Artificial Intelligence Exploration (AIE) Opportunity
DARPA-PA-19-03-04
Signal Processing in Neural Networks (SPiNN)

I. Opportunity Description

The Defense Advanced Research Projects Agency (DARPA) is issuing an Artificial Intelligence Exploration (AIE) Opportunity inviting submissions of innovative basic or applied research concepts in the technical domain of intelligent edge signal processing and application. This AIE Opportunity is being issued under the Program Announcement for AIE, DARPA-PA-19-03. All awards will be made in the form of an Other Transaction (OT) for prototype project. The total award value for the combined Phase 1 base and Phase 2 option is limited to $1,000,000. This total award value includes Government funding and performer cost share, if required or if proposed.

A. Introduction

Accurately communicating multi-dimensional complex modulated signals through non-ideal dynamic communication channels is critical to many DoD radio frequency (RF) sensing and communication applications. Conventional digital signal processing (DSP) techniques recover distorted signals by executing dedicated processing physics models to mitigate specific impairments sequentially. They assume stationary channel models with Gaussian noise, and therefore have very limited capability to process temporal dispersion, non-linear distortions, or interference and jamming artifacts. These cascaded operations constrain the signal recovery with high error rate conditions and are incapable of discovering and mitigating unknown impairments beyond established simple channel models. Such approaches are also computationally intensive, with long latency and poor size, weight, power, and cost.

Emerging machine learning techniques promise a new generation of computational approaches with reduced compute complexity and latency, as well as potentially new capabilities in discovering and processing events beyond established models. For example, recent advances in Deep Neuromorphic Network (DNN) demonstrate fast feed-forward inference to achieve good accuracy once it is trained with high-quality data sets. Currently, DNNs are trained by data sets and do not use physics-based mathematical models. Missing corner cases and other unseen events beyond the collected data sets often leads to insufficient or misinterpreted representations to cause critical mission failures. To establish a reliable and accurate DNN model, remote cloud computing facilities are needed to support a vast computational workload on a large volume of training data. This practice makes DNNs impractical for many DoD machine learning models, which must be timely in adapting their pre-training exemplars to dynamic events and unforeseen corner cases under rapidly changing field conditions.

By examining the required amount of training data to represent a physical system, the performance measures can be illustrated between physics-derived mathematical models and data-driven learning models (such as DNNs), as shown in Figure 1.

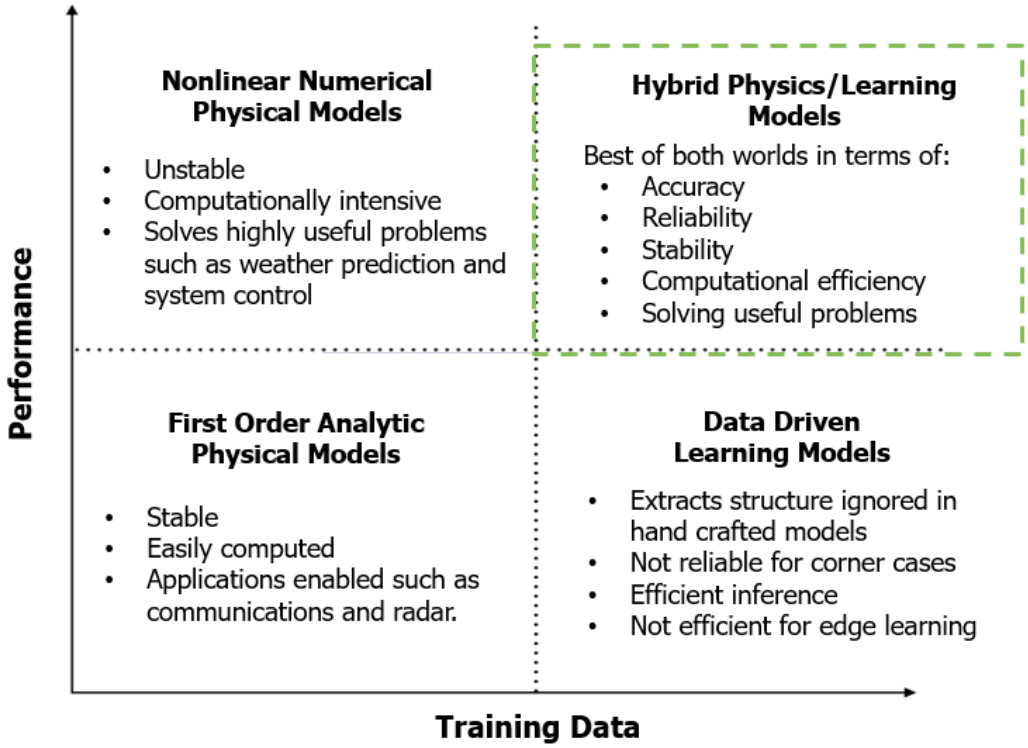

Figure 1: Motivation for Hybrid Physics/Learning Models

This relationship has led exploration into hybrid learning models to embed physics knowledge into learning algorithms, as well as to use learning capabilities to extract hidden structure in the causal systems. For example, the recent DARPA Physics of AI (PAI) program[1] is seeking "to facilitate better incorporation of AI into DoD systems [by] exploring novel AI architectures, algorithms, and approaches that 'bake in' physics, mathematics, and prior knowledge relevant to DoD application domains. PAI aims to show that embedding physics and prior knowledge into AI will help to overcome the challenges of sparse data and will facilitate the development of generative models that are causal and explanative." By embedding scattering symmetries known from physics into learning models, the PAI program has demonstrated a 100x speedup in NN computation of nonlinear turbulent systems. PAI has shown that by using known physics models of diffusion in blood vessels, scientists can more accurately predict how tumors evolve[2].

The Signal Processing in Neural Networks (SPiNN) program will develop a new set of advanced neural network (NN) computing kernels that embed established physics-based mathematical digital signal processing (DSP) models.

B. Objective/Scope

The SPiNN program will leverage the established physics-based signal processing algorithms and associated mathematical tool kits as a basis to establish a set of trained, verifiable, accurate and efficient neural network kernels. SPiNN seeks to transpose important linear and non-linear DSP function blocks such as Fast-Fourier Transform (FFT/iFFT), Multi-Input Multi-Output (MIMO), Matched Filter (MF), Kalman Filter (KF), trellis/Viterbi decoders, error-correction codes, etc. with verifiable outcome and accuracy into pretrained and low latency neural network kernel representations. As these pretrained neural network kernels will be fine-tuned according to real-world data and adapted generatively, they are expected

[1] https://www.darpa.mil/program/physics-of-artificial-intelligence
[2] https://www.youtube.com/watch?v=mw6Rza-BQj4

to outperform traditional DSP models, which lack the inherent capability to capture and process events (e.g., nonlinearities, temporal and spatial dispersions, multi-modal, non-periodic events) that are difficult to model.

SPiNN kernels will be built on these verified DSP model sets as a foundation to establish pre-trained NN discriminators to efficiently process the incoming data with known accuracy. It will then combine the trained NN discriminator block with a generative NN block and adaptive learning transform layer to form a generative-adversarial network kernel, which will capture corner cases and extract additional hidden structures beyond the known DSP models. The resulting SPiNN adaptive NN kernels will provide optimal and robust signal processing capability with accurate performance in real time when facing a dynamic real world environment. When adaptation of the NN kernel is handicapped by the poor corner cases, it would also provide verified fallback performance to the traditional model-based DSP practice. The final resulting enhanced SPiNN kernels are expected to form a foundation for a new generation of adaptive machine learning signal processors.

Data Sharing

The SPiNN research effort will depend in part upon the open exchange of data sets and common interface of emulator suites between performers involved in the effort. Once selected for SPiNN, the performers will be required to share data set and common interface to the emulation suites. Proposals should include plans for sharing data and any necessary data use approvals required for including other performers. DARPA will not be serving as a data distributor, and will not be taking possession of any of the datasets. Therefore, the use of sensitive, proprietary, and/or non-identified data is not encouraged. Proposers should list any anticipated risks and mitigation strategies associated with these requirements and their approaches. If the usage of intellectual property (IP) blocks in the designs are needed, performers will establish all appropriate agreements with others regarding restrictions on usage of IP.

C. Structure

Proposals submitted to DARPA-PA-19-03-04 in response to the technical area of this AIE Opportunity must be UNCLASSIFIED and must address two independent and sequential project phases: a Phase 1 Feasibility Study (base) and a Phase 2 Proof of Concept (option). The periods of performance for these phases are 9 months for the Phase 1 base effort and 9 months for the Phase 2 option effort. Combined Phase 1 base and Phase 2 option efforts for this AIE Opportunity should not exceed 18 months. The Phase 1 (base) award value is limited to $500,000. The Phase 2 (option) award value is limited to $500,000. The total award value for the combined Phase 1 and Phase 2 is no more than $1,000,000. This total award value includes Government funding and performer cost share, if required or if proposed.

Phase 1 studies will be evaluated to determine the feasibility of the approach and determine whether to continue the option for Phase 2.

D. SPiNN Technical Description

SPiNN will explore efficient representations of AI based communications models for edge processing by systematically linking today's DSP solutions with innovative, adaptive neural network processing. Proposers should explore the following methodology to enhanced performance beyond today's DSP solution.

SPiNN AI Exploration Methodology:

Step 1: Train the discriminator in the NN kernel and verify accuracy with the physics based model. As seen in Figure 2, the first step of SPiNN will be to train the NN kernels to accurately represent specific physics based models. The trained NN discriminator should be verified against traditional computed DSP data or DNN result for accuracy as well as corner cases. Performers will verify that accuracy of the established NN kernels and the performance enhancement in computing latency, efficiency, workload, and resource requirement as compared with their physics-based mathematical approaches.

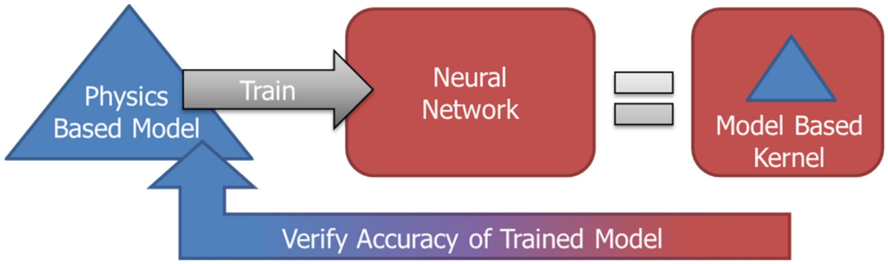

Figure 2: Training and verification of a physics based NN model

Step 2: With the discriminator in the NN kernel trained by a physics based model, an adaptive transformer will be incorporated to further enhance the signal processing kernel using the emerging generative-adversarial network (GAN) architecture. Working with the trained discriminator as an efficient primitive, the transformer layer will explore and learn complex features, such as temporal and nonlinear characteristics, provided by more complex training data and/or a generative data set. SPiNN kernels will reduce the dependence on large data sets with continual transformer learning to mitigate non-linear and non-stationary environments, to produce higher accuracy and more reliable results than conventional DSP or DNN approaches. For example, synthetic GAN data generative methods or variational auto-encoders (VAE) can enhance accuracy with the true data distribution of the system. SPiNN performers are encouraged to explore innovative neural network algorithms and apply this synthetic generation process for signal processing functions, such as Viterbi-Viterbi symbol decoder, multi-path equalizer, etc.

As seen in Figure 3a, once the discriminator is verified by the training data from physics-based model, performers will develop an additional transform layer to form an adaptive neural network kernel. Figure 3b shows a generic methodology for using copies of the trained model-based NN as both the generator and discriminator network to evolve the signal processing kernel to perform known functions with additional capabilities such as processing non-linear or non-Gaussian signal environments. *Each SPiNN performer is required to implement an end-to-end neural network with adaptive NN kernels for one of the DoD-relevant applications in communications or radar to achieve more than 10x improvements to input signal-to-noise-ratio sensitivity at an error rate of 10^{-3}, compared to SOA physics model-based approaches.*

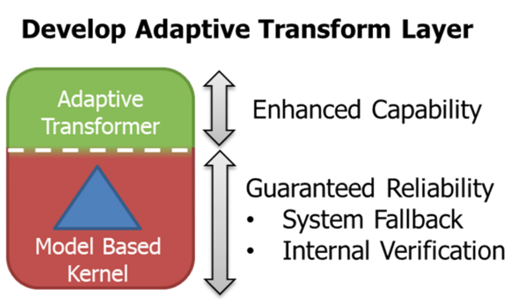

Figure 4a: Enhancement of model with adaptive transform layer

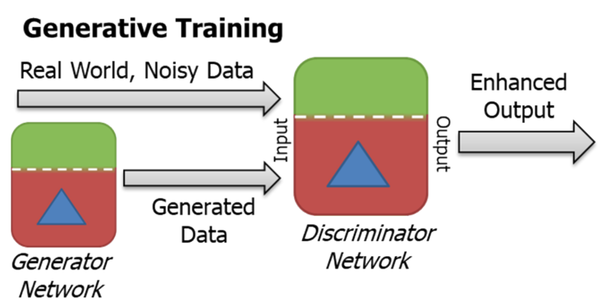

Figure 3b: Generative training with enhanced kernels.

SPiNN Program Overview:

Phase 1 Technical Approach: Proposers will develop and demonstrate signal processing kernels based on physics models and adapt the models using a transformer layer with generative training methods to produce enhanced adaptive NN kernels that have measured accuracy and reliability. Kernels will be developed for a set of basic signal processing functions for either a communications system or a radar system.

Proposers will:

1. Train the NN discriminator based on physics-based models from known DSP toolkits such as those built into Matlab. Verify accuracy, workload and latency as well as estimate required resources of the NN approach vs traditional physics-based mathematical approach.
2. Develop an adaptive transform architecture with the model based discriminator to form an adaptive NN kernel that can be used to train the network using generative-adversarial training methodology.

Key Demonstration: Demonstration of key metrics *in simulation* of a set of basic adaptive neural network kernels to perform signal processing functions for a communication or radar application.

Phase 2 Technical Approach: Proposers will develop more sophisticated signal processing kernels and implement them in FPGA emulated on cloud service to demonstrate an end-to-end signal processing function for either a communication or radar application.

Proposers will:

1. Implement an end–to-end demonstration of targeted SP application using adaptive NN kernels implemented in FPGA emulated on cloud service.
2. Demonstrate key metrics for both individual signal processing kernels as well as the complete end-to-end signal processing operation.

Key Demonstration: Demonstration and performance projections *in FPGA emulated on cloud service* to show the collaborative interaction of AI based adaptive NN kernels for enhanced end-to-end signal processing performance in a basic communication or radar system.

SPiNN Applications

Each proposal should propose **one** of the following two application areas, communications or radar, to best illustrate the measurable advantages and key goals of the SPiNN program. Only these two applications are of interest to the SPiNN program at this time and other applications will be considered non-responsive for the purposes of this solicitation. Specific information related to the two focus application areas are listed below.

For the chosen application, proposers should propose a plan to establish a dataset that does not contain personally identifiable information (PII) or classified data to validate the proposed solution. SPiNN is an unclassified 6.2 applied research effort and 6.2 fundamental research effort in the case of universities performing research on campus. Therefore, it is imperative that no PII or classified information is collected, communicated, or transferred in any way during the operation of the program.

As the SPiNN technical area involves cutting edge knowledge of emerging AI techniques as well as domain specific knowledge in the application areas, teaming is encouraged if it will lead to a stronger proposal to the solicitation.

Application 1: RF Communications

Proposer should select one specific application relevant to RF communications, such as mobile phones, IoT, point-to-point link, cognitive radios, etc., to implement and demonstrate the benefits of using innovative SPiNN techniques.

In Phase 1, proposers must develop four individual signal processing neural network elements (e.g. Modem, Channel equalization (EQ), De-Noise Filters, and Error Correction), that would form a basic communication link using the SPiNN methodology. Proposers will have the flexibility to implement associated signal processing kernels to optimize their application. These kernel can be simulated in Matlab or other tools that would enable the comparison of physics-based mathematical processing elements in a signal processing suite against the SPiNN adaptive model-based neural network kernels. The following four basic elements should be addressed in the proposal, although proposers may introduce alternative or new kernels for the element if they would enhance the functionality and performance for the selected applications:

RF Communication Elements

A. Modem
   - Phase 1 & 2: Modulation - 12-channel 16QAM OFDM or other more advanced waveform.
     - Examples of signal processing kernels:
       - Constant Modulus Algorithm
       - Viterbi-Viterbi Algorithm
       - Frequency-Phase recovery

B. Channel EQ
   - Phase 1: 1-channel point-point Rayleigh fading
   - Phase 2: Array equalizer
     - Examples of signal processing kernels:
       - Beam Forming Array
       - MIMO Array

C. De-Noise Filters – AWGN with In-Band Interferences
   - Phase 1: Linear filter: 16th order FIR & IIR
   - Phase 2: De-noise filters
     - Examples of signal processing kernels:
       - Matched filter
       - Kalman filter
       - Cyclostationary Detection filter

D. Error Correction
   - Phase 1: Reed-Solomon code (255, 239)
   - Phase 2: Advanced error correction
     - Examples of signal processing kernels:
       - Viterbi algorithm
       - LDPC

In Phase 2, proposers are required to implement a communication signal processing neural network suite for simulation and FPGA emulation on cloud service as shown in Figure 4 below.

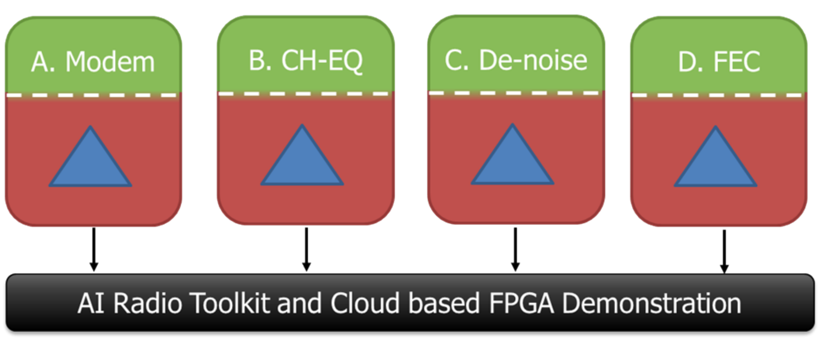

Figure 4: Notional communications signal processing neural network suite with required SP elements for Phase 2 demonstration in FPGA emulated on cloud service.

The goal of this implementation is to integrate various SPiNN NN kernels together for optimal end-to-end communication signal processing performance to achieve at least 10x improvement in the input signal-to-noise ratio at an error rate of 10^{-3} when compared to a traditional approach. The enhanced signal processing performance and new capabilities should be analyzed and illustrated.

Application 2: Radar Sensing

Proposer will select one specific application relevant to radar sensing, such as surveillance radar, moving target indicator radar, synthetic aperture radar, automotive radars, etc., to implement and demonstrate the benefits of using innovative SPiNN techniques.

In Phase 1, proposers must develop four individual radar signal processing neural network elements (e.g. Adaptive Waveform Generator & Space Time Adaptive Processor, Spectrum Management, Down Conversion, Filtering, Waveform Filtering, and Doppler & Clutter Processing, Range & Doppler gates) that would form a basic radar sensor using the SPiNN methodology. Proposers will have the flexibility to implement associated signal processing kernels to optimize their application. These codes can be simulated in Matlab or other useful tool that would enable the comparison of physics-based mathematical processing elements in the radar signal processing suite against the SPiNN adaptive model-based neural network kernels. The following basic elements should be addressed in the proposal, although proposers may make enhancements or add new elements if they feel that it would enhance the strength of their proposal:

Radar Processing Elements
For Phase 1 and 2: Monostatic radar

A. Adaptive Waveform Generator & Space Time Adaptive Processor (STAP)
    - Phase 1 Conventional Waveform Generator and STAP Beam
    - Phase 2 Adaptive waveform processing to enhance interference suppression with signal-to-interference ratio (SIR) > 10 dB enhancement
        - Examples:
            - Adaptive waveform and Beam Pattern management
            - Interference excision techniques

B. Spectrum Management, Down Conversion, Filtering
    - Phase 1 Conventional Spectrum & Instantaneous Bandwidth
    - Phase 2 Spectrum & filtering enhance SIR > 10 dB

C. Waveform Filtering
    - Phase 1 Conventional compression & matched filtering
    - Phase 2: Advanced filters
        - Examples:
            - Cyclostationary detection filter
            - Target specific optimized filters

D. Doppler & Clutter Processing, Range & Doppler gates
    - Phase 1 Conventional Doppler & Clutter filters & range gates
    - Phase 2 Tracking based optimizations

In Phase 2, proposers are required to implement a radar signal processing neural network suite for simulation and FPGA emulation on cloud service as shown in Figure 5 below.

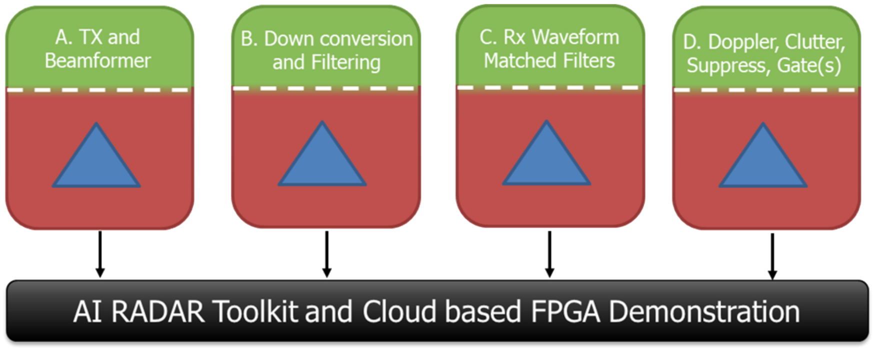

Figure 5: Notional pipeline for radar sensing neural network suite with required SP elements for Phase 2 demonstration in FPGA emulated on cloud service.

The goal of this implementation is to integrate various SPiNN adaptive NN kernels together for optimal end-to-end radar signal processing performance to achieve at least 10x improvement in the input signal-to-interference ratio at a target missed detection error rate of 10^{-3} when compared to a traditional approach. The enhanced signal processing performance and new capabilities should be analyzed and illustrated.

E. Schedule/Milestones

Proposers must address the following Research Project Objectives, metrics, and deliverables, along with fixed payable milestones in their proposals. The task structure must be consistent across the proposed schedule, Task Description Document (TDD), and the Volume 2 - Price Volume. Proposers must complete the "Schedule of Milestones and Payments" Excel Attachment provided with this AIE Opportunity as part of submitting a complete proposal and fulfilling the requirements under Vol. 2 Price Volume. If selected for award negotiation, the fixed payable milestones will be directly incorporated into Attachment 2 of the OT for Prototype agreement ("Schedule of Milestones and Payments"). Proposers are encouraged to use the TDD template provided with the Program Announcement DARPA-PA-19-03, which will be Attachment 1 of the OT agreement.

For planning and budgetary purposes, proposers should assume a program start date of April 1, 2020. Schedules will be synchronized across performers, as required, and monitored/revised as necessary throughout the effort. Proposals must include delivery schedules for Phase 1 and Phase 2 that include timelines for preliminary (to facilitate inspection by the Program Manager) and final (to facilitate evaluation) release of deliverables. Research objectives and fixed milestones for this program should include:

Phase 1 Research Objectives:

Develop the adaptive model-based neural network signal processing kernels to demonstrate the selected communications or radar application, following the methodologies and approaches outlined in Section C and D. This will implement the physics model-based neural network processing element, adaptive transform architecture, data representation, generative-adversarial training methodology, algorithms, compute primitives, neural network topology, dataset and simulation to show the initial accuracy and power projections for the proposed application.

The outcome of Phase 1 efforts will result in the development of the SPiNN methodology and validate the approach through development of advanced adaptive neural network kernels toward establishing integrated signal processing neural network suites for end-to-end neural network signal processing solutions in Phase 2.

Phase 1 Milestones:

- Milestone #1 (Month 2): Establish simulation platform including physics-based mathematical models, dataset, architecture, and algorithms for proposed neural network kernels. Provide a status update (milestone telecon and PowerPoint) on the updated approach.
- Milestone #2 (Month 5): Train NN discriminator kernels using the data from physics-based mathematical models and verify the accuracy of kernels with data not in the training set. Estimate the reduction in FPGA resources and evaluate improvement in computation efficiency and required resources of the model-based NN kernels versus the traditional physics-based mathematical approach. Provide interim written report on finding and simulation results. Participate in a milestone telecon with PowerPoint presentation.
- Milestone #3 (Month 8): Develop transformer layer for generative trained models. Training using generative-adversarial network methodology and the model-based discriminator network. Estimate the additional FPGA resources required to implement the transformer layer and generative network. Demonstrate the enhanced performance of the SPiNN enhanced kernels for SP against the Phase 1 metrics. Provide a final written report for Phase 1 detailing the technical findings as well as possible modifications to execute Phase 2 research. Participate in a milestone telecon with PowerPoint presentation.

Phase 2 Research Objectives:

Utilize adaptive physics based kernels from Phase 1 to produce more challenging kernels in Phase 2. Optimize the inputs and outputs of each of the kernels in the integrated suite to demonstrate a full signal processing neural network solution in simulation and to implement the solution in an FPGA emulated on cloud service.

The outcome of this Phase 2 integrated SPiNN signal processing suite and its end-to-end test bench is to demonstrate final program goals and metrics for the proposed applications. The milestones, deliverables and metrics are outlined below.

Phase 2 Milestones:

- Milestone #4 (Month 11): Establish physics model, dataset, and architecture for more complex physics model-based NN kernels. Provide a status update (milestone telecon and PowerPoint) on the updated approach.
- Milestone #5 (Month 14): Train advanced discriminator NN kernels using the data from physics-based mathematical models and verify the accuracy of kernels with untrained data. Estimate the reduction in FPGA resources and evaluate improvement in computation efficiency and required resources of the model-based NN kernels versus the traditional mathematical physics-based approach. Provide interim written report on finding and simulation results. Participate in a milestone telecon with PowerPoint presentation.
- Milestone #6 (Month 17): Integrate combined kernels into signal processing neural network suites to demonstrate optimal end-to-end integrated signal processing capabilities for communication or radar application. Perform generative training to optimize the end to end performance of the neural network suite in FPGA emulated on cloud service implementation. Include performance analysis of application-specific metrics for the integrated NN signal processing solution when compared to a traditional physics-based mathematical implementation. Provide the Phase 2 report documenting technical approach, physics models, transformer and kernel architectures, data sets, and simulated results. Provide a final report to summarize the technical approach, comparison of the performance and compute resources between the SPiNN adaptive kernels and traditional model-based approaches.

All proposals must include the following meetings and travel in the proposed schedule and cost:

- To foster collaboration between teams and disseminate SPiNN AIE opportunity developments, a two-day Principal Investigator (PI) meeting will be held approximately every six months after the kickoff meeting. For budgeting purposes, plan for four two-day meetings over the course of 18 months: two meetings in the Washington, D.C. area (Month 1 and Month 12) and two meetings in the San Francisco, CA area (Month 6 and Month 18).

- Regular teleconference meetings will be scheduled with the Government team for progress update, as well as for problem identification and mitigation. These are expected to occur within 2 weeks of each milestone due date and will be attended by the government Agreement Officer's Representative (AOR) as a requirement to approve milestone payment.
- Proposers should anticipate at least one site visit per phase by the DARPA Program Manager. At that time, performers will have the opportunity to update and demonstrate progress towards agreed-upon milestones.
- Proposers should plan travel for GOMACTech 2021 or a similar technical conference to participate in a joint SPiNN session detailing their non-proprietary results for interested government transition partners.
- Proposers may plan travel for one technical conference per year.

F. Metrics/Deliverables

The following metrics will be used to verify that the SPiNN methodology is successful in terms of producing a new approach to signal processing with AI.

SPiNN Program Metrics:

Phase 1:

- 10x reduction in computational effort (combined latency/efficiency) for inference to calculate physics based NN SP models in FPGA emulated on cloud service when compared with a traditional model based approach.
- Enhanced performance, accuracy and new capabilities of SPiNN neural network kernels when compared with a traditional physics-based mathematical approach.

Phase 2:

- 10x improvement in signal-to-noise ratio (SNR) and signal-to-interference ration (SIR) for a given communications or radar processing solution when compared against a traditional physics-based mathematical model implementation.
- 10x reduction in end-to-end computational effort (combined latency/efficiency) for a given communications or radar processing solution when compared against a traditional model based implementation.

The SPiNN program metrics must be achieved while maintaining certain standards for the chosen application as listed below:

Application 1: Communications Metrics

Function*	Phase 1	Phase 2
Modem	<10E^{-3} Output BER (@Eb/No =20dB)	<10E^{-3} Output BER (@Eb/No =12dB)
EQ Channel[1]	ΔEb/No = >3dB	ΔEb/No = >6dB
EQ Noise[2]	JSR = 15dB	JSR = 20dB
Error Correction[3]	Output BER < 10^{-5}	Output BER < 10^{-6}
FPGA Resources	-	5x less than model-based implementation

Table 1: Communications Metrics

** Metrics are based on LTE 12-channel 16QAM OFDM or other more advanced waveform.*

1. *EQ can enhance performance of a multipath Rician or Rayleigh channel by greater than 3db, 6 dB respectively*
2. *Enhance performance of an interference channel by the equivalent of reducing interference by 15, 20 dB respectively*
3. *FEC functions can enhance performance of a of a by reducing bit errors for an input BER of 10^{-3} to **less than** 10^{-5} and 10^{-6} respectively*

Application 2: Radar Metrics

Function	Phase 1	Phase 2
Adaptive Processing	STAP	Interference Suppression > 10 dB
Waveform Generator	Standard Chirp or LFMOP	Parametric Adaption SJR enhancement > 10 dB
Spectrum, Down Conversion & Filter	Adaptive Passband	Interference suppression > 10 dB
Matched Filtering	Optimized windows, sidelobe management	Context Adaptive Mismatched filtering [1]
Doppler & Clutter Filter	Detect Doppler Foldover	Joint optimization of waveform and Doppler processing enhanced Discrimination [1]
Range Gate & Doppler Gate	Awareness of range Doppler resolution status	Joint Parametric optimization with Waveform, Target ID & Track [1]

Table 2: Radar Metrics

** Metrics based on Monostatic Ground Moving Target Indicator (GMTI) radar or automotive radar*

1. *Vendor defined performance parameters*

Performers will be required to provide at a minimum the following deliverables:

Phase 1:

- Simulation of SPiNN signal processing kernels and code to demonstrate successful application of the SPiNN methodology to signal processing for the selected application.
- Final Phase 1 Report summarizing targeted application, technical approach, physics models, transformer and kernel architectures, data sets, and simulated results. Documentation of the performance comparison to alternative SOA methodology, quantification of accuracy, quantification of robustness.
- All other reports and data as required by the individual Phase 1 milestones.

Phase 2:

- Live end-of-phase demonstration of proposed communication or radar pipeline showing real-time processing of stored data on an end-to-end demonstration in FPGA emulated on cloud service.
- Report of FPGA code, resources, dataset, and COTs components (if applicable).

- Final Phase 2 report documenting technical approach, physics models, transformer and kernel architectures, data sets, and simulated results.
- Final report to summarize technical approach, accomplishments and the performance comparison to alternative SOA methodology, quantification of accuracy, quantification of robustness, computing resource and emulated utilization of FPGA.
- All other deliverables as required by the individual Phase 2 milestones such as registered reports on milestones, experimental protocols, publications, IP inventions, intermediate and final versions of hardware and hardware designs, software libraries, code, and APIs, including documentation and user manuals, and/or a comprehensive assemblage of design documents, models, modeling data and results, and model validation data.

Unless otherwise specified, all deliverables are expected to be released with unlimited rights.

II. Award Information

Selected proposals that are successfully negotiated will result in award of an OT for prototype project. See Section 3 of DARPA-PA-19-03 for information on awards that may result from proposals submitted in response to this notice.

Proposers must review the model OT for Prototype agreement provided as an attachment to DARPA-PA-19-03 prior to submitting a proposal. DARPA has provided the model OT in order to expedite the negotiation and award process, and ensure DARPA achieves the goal of AIE, which is to enable DARPA to initiate a new investment in less than 90 days from idea inception. The model OT is representative of the terms and conditions that DARPA intends to award for all AIE Awards. The task description document, schedule of milestones and payments, and data rights assertions requested under Volumes 1, 2, and 3 will be included as attachments to the OT agreement upon negotiation and award.

Proposers may suggest edits to the model OT for consideration by DARPA and provide a copy of the model OT with track changes as part of their proposal package. Suggested edits may not be accepted by DARPA. The Government reserves the right to remove a proposal from award consideration should the parties fail to reach agreement on OT award terms and conditions. If edits to the model OT are not provided as part of the proposal package, DARPA assumes that the proposer has reviewed and accepted the award terms and conditions to which they may have to adhere and the sample OT agreement provided as an attachment, indicating agreement (in principle) with the listed terms and conditions applicable to the specific award instrument.

In order to ensure that DARPA achieves the AIE goal of award within 90 days from the posting date **(January 2, 2020)** of this announcement, DARPA reserves the right to cease negotiations when an award is not executed by both parties (DARPA and the selected organization) on or before **March 31, 2020**.

III. Eligibility

See Section 4 of DARPA-PA-19-03 for information on who may be eligible to respond to this notice.

IV. AIE Opportunity Responses

A. Proposal Content and Format

All proposals submitted in response to this notice must comply with the content and format instructions in Section 5 of DARPA-PA-19-03. All proposals must use the templates provided as Attachments to the PA and the "Schedule of Milestones and Payments" Excel Attachment provided with this AIE Opportunity and follow the instructions therein.

Information not explicitly requested in DARPA-PA-19-03, its Attachments, or this notice may not be evaluated.

B. Proposal Submission Instructions

Responses to DARPA-PA-19-03-04 shall be submitted through electronic upload to DARPA's BAA Portal (https://baa.darpa.mil).

DARPA will acknowledge receipt of complete submissions via email and assign identifying numbers that should be used in all further correspondence regarding those submissions. If no confirmation is received within two business days, please contact
SPiNN@darpa.mil to verify receipt.

When planning a response to this AIE Opportunity, proposers should take into account the submission time zone and that some parts of the submission process may take from one business day to one month to complete (e.g., registering for a Data Universal Numbering System (DUNS) number or Tax Identification Number (TIN).

Electronic Upload

First time users of the DARPA BAA Portal must complete a two-step account creation process. The first step consists of registering for an extranet account by going to the URL listed above and selecting the "Account Request" link. Upon completion of the online form, proposers will receive two separate emails; one will contain a user name and the second will provide a temporary password. Once both emails have been received, the second step requires proposers to go back to the submission website and log in using that user name and password. After accessing the extranet, proposers may then create a user account for the DARPA Submission website by selecting the "Register your Organization" link at the top of the page. Once the user account is created, proposers will be able to see a list of solicitations open for submissions, view submission instructions, and upload/finalize their proposal.

Proposers who already have an account on the DARPA BAA Portal may simply log in at https://baa.darpa.mil, select this solicitation from the list of open DARPA solicitations and proceed with their proposal submission. Note: proposers who have created a DARPA
Submission website account to submit to another DARPA Technical Office's solicitations do not need to create a new account to submit to this solicitation.

All full proposals submitted electronically through the DARPA Submission website must meet the following requirements: (1) uploaded as a zip file (.zip or .zipx extension); (2) only contain the document(s) requested herein; (3) only contain unclassified information; and (4) must not exceed 100 MB in size. Only one zip file will be accepted per full proposal. Full proposals not uploaded as zip files will be rejected by DARPA. Technical support for the DARPA Submission website is available during regular business hours, Monday – Friday, 9:00 a.m. – 5:00 p.m. Requests for technical support must be emailed to BAAT_Support@darpa.mil with a copy to SPiNN@darpa.mil. Questions regarding submission contents, format, deadlines, etc. should be emailed to SPiNN@darpa.mil. Questions/requests for support sent to any other email address may result in delayed/no response.

Since proposers may encounter heavy traffic on the web server, DARPA discourages waiting until the day proposals are due to request an account and/or upload the submission. Note: Proposers submitting a proposal via the DARPA Submission site MUST (1) click the "Finalize" button in order for the submission to upload AND (2) do so with sufficient time for the upload to complete prior to the deadline. Failure to do so will result in a late submission.

C. Proposal Due Date and Time

Proposals in response to this notice are due no later than **at 4:00 PM on January 31, 2020**. Full proposal packages as described in Section 5 of DARPA-PA-19-03 must be submitted per the instructions outlined in this AIE Opportunity *and received by DARPA* no later than the above time and date. Proposals received after this time and date will not be reviewed.

Proposers are warned that the proposal deadline outlined herein is in Eastern Time and will be strictly enforced. When planning a response to this notice, proposers should take into account that some parts of the submission process may take from one business day to one month to complete.

V. Proposal Evaluation and Selection

Proposals will be evaluated and selected in accordance with Section 6 of DARPA-PA-19-03. Proposers will be notified of the results of this process as described in Section 7.1 of DARPA-PA-19-03.

VI. Administrative and National Policy Requirements

Section 7.2 of DARPA-PA-19-03 provides information on Administrative and National Policy Requirements that may be applicable for proposal submission as well as performance under an award.

VII. Point of Contact Information

Dr. Young-Kai Chen, Program Manager, DARPA/MTO, SPiNN@darpa.mil

VIII. Frequently Asked Questions (FAQs)

All technical, contractual, and administrative questions regarding this notice must be emailed to SPiNN@darpa.mil. Emails sent directly to the Program Manager or any other address may result in delayed or no response.

All questions must be in English and must include name, email address, and the telephone number of a point of contact. DARPA will attempt to answer questions publically in a timely manner; however, questions submitted within 7 days of the proposal due date and time listed herein may not be answered.

DARPA will post a SPiNN FAQ list under the AIE Opportunity on the DARPA/MTO Opportunities page at https://www.darpa.mil/work-with-us/opportunities?tFilter=&oFilter=3&_sort=name. The list will be updated on an ongoing basis until one week prior to the proposal due date. In addition to the FAQ specific to this notice, proposers should also review the Program Announcement for AIE General FAQ list on the DARPA/DSO Opportunities page under the Program Announcement for AIE (DARPA-PA-19-03).