Closed-Loop Neuromorphic Artificial Intelligence for Decision Support

Daniel Barber and Lauren Reinerman-Jones
Southwest Research Institute
San Antonio, TX
daniel.barber@swri.org, lauren.reinerman-jones@swri.org

ABSTRACT

Recent developments in autonomous fighter jets and concepts such as an autonomous wingman are pushing the boundaries for the integration of artificial intelligence (AI) into high-risk military flight operations. Many of these concepts explore how autonomy can accelerate pilot decision-making to reduce cognitive demands, improve data processing, and enhance electronic warfare capabilities. Meeting the demands of high speed, high risk, and hardware constrained environments typical of military flight has inspired exploration of neuromorphic technology. Neuromorphics offers low size, weight, and power (SWaP) and high-speed parallel processing capable of delivering advanced AI abilities. Neuromorphic solutions, which pair biologically inspired hardware and spiking artificial neural networks, are still a nascent technology with many challenges to overcome before they can be fully integrated and flight worthy. One key challenge is demonstrating performance greater than or equal to existing artificial neural networks running on conventional hardware (e.g., Graphics Processing Unit) in an ecologically relevant use case. In this paper, we discuss the exploratory research to develop a closed-loop decision support aid for pilots that combines a spiking neural network (SNN) running on Intel's Loihi 2 neuromorphic chip with the U.S. Air Force's Advanced Framework for Simulation, Integration, and Modeling (AFSIM). Specifically, this paper describes the architecture for a closed-loop framework supporting use of different AI models to generate tactical decision recommendations in an adversarial flight scenario. Finally, an approach for the creation of an SNN, executed on Intel's Loihi 2 neuromorphic chip, using a pre-trained deep neural network (DNN) is described along with a comparison of SNN and DNN performance in a closed-loop simulation.

ABOUT THE AUTHORS

Dr. Daniel Barber has 15+ years of experience conducting and applying interdisciplinary research within the fields of human-machine interaction, robotics, machine-learning, modeling and simulation, augmented cognition, physiological assessment, control systems, path-planning, communication frameworks, and environment modeling. In the execution of these efforts, Dr. Barber has also developed multiple prototype LVC systems and autonomy enabling human machine interaction in the domains of counter small-unmanned air systems (C-sUAS), mounted and dismounted human robot teaming for the DoD, driverless cars, and nuclear power plant main control room operations. He has extensive experience executing human-in-the-loop experiments both in the field and laboratory to evaluate human-machine team performance and cognitive demands. Dr. Barber is currently researching neuromorphic artificial intelligence, including the development of spiking neural networks for hardware constrained environments.

Dr. Reinerman-Jones is currently an Acting Section Manager and Principal Analyst at Southwest Research Institute (SwRI) leading the neuromorphics team and Dayton satellite office. She has vast experience in human factors psychology, modeling, simulation, experimental design, technology, training, and AI enabling her to be a strong leader in synergistic solutions to problem spaces. Her extensive experience includes working with the Army, Air Force, Navy, Marines, Special Forces, DARPA, IARPA, FBI, and Nuclear Regulatory Commission. Her internal driving force is to deliver revolutionary scientific discoveries and advanced capabilities leaving a legacy of greatness for generations to come.

Closed-Loop Neuromorphic Artificial Intelligence for Decision Support

Daniel Barber, Lauren Reinerman-Jones

Southwest Research Institute

San Antonio, TX

daniel.barber@swri.org, lauren.reinerman-jones@swri.org

INTRODUCTION

Advances in autonomous fighter jets and concepts such as an autonomous wingman are pushing the boundaries for the integration of artificial intelligence (AI) into high-risk military flight operations. DARPA broadcast the Alpha Dogfight Trials live on YouTube in the middle of the pandemic, whereby AI controlled aircraft faced-off against pilot-controlled aircraft in a virtual environment with the AI winning, (DARPA, 2020). These trials fed the DARPA Air Combat Evolution (ACE) program, which aimed to advance and transition higher-level cognitive AI functions from simulation to real aircraft with a focus on objective operator trust calibrated with the dogfighting agents' intent, (DARPA, 2020). In 2024, this program achieved the first-ever in-air tests of AI algorithms autonomously flying an F-16 against a human-piloted F-16 in within-visual-range combat scenarios, ushering in the future of autonomy in tactical aerospace and paving the way for the next generation of AI, (DARPA, 2024). Other AI concepts outside of autonomous dogfighting capabilities for military aircraft include acceleration of pilot decision-making to reduce cognitive demands, improve data processing, and enhanced electronic warfare capabilities.

Meeting the growing demand for integration of new AI capabilities into high speed, high-risk, and hardware constrained environments typical of military flight has inspired the exploration of neuromorphic technology. Neuromorphic AI seeks to build computational systems consisting of biologically inspired architecture, which involves creating spiking neural networks (SNNs), a type of artificial neural network (ANN), paired with specialized hardware and software that emulates the behavior of biological neurons. Pairing of SNNs with biologically inspired hardware offers low size, weight, and power (SWaP) and high-speed parallel processing capable of delivering advanced AI capabilities to new and existing platforms that cannot support graphical processing units (GPUs) used in current state-of-the-art AI. As a nascent technology, a key challenge for neuromorphics is demonstrating performance greater or equal to existing ANNs running on conventional hardware in ecologically relevant use cases. In this paper, we discuss the trials in developing a closed-loop decision support aid for pilots that combines a SNN running on Intel's Loihi 2 neuromorphic chip with the U.S. Air Force's Advanced Framework for Simulation, Integration, and Modeling (AFSIM). Specifically, this paper describes architecture and performance of a closed-loop framework supporting use of different versions of a deep neural network (DNN) (PyTorch on GPU vs SNN on Loihi 2) to generate tactical decision recommendations in an adversarial flight scenario.

DECISION SUPPORT TASK FOR HUMAN-MACHINE TEAMING

The U.S. Air Force (USAF) Doctrine Note 25-1, Artificial Intelligence, describes AI and anticipated roles in air operations across the competition continuum, (U.S. Air Force, 2025). This doctrine states the USAF approach to AI employment will be human-machine teaming (HMT) with AI systems augmenting the performance of Airmen and the execution of USAF operations. Common constructs in HMT are human-in-the-loop, where a machine (or AI) makes recommendations and the person makes the decision, and human-on-the-loop when the machine makes recommendations and implements them unless the person vetoes the machine action. Based on the direction the USAF is taking for AI, complemented in programs like DARPA ACE and defined in doctrine, the Southwest Research Institute chose to simulate a human-in-the-loop use case to explore the capabilities of neuromorphic technology and compare them to traditional AI hardware.

The first step in development of a deep neural network (DNN) and spiking neural network (SNN) for evaluation of neuromorphic technology requires a relevant dataset for training the models and an environment to assess performance. The research team selected the USAF's Advanced Framework for Simulation, Integration, and Modeling

(AFSIM) to meet this need. AFSIM is a government owned, open source, community-informed, military simulation framework supporting multi-domain capabilities from sub-surface to space, (DCS Corporation, 2025). AFSIM includes a set of out-of-the-box models and scenarios and JavaScript-like scripting language with Python language support. This combination of features provided a framework to simulate ecologically valid scenarios for a decision-support capability that enables human-in-the-loop HMT. Using AFSIM, Southwest Research Institute (SwRI) researchers modeled a generic fighter jet with the goal of flying a route and bombing a ground target (i.e., surface-to-air missile site) while surviving any encounters with other enemy fighters, Figure 1. An AFSIM scripted behavior tree implemented the decision-making logic for the simulated aircraft which included three primary tasks: ROUTE, FIGHT, and EGRESS. The ROUTE task instructed the aircraft to follow pre-planned routes towards the ground target, FIGHT resulted in engaging enemy aircraft, and EGRESS resulted in the aircraft returning home. As proof of concept for exploration of neuromorphic technology, the goal was to develop a DNN and SNN modeling the rules of this behavior tree based using the same information it uses which includes state of the aircraft, distance to target, weapon payload status, and enemy fighters that RADAR sensors detect.

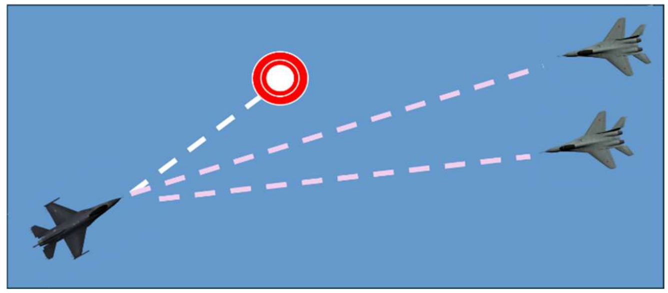

Figure 1. Example illustration of decision support scenario, with fighter attempting to bomb a target while surviving any enemy engagements.

Dataset Capture and Simulating Decision Support Interaction

To simulate human-in-the-loop interaction between an AI agent and piloted generic fighter aircraft, SwRI engineers created re-usable Python modules enabling bi-directional communication between AFSIM simulations scripts and external applications using the industry standard and open-source messaging library ZeroMQ, (ZeroMQ, 2024). This bi-directional interface supported two primary functions: 1) capture of aircraft state and sensor perception data and 2) sending tactical decision recommendations to the simulated aircraft. A standalone Python script executed AFSIM simulations in batched runs to capture data for constructing training, validation, and test datasets for developing DNN and SNN models. Developers created multiple variations of the previously described bombing scenario in AFSIM to produce variation across scenarios while balancing the number of each decision type to minimize bias within the datasets, Table 1. Manipulations to scenarios include, but not limited to, routes followed, distances to bombing target, initial speed and direction, starting locations, and number of enemy aircraft. Barber & Reinerman-Jones, 2025, provides further details on the creation and validation of the datasets.

Table 1. Sample scenarios and the amount of time (%) each decision type is represented.

# Enemy	Scenario Description	ROUTE	FIGHT	EGRESS
0	ROUTE to target, drop bomb, then EGRESS home.	50%	0%	50%
1	ROUTE to target. Switch to FIGHT defeat enemy aircraft, resume ROUTE, then EGRESS home after bombing target.	33%	34%	33%
1	ROUTE to target. Switch to FIGHT defeat enemy aircraft, EGRESS when/if bingo fuel/weapons.	32%	50%	18%
2	ROUTE to target. EGRESS if unable to bomb without engagement of two (2) enemy aircraft.	25%	0%	75%
Totals		32%	36%	32%

A second Python application, hence-forth called the “AI Agent”, supported execution of either a DNN or SNN model to generate tactical decision recommendations and transmit them to an AFSIM simulation, Figure 2.

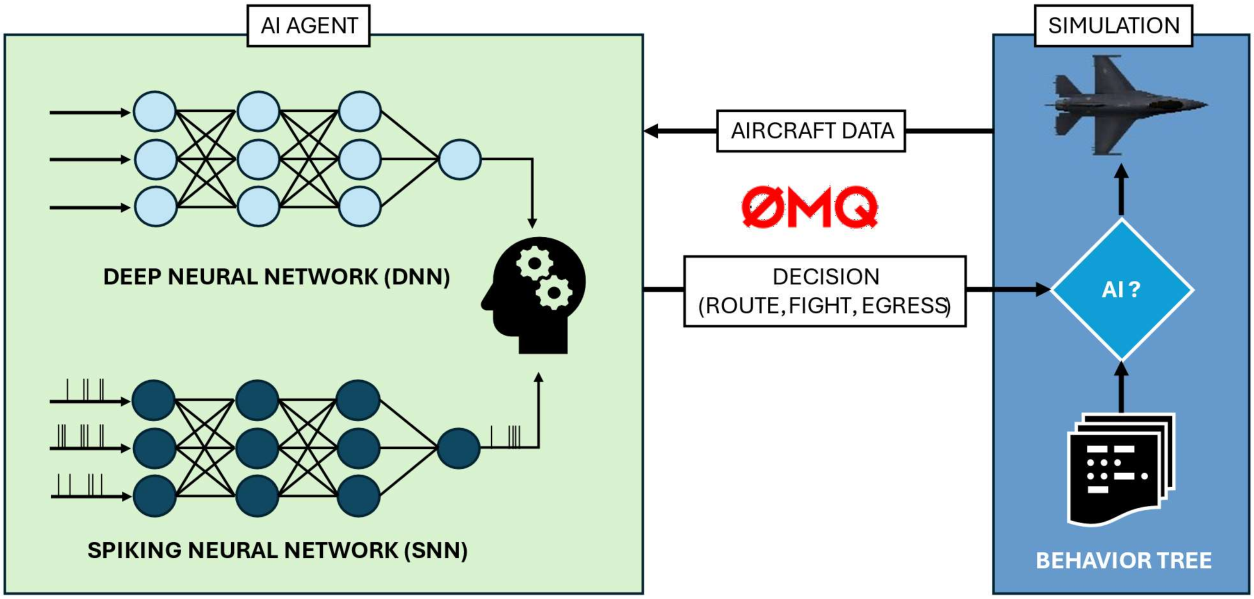

Figure 2. High-level diagram of simulated human-in-the-loop interaction between AI agent and generic fighter aircraft. The AI agent application loads either a DNN or SNN to generate a decision recommendation from aircraft state and sensor detections received. The simulated aircraft follows scripted behavior tree rules unless it receives a decision recommendation from the agent.

Using different program arguments at startup, the AI Agent application loads model weights and parameters to instantiate either a DNN implemented in PyTorch or an SNN implemented in Intel's NxKernel libraries for the Loihi 2 neuromorphic chip, (Intel, 2021). The application then attempts to connect to an active AFSIM simulation using the previously described interfaces, and once connected, feeds captured aircraft data to the selected artificial neural network (ANN) and transmits recommendation outputs back to AFSIM for execution. The neuromorphic hardware used for the SNN was a 3U VPX board containing 16 Loihi 2 Chips and a Cyclone V System on Chip (SoC) FPGA with an ARM-based processor running Ubuntu 20.04 LTS. The computer running AFSIM and the AI Agent, called the "Super Host," connects to the 3U VPX board with 1 GB ethernet. The NxKernel libraries load the SNN onto the neuromorphic hardware using drivers running on the Ubuntu 20.04 LTS ARM-based "Host." NxKernel libraries also provide input/output streaming to the Loihi 2 chips and conversion of integer data to and from spike signals, see Figure 3.

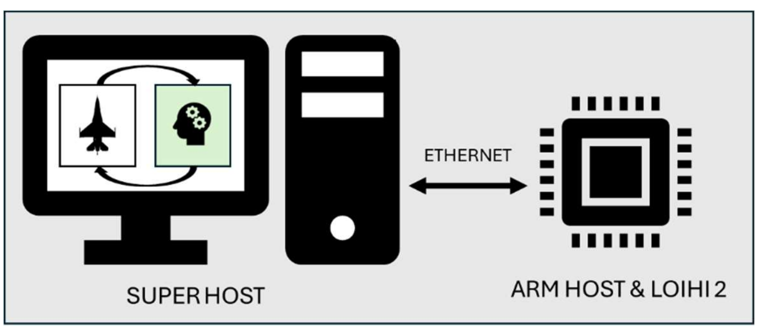

Figure 3. Hardware configuration for SNN execution.

SPIKING NEURAL NETWORK APPROACH

Two equivalent models were created for comparison of accuracy between a deep neural network (DNN) running on conventional hardware and a spiking neural network (SNN) on the Loihi 2 for the decision support task. The approach for an "apples-to-apples" comparative analyses chosen included three steps: 1) develop a DNN which supported conversion to an SNN, 2) convert the DNN from a floating-point representation to fixed-point integer values to support SNN conversion, and 3) convert the DNN to an SNN using NxKernel libraries and the pre-trained weights and bias values. This approach is beneficial for comparing the DNN and SNN models as the overall structure of the networks are equivalent and it bypasses challenges associated with supervised learning in SNNs, (Lyashenko, 2024; Xiangwen, Xianghong, & Xiaochao, 2020).

DNN Model Design

As indicated above, the first step in the development of an SNN was the creation of a DNN capable of correctly modeling the behavior tree decision outputs scripted in AFSIM. This DNN must also support quantization and translation to an SNN. Quantization is a technique used to reduce the size and computation costs of models by replacing high-precision floating-point numbers (like 32-bit floating-point) with lower precision 16-bit floating-point or 8-bit integers. The Loihi 2 neuromorphic chip does not support floating-point data, therefore this quantization step is critical, as 8-bit integer values are needed for generation of a spike train inputs. Therefore, a DNN model was created and trained using PyTorch, then quantized using the built-in PyTorch inference conversion tools. The resulting DNN included 8 blocks, each block containing a linear layer, batch norm layer, and a rectified linear unit (ReLU) activation function. A dropout layer was included in each block for training. The head contains the final output layers with output class (i.e., decision type) determined from the maximum value in the layer, Figure 4. The 32-bit floating point model was trained using supervised learning via gradient descent and achieved an overall accuracy of 99.9% on the validation dataset. Accuracy is defined as the percentage of model outputs matching ground truth on the data set (i.e., what the task decision should be given current state of the environment).

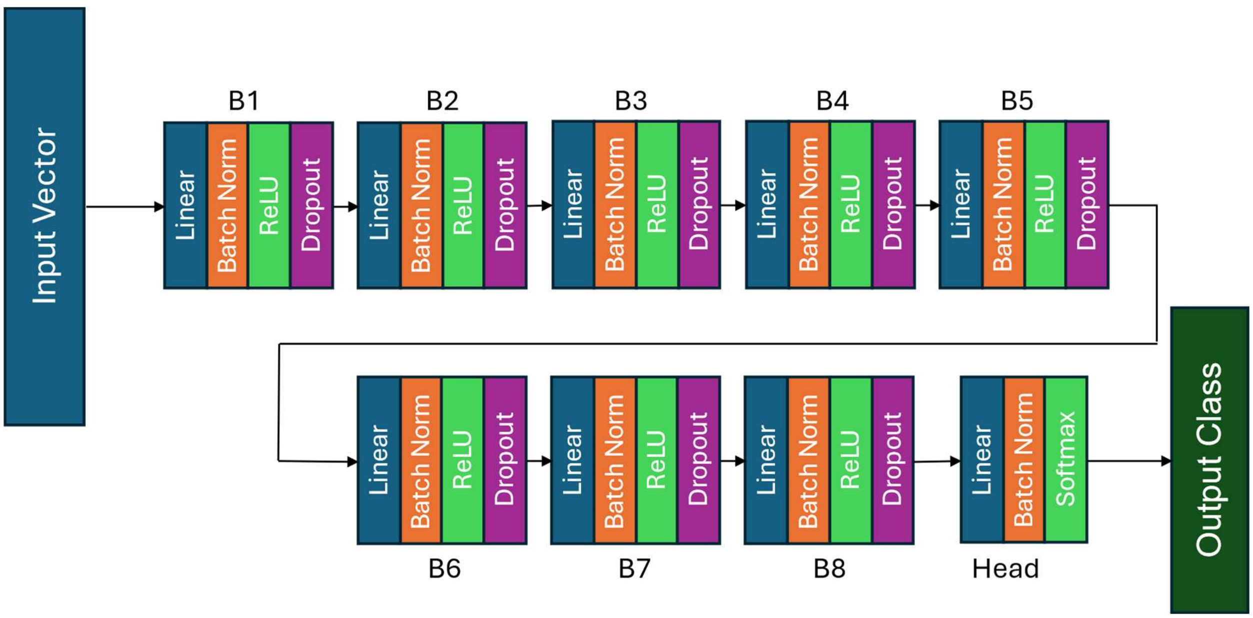

Figure 4. Deep Neural Network (DNN) topology.

Conversion to SNN Model

Following the previously described steps, the 32-bit floating point DNN was pruned and quantized. Near zero weights do not contribute significantly to network output and the Loihi 2 is able to take advantage of pruned weights to reduce latency and increase throughput. PyTorch tools pruned the model with a factor of 0.8 (80%) resulting in no significant loss in accuracy on the validation dataset, demonstrating 99.9% accuracy. After confirming the pruned model had not lost performance, the weights and bias values of the model were quantized using PyTorch quantization tools. Additionally, the weights for the linear function and batch normalization functions of each block were combined into a single linear layer to support SNN conversion in NxKernel. Finally, each linear function does not have an associated bias so the bias from the batch normalization layer was used instead, see Figure 5.

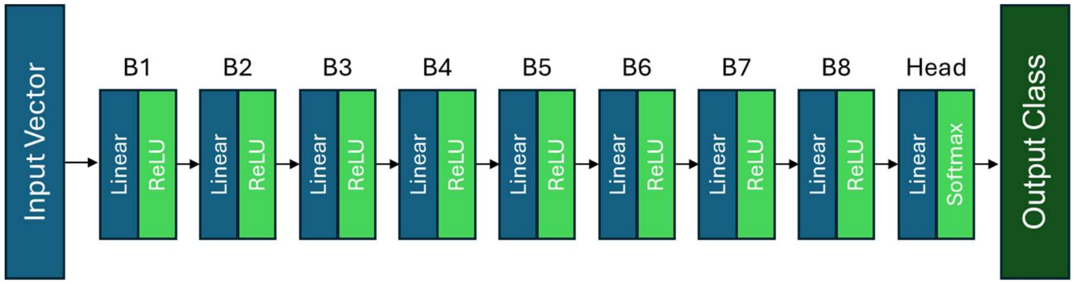

Figure 5. Quantized Deep Neural Network (DNN) topology.

The resulting quantized model was assessed on the validation dataset to check for any significant loss in performance and demonstrated 99.8% accuracy. Based on this finding, the PyTorch quantized model was converted into an SNN using equivalent structures for representation of linear layers and ReLU activation functions within Intel's NxKernel Python libraries. The fully implemented SNN and original 32-bit floating point DNN were then evaluated on the combined training and validation datasets for the decision support task demonstrating accuracy of 99.8% and 99.9% respectively, see Figure 6 and Figure 7.

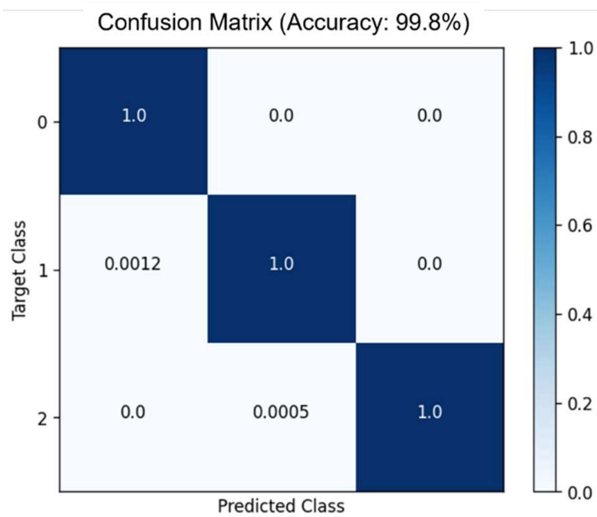

Figure 6. SNN accuracy. Target classes are defined as 0 = ROUTE, 1 = FIGHT, 2 = EGRESS.

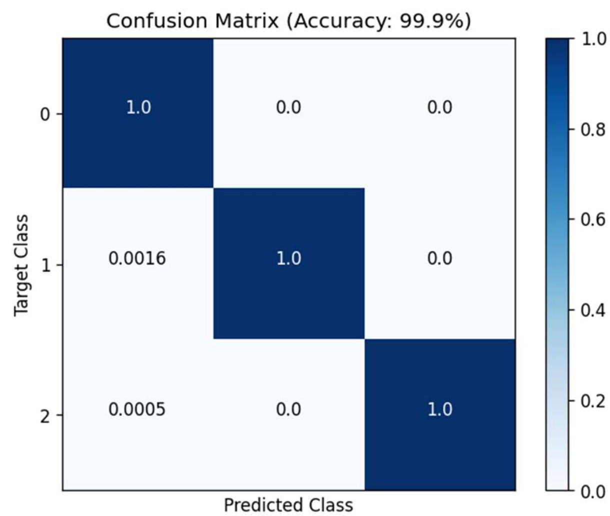

Figure 7. FP32 DNN accuracy. Target classes are defined as 0 = ROUTE, 1 = FIGHT, 2 = EGRESS.

CLOSED-LOOP INTEGRATION

After completing development of both the DNN and SNN models they were integrated into the previously described AI Agent Python application to assess performance on test scenarios running in AFSIM real-time. Both SNN and DNN models were first assessed on a test dataset captured from the test scenarios prior to running closed-loop with AFSIM and demonstrated accuracy identical to the combined test and validation datasets, 99.8% and 99.9% respectively. Next, the AI Agent application was connected to live instances of AFSIM to assess the ability of each model to demonstrate the same accuracy when executed in closed-loop with AFSIM vs on recorded data. A Graphical User Interface (GUI) was also added to the application to help visualize the decision recommendations in real-time, Figure 8.

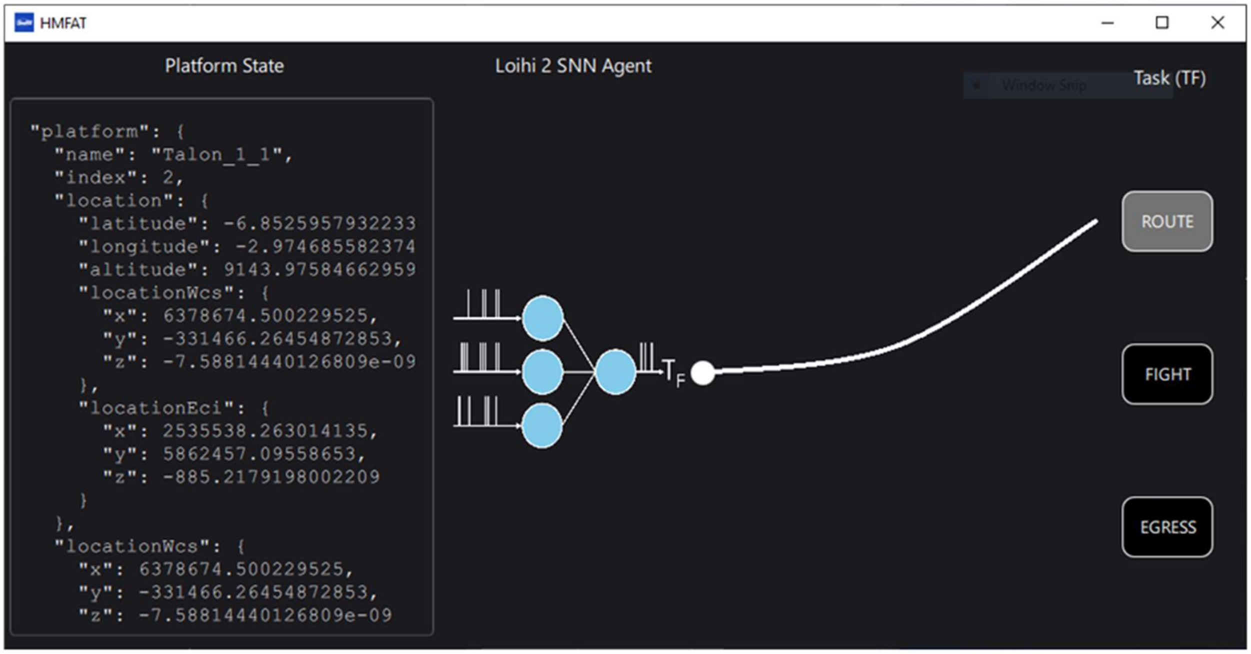

Figure 8. AI Agent application executing an SNN on the Loihi 2 to generate decision support recommendations.

When running in a closed-loop, the original behavior-tree generating in AFSIM is still active and generating tactical decisions. However, as previously described, this output is overridden when a task output is received from the AI Agent, simulating a pilot implementing the recommendation. Both values are recorded during scenario application for comparison against the pre-recorded data and revealed identical results for both DNN and SNN models, Table 2.

Table 2. Accuracy of models on test scenarios.

Execution Type	**Pre-Recorded Data**		**Closed-Loop (Real-Time)**	
Model	**FP32 DNN**	**SNN**	**FP32 DNN**	**SNN**
Accuracy	99.9%	99.8%	99.9%	99.8%

CONCLUSION

This paper describes early exploratory research to investigate the feasibility of using neuromorphic computing for a decision support task within an ecologically valid use-case. Although the models implemented here demonstrated extremely high accuracy, it should be noted that dataset and number of decision classes modeled were relatively simple in nature. The number of decision types and scenarios were limited for this effort to focus on the process of implementing spiking neural networks (SNNs) and integrating them within a relevant closed-loop application. That being said, a major finding from this effort is the successful conversion of a pre-trained deep neural network (DNN) to an SNN implementation with no significant loss in accuracy for a classification task. The second key outcome for this effort was the successful closed-loop execution of an SNN on neuromorphic hardware. These results show significant promise for neuromorphic technologies to deliver AI with greater than or equal performance to AI running on conventional hardware (e.g., GPUs) that is capable of running in low SWaP environments. Future work in this area will further investigate and compare latency, throughput, and power of SNNs running neuromorphic hardware, like the Loihi 2, in comparison to desktop GPUs. Intel's has recently released updates to their NxKernel libraries to support the collection of these metrics in a standardized way which was not available at the time of this effort. In addition to including more standardized benchmarks when comparing neuromorphics hardware to GPUs, future efforts in this domain should explore increasing task complexity to include a higher number of tactical decisions at different hierarchies. For example, this application selected relatively high-level decisions (e.g., follow route, fight), but could include more detailed decisions such as use of different weapons or engagement tactics as well.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the contributions of Brian Millikan and Howard Yanxon in the advancement of this effort.

REFERENCES

Barber, D., & Reinerman-Jones, L. (2025). Simulation for Artificial Intelligence Modeling and Assessment. *16th International Conference on Appled Human Factors and Ergonomics.* Orlando, FL: AHFE International.

DARPA. (2020). *ACE: Air Combat Evolution*. Retrieved 06 01, 2025, from DARPA: https://www.darpa.mil/research/programs/air-combat-evolution

DARPA. (2020, August 26). *AlphaDogfight Trials Foreshadow Future of Human-Machine Symbiosis*. Retrieved June 10, 2025, from DARPA: https://www.darpa.mil/news/2020/alphadogfight-trial

DARPA. (2024, 04 17). *ACE Program Achieves World First for AI in Aerospace*. Retrieved 06 03, 2025, from DARPA: https://www.darpa.mil/research/programs/air-combat-evolution

DCS Corporation. (2025). *DCS*. Retrieved 06 23, 2025, from AFSIM: https://www.dcscorp.com/home/what-we-do/afsim/

Indiveri, G., Linares-Barranco, B., Hamilton, T. J., van Schaik, A., Etienne-Cummings, R., Delbruk, T., . . . Boahen, K. (2011). Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience, 5*. doi:10.3389/fnins.2011.00073

Intel. (2021, 09 01). *Neuromorphic Computing: Intel Lab's Loihi 2 Chip*. Retrieved 06 01, 2025, from Intel Labs: https://www.intel.com/content/www/us/en/research/neuromorphic-computing-loihi-2-technology-brief.html

Lyashenko, V. (2024, 10 22). *Basic Guide to Spiking Neural Networks for Deep Learning*. (Intel Tiber AI Studio) Retrieved 10 22, 2024, from https://cnvrg.io/spiking-neural-networks/

Schuman, C., Potok, T., Patton, R., Birdwell, D., Dean, M., Rose, G., & Plank, J. (2017, May 19). *A Survey of Neuromorphic Computing and Neural Networks in Hardware.* Retrieved 01 09, 2025, from https://arxiv.org/abs/1705.06963

U.S. Air Force. (2025). *Air Force Dcontrine Note 25-1, Artificial Intelligence (AI).* U.S. Air Force. Retrieved from https://www.doctrine.af.mil/Portals/61/documents/AFDN_25-1/AFDN%2025-1%20Artificial%20Intelligence.pdf

Xiangwen, W., Xianghong, L., & Xiaochao, D. (2020). Supervised Learning in Spiking Neural Networks: A Review of Algorithms and Evaluations. *Neural Networks, 125*, 258-280. doi:10.1016/j.neunet.2020.02.011

ZeroMQ. (2024). *ZeroMQ*. Retrieved 02 12, 2025, from ZeroMQ: https://zeromq.org/