forms, this strategy is motivated by a philosophy concerning the true nature of intelligence. In recent years machine learning has become the most prevalent form of AI being studied in both academic and professional circles. It has yielded some amazing results such as Tesla's self-driving car and even granted scientists an important data analysis tool in areas such as astrophysics. We do not seek to discredit SML as a tool for interpreting large data sets, however, it is not clear that big data interpretation in and of itself will ever constitute a truly intelligent agent. Broken down to its simplest description, SML using neural networks such as Tensor Flow classifiers are matching tools which take varying inputs and matches them to its most likely output. However, it could be argued that an agent which is unable to explain why it makes specific decisions is not intelligent. This begs the question of what is intelligence. Compared to the most obvious example, human beings certainly operate with a great deal of pattern recognition very similar to a SML algorithm. We can interpret images and other patterns in the world and automatically associate them with concepts we have previously encountered. This by itself is insufficient to operate on the level we do however. In addition to taking in large amounts of data and interpreting it, we also perform at a symbolic level, where we make logical decisions based on reason. It is the combination of both of these things that allow humans to perform at such a high level in the real world. It stands to reason that if we base our current understanding of intelligence off human beings, that multiple AI techniques will need to combine to form working agents.

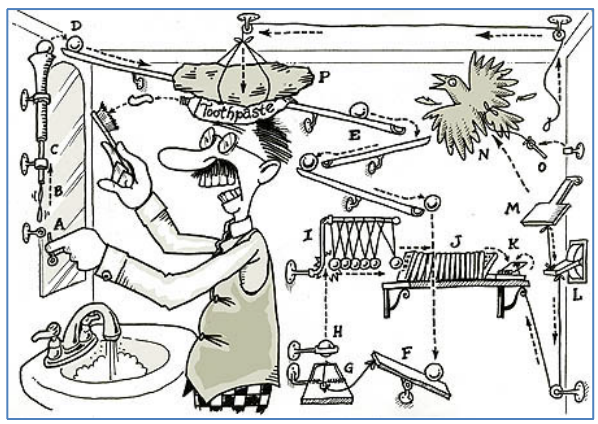

Figure 1: A Rube Goldberg example of Qualitative Physics

One sub area of symbolic artificial intelligence that was started in the late 1970s is the field of Qualitative Physics (see Hayes, 1979). The idea behind qualitative physics is that humans are able to reason about physics of the world in a qualitative way. For example, one knows that when water is poured into a vessel it will start to fill that vessel. One also can envision what happens when a ball is released on an inclined surface. The theory, however, is that one does not precisely compute these "envisionments" quantitatively but rather qualitatively. So, one understands that the volume of liquid in the vessel will increase but not exactly by how much or in what amount of time. These symbolic representations have been leveraged by Forbus in his Qualitative Process Theory (Forbus 1984,1989, 2007) which is where the idea of "envisionments". Our mental simulations of the world are envisionments. They are ways of looking out into the future to see what is going to happen. They are the basis for projection in what makes one situationally aware. The Rube Goldberg picture in Figure 1 is only funny because we can qualitatively simulate what happens at each stage.

In past projects, DMI has built qualitative simulations that enable our agents to project future states of the world using a kind of "envisionment". The difference, however, is that our envisionments of the future also involved the actions and interactions with other agents including adversaries. This made these envisionment computationally intractable allowing only a glimpse into the future. In this paper, we are exploring how a hybrid approach leveraging both symbolic and neural networks for projective simulation can reduce the complexity through the mathematics of quantum mechanics.

1.4 ENSEMBLING

An important tool for consideration is ensembling. It is the process of combining several different neural networks with the same input and output variables (Brownlee, J., 2020). By using a KR in combination with an ensemble of networks, the program is able to determine which of the networks applies most to the situation. In traditional ensembling approaches, several neural nets are applied simultaneously in order to provide a more accurate guess. In our case, we apply neural nets based on the results of the KR.

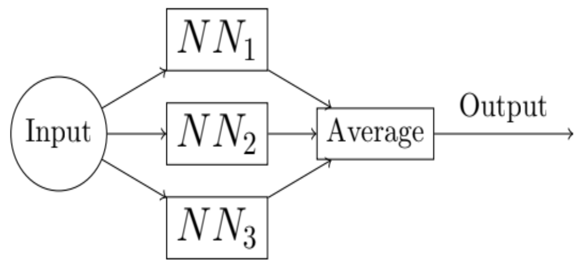

Figure 2: Example Ensemble of Three Neural Networks

In model average ensembling, a series of neural networks are simultaneously applied, and the outputs of all are averaged equally to provide the final result (Brownlee, J., 2020). This can be useful when using several networks which work in tandem to tackle one particular problem. In a weighted average ensemble however, the final outputs are averaged unequally, and can be adjusted based on data in the training dataset (Brownlee, J., 2020). This can be data intensive, so we have made it out goal to mitigate the need for additional training to find these weights. This task instead falls to the KR.

Our most recent efforts have leveraged a relatively new technique for combining neural networks called Transformers which allows us to bring in much larger trained neural networks such as wav2vec and apply what is called one-shot learning to adapt the speech recognition to the domain. These larger networks are trained based on data from big tech giants such as Meta, Amazon, or Google but are trained for “everyday” language rather than language used in specialized domains.

2.1 WHERE WE FOUND OUR DATA

Our working model was centered around data and commentary from a game called League of Legends. We chose this application for several reasons:

1. It is a specialized domain with unique speech recognition requirements.
2. It is one of the most popular games in the world, making the volume of data available very large.
3. All of the data that we used is freely available to the public, thus making this data-set accessible.
4. Riot (the creator of the game) saves rework game data and has a fairly verbose API easily accessible with a personal developer key.
5. Many players will record their games with commentary to put on YouTube, which is then transcribed automatically by YouTube, thus allowing for ease in receiving expert commentary.

Our work sought to combine this data from games with transcripts of YouTube commentary to train it effectively. It was our ultimate goal however to attempt this with military data, in which words from a Joint Terminal Attack Controller (JTAC), for example, could be leveraged by an AI A10 pilot agent for doing Combat Air Support (CAS) missions.

2.2 CONTEXT ANALYSIS

It is useful to consider what we mean by context. For the purposes of this paper, context is any additional data about or occurring at the same time as the data being analyzed. Take the toy but practical example of a call center. The purpose of call centers is to efficiently direct callers into categories for their particular issue (payment, cancellation,

Figure 3: League of Legends Tournament (https://leaguefeed.net/biggest-lol-tournaments/)

problems with service, etc.) that they are dealing with. A context data point might be 'time of day given time zone' or 'length of time on the line'. These variables can be combined with other variables into a parameter called distress to describe the caller, and thus will hint at the language and goals of the call. It is useful to also consider the fact that context is not by itself a determining factor in the classification. Perhaps with the risk of sounding too obvious, any program is finite. As a result, any knowledge representation of an environment has a finite (and often small) number of parameters it represents for the situation. Thus, any environmental features captured or calculated by a KR must therefore be able to be represented entirely in a compact coordinate form. This is important because machine learning can only manipulate data transformed into a numerical format.

2.3 PPNN

Our network structure has a very standard format. Input layers are created for both contextual and language data, with the contextual data being interpreted by the KR before being fed in a compatible way to the NN. This is how we coined the term Parallel Processing Neural Net, as data is being processed simultaneously by two networks which are inputted different data. An important conclusion to draw from this is that as data demands increase, it is possible to incorporate additional data in alternate locations into the analysis. This is at the root of our contribution to speech recognition.

3 RESULTS

While our model was simplistic in nature, we did at least achieve a result which was effective in analyzing basic speech in League of Legends which is a competitive real time strategy game that has professional teams and a fan base (see Figure 3). It is somewhat unclear if this result indicates an overall improvement in speech recognition, however this shows that the concept can be implemented in a practical environment. We created these techniques in the hope of applying them in more complex and data rich environments, such as for military training purposes. This work is currently being continued by employees of Discovery Machine®, Inc. The main piece of data we were able to use to contextualize the speech was when a kill occurred. We were able to determine whether the kill was on a champion (a player), a building or structure, or an elite monster which the player gains bonus's for killing. In league of legends, all of these events hold special significance for the game and vary in their consequences.

4 EXPLORATION OF QUANTUM AI AND CURRENT ENDEAVORS

In the realm of AI, it has been discovered that the mutual dependency of machine learning and quantum theory on tensor manipulation promises to produce superior results in both areas. It was our feeling that if quantum theory could be leveraged to provide improvements for intelligent agents, that this could be a major win for training software. The potential combinations of these fields break into four categories which are nicely summarized by Dunjko and Briegel (2017). They are as follows: Using machine learning to produce superior research techniques in quantum physics, enhancing ML with quantum theory, ML of quantum data, and creation of agents which learn in a quantum environment. Our focus was originally in quantum enhancements of ML and the exploration of the use of KR to enhance these quantum ML algorithms, especially for the purpose of better quantum semantics (use of quantum theory to interpret meaning of speech). More recently, we have also explored introducing KR into quantum projective simulation algorithms. Though we explored multiple ideas, we found that the tensor dependent nature of ML and quantum theory made combining KR into the mix difficult in some applications. However, we hope to provide an overview of the approaches we tried as well as other potential avenues of exploration.

4.1 KNOWLEDGE REPRESENTATIONS TO ADAPT A NEURAL NETWORK

Our first idea and one we spent a great deal of attention on was the possibility of creating a knowledge representation which was able to independently tune the parameters of a neural network to produce the desired result. The main issue with this concept we found is that the creation of neural networks can often be more art than science so to speak, and therefore creating any meaningful process the computer would have "knowledge" of becomes intractable. We also explored the idea of vectorizing contextual data and feeding this back into the neural network as another component. This became difficult because it is unclear how this vectorization could actually be accomplished.

In the mid-90's Mica Endsley (1995) did some groundbreaking work on situation awareness. Her idea was that an agent's awareness occurs in levels: raw information from sensors being the lowest level (level 1), the identification of objects or things being the next level (level 2), and the projection of those objects into the future being the highest level (level 3)[1]. So for a Football Quarterback, the visual sensor information is the first level, identifying the players, receivers and defenders, is the second, while projecting where those players will be in 5 seconds is the third. In the DoD community these ideas were expanded upon in the area of "sensor fusion" by systems tracking information within theaters of operation. For example, Discovery Machine® has automated entities for training that use sensors to pick up something (a track) that is unidentified (level 1), and then by combining it with other sensory information the agent identifies that thing as an adversary (level 2). Once the adversary is identified, the AI entity can project where it expects that adversary to go (level 3).

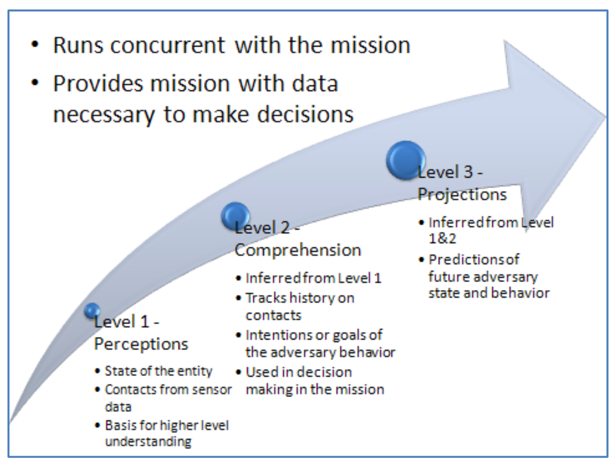

Figure 4: Situational Awareness Processing

The next section of this paper explores this third level of situational awareness by leveraging quantum theory to do projective simulation.

4.2 PROJECTIVE SIMULATION AND QUANTUM THEORY

The method most relevant to this conference in our current area of study is the possibility of quantum algorithms in projective simulation. Projective simulation is a method of creating agents which anticipate future events and simulate them in order to determine which action is ideal. Work by Briegel and De las Cuevas (2012) and Tiersch, Ganahl and Briegel (2015) theorize that implementing a memory space (or series of episodic memories) which the agent uses to

[1] The levels have been described in many ways from 0 to 5. For our purposes, the fact that there are different levels is sufficient.

simulate future events might present advantages. This works by initiating a random memory walk based on probability through the memory space, and initiating an action from that space if it produces a desired result. This involves using a weighted matrix which is altered whenever an event occurs. The successful completion of an action and the measurement of whether its result was good or bad alters the weighted matrix accordingly, adjusting the probabilities for specific nodes in the memory space. There is also a dissipation effect which reduces the weights of all clips over time, simulating “forgetfulness”. This means that sequence of events is roughly as follows: an event occurs, this triggers a random memory walk, this may arrive at an action which the algorithm determines to be good or bad, and the action is initiated only if the simulated result is good. The matrix then alters the weights based on the result. This can be written mathematically as follows:

$$h(n+1) = h(n) - \gamma(h(n) - 1) + \delta\lambda$$

Where h is the matrix, γ is the dissipation and δ is a dirac delta determining if the result was good or bad (Briegel, De las Cuevas 2012).

The example Briegel and De las Cuevas (2012) use to demonstrate this algorithm is an invader game. The agent is tasked with preventing an invader from entering a wall which has equally separated spaces. The invader can move left one or right one, and the agent must choose to defend left or right based on one of two symbols (arrows generally) which are presented to it. The agent is rewarded when it chooses the correct direction. The previously mentioned dissipation factor comes into play if the meaning of the symbols is changed, and the program must relearn their meaning. This demonstrated an increased efficiency compared to standard methods when the situation surrounding the agent changes. In their work they also detail how this episodic memory could be stored in a quantum memory space. Instead of a series of clips in as classical variables, they become quantum states in a complex vector space known as a Hilbert space. This in turn makes the random memory walk a quantum random walk.

A major theme of our work is determining if symbolic approaches may offer improvements in neural networks and simulations. The basis of our exploration into this work relates to how people might recall memories and use them to project into the future vs how this program currently works. Right now, it seems to operate on good and bad evaluators, which is certainly useful when training an agent to perform better over time. However, it is our hypothesis that the information one gains by performing this type of projective simulation has the potential to be far more complex than good result bad result. It is our hypothesis that by encoding more information about the result of an action as well as a means to interpret the result in a symbolic manner could present significant advantages.

For example, consider a scenario in which a chess player encounters a particular strategy or play. He or she might recall how a particular response to said strategy worked and whether the result was good or bad, but more likely the player also recalls particulars of <u>how</u> said response effected the game. This might include possible responses the opponent might employ to the counter attack, how the action might affect the game at a later time etc.. The example used by Briegel and De Las Cuevas (2012) accounts for a simple task in which there are only two option, but as the complexity increases, additional factors are created which can alter a decision.

Our proposed method of implementing this is that rather than returning a 0 or 1 as the good or bad parameter, a tensor is created which is returned which encodes data along a series of parameters. Examples might be hot/cold or heavy/light. This work is theoretical at the moment and still needs a time to develop, however with additional work we believe it may offer benefits. The approach to this might vary depending on whether a quantum memory space or classical variables are used since classical and quantum information are encoded quite differently, however exploring both of those options is certainly a matter for future exploration.

5 SUMMARY OF EFFORTS

The primary theme of our work was to determine if a hybrid neuro-symbolic approach to creating intelligent agents was effective and could provide improvements to existing agents. Our focus was in the improvement of two areas: Projective simulation and speech recognition, both of which are critical for the advancement of modeling and simulation. Additionally, this approach is founded upon a philosophy that it will require more than either of the two approaches alone to create agents with abilities that mimic true intelligence. While this work is in many ways preliminary, we believe that these results indicate a great potential for future work of this nature. We have shown with a limited dataset that these ideas are at least possible to apply to practical problems in speech recognition.

Additionally, while we have not created a practical use for a neuro-symbolic algorithm in the realm of quantum AI, we have determined a potential application for such an approach in projective simulation.

6 ACKNOWLEDGEMENTS

We would like to acknowledge RIOT games for exposing League of Legends game data and Northrop Grumman for sponsoring work to transition our efforts for A10 CAS missions.

7 REFERENCES

Chandrasekaren, B. (1986) Generic Tasks in Know-ledge-Based Reasoning: High-Level Building Blocks for Expert System Design, *IEEE Expert*. Fall, 1986.

deKleer, J. & Brown, J. S. (1982) Foundations of envisioning. Proceedings of the Second National Conference on Artificial Intelligence. 434-437.

Endsley, M.R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, *37(1)*, 32-64.

Forbus, K. (1984) Qualitative Process Theory. *Artificial Intelligence*, 24:85-168.

Forbus, K. (1989) Introducing actions into qualitative simulation. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1273-1278.

Forbus, K. (2007) Qualitative modeling. In Harmelen, F., Lifschitz, V., and Porter, B. (Eds) *Handbook of Knowledge Representation*. Elsevier.

Garcia, C.J. & Griffith, T.W. (2005) *A Composable Behavior Modeling System for Rapidly Constructing Human Behaviors,* Proceedings of I/ITSEC 2005.

Griffith, T.W. (2006) *Domain Specific Knowledge Capture Interfaces for Behavior Modeling,* Proceedings of I/ITSEC 2006

Hayes, Patrick, J. (1979) "The Naive Physics Manifesto", in D. Michie, ed., *Expert Systems in the Micro-Electronic Age*, Edinburgh: Edinburgh University Press, 242-70, as repr. in Boden, ed., 171-205.

Hegde, M., Allison, D. & Griffith, T. (2014). A Framework for Enabling Virtual Observer Controllers in Synthetic Training. *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference.*

Hegde, M., Allison, D. & Griffith, T. (2014). Reflecting on Cognitive Process Models to Enhance Adaptive Training. *MODSIM World Proceedings.*

Piedeleu, R., Kartsaklis, D., Coecke, B., & Sadrzadeh, M. (2015, February 4). *Open system categorical quantum semantics in natural language processing*. arXiv.org. Retrieved from https://arxiv.org/abs/1502.00831

Potts, J., Griffith, T., Sharp, J., & Allison, D. (2011). Subject Matter Expert-Driven Behavior Modeling Within Simulation. *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2011*.

Potts, J.R., Griffith, T., Roth, K., & Snyder, J.. (2012) Customizable Speech Centers for Automated Entities within Simulation. *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference.*

Sharp, J. J., & Potts, J. R. (2011). Improving Trainee Engagement Levels through Adaptive Entity Behaviors. *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2011*.

Snyder, J. K., Morse, S. R., Potts, J. R., & Griffith, T. (2013) Cognitive Projection of Future States by Autonomous Entities. *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*.

Kahnemen, D. (2011) *Thinking, fast and slow: Combining vector ... - arxiv.org*. (2017, August 21). Retrieved from https://arxiv.org/pdf/1708.03310.pdf

Briegel, H. J., & De las Cuevas, G. (2012, May 15). *Projective simulation for artificial intelligence*. Nature News. Retrieved from https://www.nature.com/articles/srep00400

Tiersch, M., Ganahl, E. J., & Briegel, H. J. (2015, August 31). *Adaptive quantum computation in changing environments using projective simulation*. arXiv.org. Retrieved from https://arxiv.org/abs/1407.1535

Dunjko, V., & Briegel, H. J. (2017, September 8). *Machine learning \& artificial intelligence in the quantum domain*. arXiv.org. Retrieved from https://arxiv.org/abs/1709.02779

Brownlee, J. (2020, August 25). *How to develop a weighted average ensemble for Deep Learning Neural Networks.* Machine Learning Mastery. Retrieved from https://machinelearningmastery.com/weighted-average-ensemble-for-deep-learning-neural-networks/