Subject

Intel's Neuromorphic Computing Lab (NCL) seeks proposals for research projects that advance neuromorphic technology towards real-world state-of-the-art applications. We are especially interested in projects which use and improve the open-source Lava framework as well as neuromorphic algorithms and application demonstrations targeting Intel's Loihi 2 chip. We encourage researchers to publish openly and help expand the neuromorphic research community and accelerate the commercial adoption of neuromorphic technology.

Accepted proposals will receive access to Intel's cloud-based or on-premises neuromorphic systems depending on project requirements and can be considered for upcoming funding opportunities.

Table of Contents

Key Dates

Info Session and Feedback from Intel:

Principal Investigators may attend an Info Session about this Request for Proposals on 7 February 2023 at 8:00 – 9:00 AM PST.

INRC Spring Workshop:

Intel will host a Spring 2023 workshop featuring the latest research from Intel Labs and members of the INRC. The workshop will provide in-depth technical training on Intel Loihi and Lava. Visit neuromorphic.intel.com for more info and to plan your submission.

Proposal Submission Deadline:

Proposals for funding are due before 1 March 2023.

Proposals for hardware access only will be considered at any time.

How to Submit:

Proposals should use the latest INRC Proposal Template available on the Join the INRC webpage. Completed proposals should be emailed to inrc_project_proposals@intel.com.

If you are not a current member of the INRC, please apply to join the INRC by following the above link.

Contact for Questions:

For all inquiries related to this RFP or the INRC, please contact inrc_interest@intel.com.

Introduction to the INRC

In 2018, Intel's Neuromorphic Computing Lab launched the **Intel Neuromorphic Research Community** (INRC). This collaborative research program is open to all academic, government, and industry research groups interested in exploring neuromorphic architectures for mainstream computing applications. In support of INRC projects, Intel provides Loihi research hardware and software to INRC members to evaluate the capabilities and advantages of neuromorphic approaches in a rigorous manner with real-world measurements and demonstrations.

In support of INRC research, Intel offers remote login access to Loihi systems and software development tools. Intel also loans physical hardware systems to teams that require physical access for their proposed research.

Intel hosts regular workshops and meetings open to INRC project participants to share results, discuss challenges, and provide hands-on training. A monthly online forum features presentations from Intel and invited researchers to share progress and new developments. Semi-annual INRC workshops bring the community together to meet, share progress, and discuss future directions. The INRC Spring 2022 workshop attracted over 700 registered attendees.

Intel's publication, "Advancing Neuromorphic Computing With Loihi: A Survey of Results and Outlook," [1] summarizes the findings of the first three years of research with Loihi.

To encourage growth of the community and convergence at the software level, Intel launched the Lava software framework as an open-source project on GitHub with permissive licensing. Lava supports cross-platform execution on Loihi hardware and x86 CPUs, as well as model training on GPUs, and can be ported to other platforms. Lava builds on a foundation of channel-based event-driven parallel processing with the goal of supporting a wide range of neuromorphic programming paradigms spanning deep learning to online learning to dynamics-based computing. Lava is modular, extensible, and easy to use. Projects supported by this RFP are expected to use, and preferably contribute to, the Lava framework.

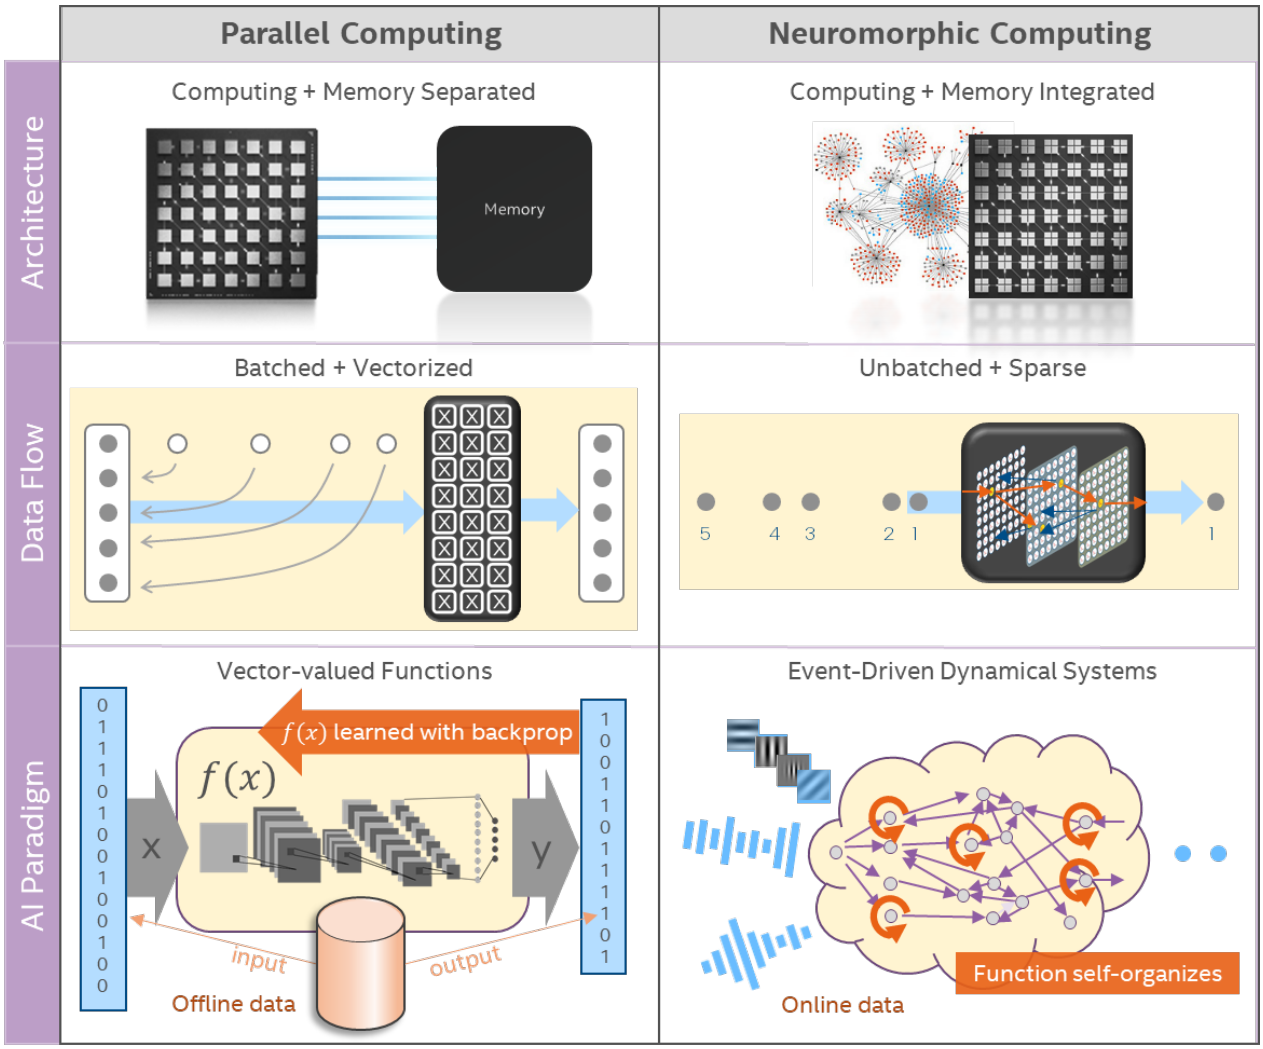

Figure 1. SIMPLIFIED VIEW CONTRASTING TODAY'S PARALLEL COMPUTING AND NEUROMORPHIC COMPUTING, WITH AI PARADIGMS CORRESPONDING TO EACH ONE.

Program Funding

To date, INRC projects have collectively received several million USD in funding from Intel's Corporate University Research Office (CUR). Corporate and government members of the INRC have funded over a dozen projects and regularly look to this community for partners, demonstrations, and promising emerging research directions.

Intel is now inviting proposals for 1- to 2-year projects that address at least one of the key INRC research vectors (RV1 - 5, see below) and align with Intel's strategic priorities in neuromorphic computing. Based on the available funding, Intel expects typical projects to fund one to two students or postdoctoral researchers. Proposals must justify the proposed budget in terms of the resources needed to carry out the proposed work.

Intel may share relevant submitted proposals with corporate and government members of the INRC interested in sponsoring research relevant to their application interests. Any group may opt out of this broader consideration by indicating so in their proposal.

Due to the limited grant funds available, we highly encourage researchers to leverage INRC support and membership to secure funding from other sources. Proposals outlining specific co-sponsorship opportunities will be considered favorably. Intel will provide letters of support for external funding applications given two weeks advance notice and a copy of the relevant proposal.

Background on Neuromorphic Computing

Neuromorphic computing aims to apply insights from neuroscience to create a new class of computing technology that follows the form and function of biological neural networks. The goal is to discover a computer architecture that is inherently suited for the kinds of intelligent information processing that living brains effortlessly support.

Interest in neuromorphic computing has intensified in recent years due to several developments.

First, the success of artificial neural networks in the form of deep learning inspires confidence that biological insights can lead to great practical gains in computing and AI. While the breakthroughs coming from the deep learning approach are impressive and of tremendous practical value, deep learning models are facing limits in application scope because of their large data, power, and latency requirements.

Second, the golden era of process scaling that provided conventional architectures with steady and massive gains in computing power has passed. While process scaling continues to shrink transistor sizes, conventional CPU and GPU architectures struggle to use ever-increasing transistor counts to deliver commensurate gains in application performance and energy efficiency. This motivates new architectural approaches that can deliver greater application-level performance using smaller but slower circuits.

Finally, the pace of progress in neuroscience has accelerated dramatically in recent years, providing a wealth of new understanding and insights about the functioning of brains at the neuron level.

Neuromorphic computing represents a fundamental re-thinking of computer architecture at the transistor level. Compared to conventional architectures, it is massively parallel, with the fundamental unit of computation being a neuron with time-dependent dynamics, compared to processors executing sequential instruction streams in conventional architectures. The computation in the brain and in most neuromorphic algorithms is an *emergent phenomenon*, the result of collective interactions between simple neural units. In contrast, computation in conventional CPU, GPU, and matrix arithmetic processors is a precisely sequenced procedure accessing state from a shared address space. Communication in a neuromorphic architecture occurs in a peer-to-peer multicast fashion, with

Box 1. Application properties necessary for realizing gains on neuromorphic architectures compared to conventional computing architectures.

- **Streaming input data** with temporal information structure such as audio, video, or any signals changing on microsecond-to-second time scales, especially when events of interest arrive infrequently and unpredictably.
- **A need for fast pattern matching, search, and optimization** supporting sub-symbolic processing, combinatorial recognition in high-dimensional spaces, and optimizing behavior through emergent neural dynamics.
- **A need for adaptation, fine-tuning, or associative learning** in response to changing sensory information or environmental conditions.
- **A need for low latency responses**, such as in closed-loop control applications, where the time and resource cost of *batching* and *vectorization* in conventional architectures is unacceptable.
- **Power constrained**, where often conventional architectures can achieve low latency at the expense of high power consumption. For suitable applications, neuromorphic architectures support both low latency and low power operation.
- **Small to medium AI models**. Compared to conventional computing systems, neuromorphic systems contain a relatively small amount of aggregate memory, the result of its compute/memory-integrated architecture. This makes Loihi systems generally unsuitable for architectures such as *Large Language Models* or very deep *ResNets*.

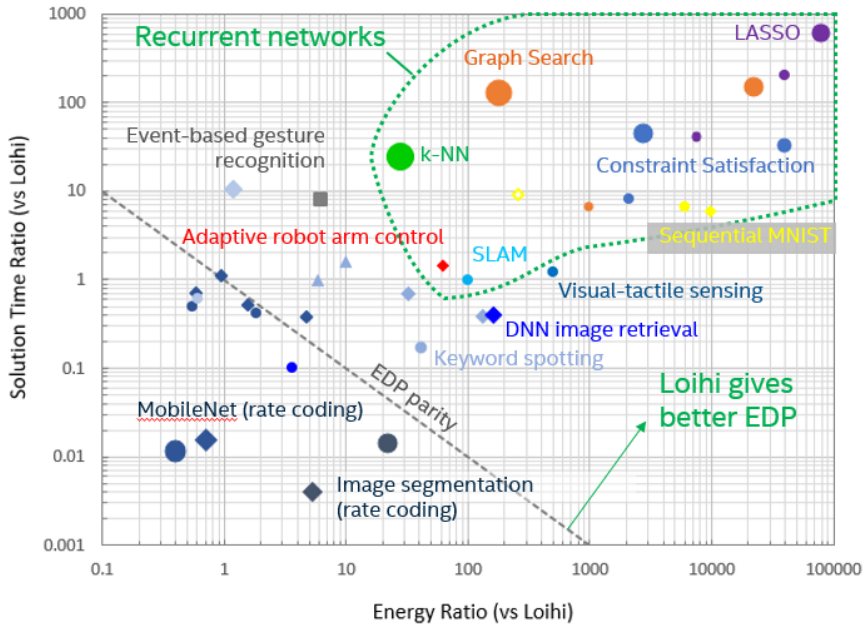

Figure 2. LOIHI RESULTS SHOWING RELATIVE GAINS IN SOLUTION ENERGY AND LATENCY VERSUS REFERENCE ARCHITECTURES. SEE [1] FOR DETAILS. RESULTS MAY VARY.

each neuron communicating scalar information to a diverse distribution of other neurons. Communication patterns are asynchronous, event-based, and extremely sparse in both time and space.

In a conventional processor, the top-level partitioning of the architecture between main memory and execution units leads to wide, vectorized datapaths that must stream data through the system at high bandwidth to achieve maximum efficiency. In a neuromorphic architecture, neural network weights and parameters are always stationary while only sparse data samples travel through the silicon. Data representations are low precision, often one bit, and all state changes, including weight changes, are the result of interactions between locally available quantities. Properties like noise, time-coding of information, and high-dimensional distributed data representations are used to achieve efficiency, robustness, and other surprising computational capabilities.

Figure 1 contrasts modern parallel computing and neuromorphic computing from an abstract architectural perspective. From this abstract view, ignoring for now any biological motivation, one can appreciate the bottom-up promise of the technology: low latency as a result of sparse, unbatched, and event-based data processing; resource-efficient processing of time-varying sensor input as a result of recurrent state updated locally per neuron, highly efficient online adaptation and learning as a result of fully localized state changes, and overall very low power as a result of its pervasive sparsity and activity-gating feedback paths. On the other hand, conventional parallel architectures supporting high precision matrix arithmetic are far better suited for offline training of differentiable and feed-forward models where sufficient pre-collected data is available.

As realized by chips such as the Intel Loihi 2 neuromorphic research processor , neuromorphic technology provides value for applications characterized by specific properties, shown in Box 1. These loosely correspond to the ecological needs that shaped brain evolution in nature. Rapid responses to sensory information allow mobile organisms to evade threats and capture prey, while fast learning allows organisms to respond to changing environmental conditions and outperform competition. Brain matter in nature is extremely expensive in both energy and material resources, just as we find in computing, so evolutionary pressures have led to designs that minimize resource consumption while maximizing behavioral objectives.

For these applications, Loihi has shown gains in latency and energy compared to conventional solutions into the orders of magnitude. These results are surveyed in [1] and summarized in Figure 2. Notable examples include constraint satisfaction, achieving up to 100,000x gains in energy-delay-product compared to conventional solutions.

Conversely, applications that do not exhibit the properties listed in Box 1 are unlikely to run better on neuromorphic architectures available during the time frame of research projects funded by this RFP.

While a considerable body of results now exists pointing to the advantages of neuromorphic technology, the algorithmic methods and programming tools needed to realize this value for real-world applications continue to limit progress. To enable commercially relevant applications and attract increased investment to the field, more attention should be directed to the most pressing of these near-term challenges. That is the objective of this RFP.

Technical Objectives of Research

Figure 3 defines the complete scope of neuromorphic computing research vectors. Although Intel pursues all these vectors, we prioritize support for external groups working in vectors one through five, with the greatest emphasis on vectors two through four.

Research on neuromorphic algorithms, applications, and programming models offer the most promising directions for real-world neuromorphic technologies and near-term commercial value, as detailed in the sections that follow. Most

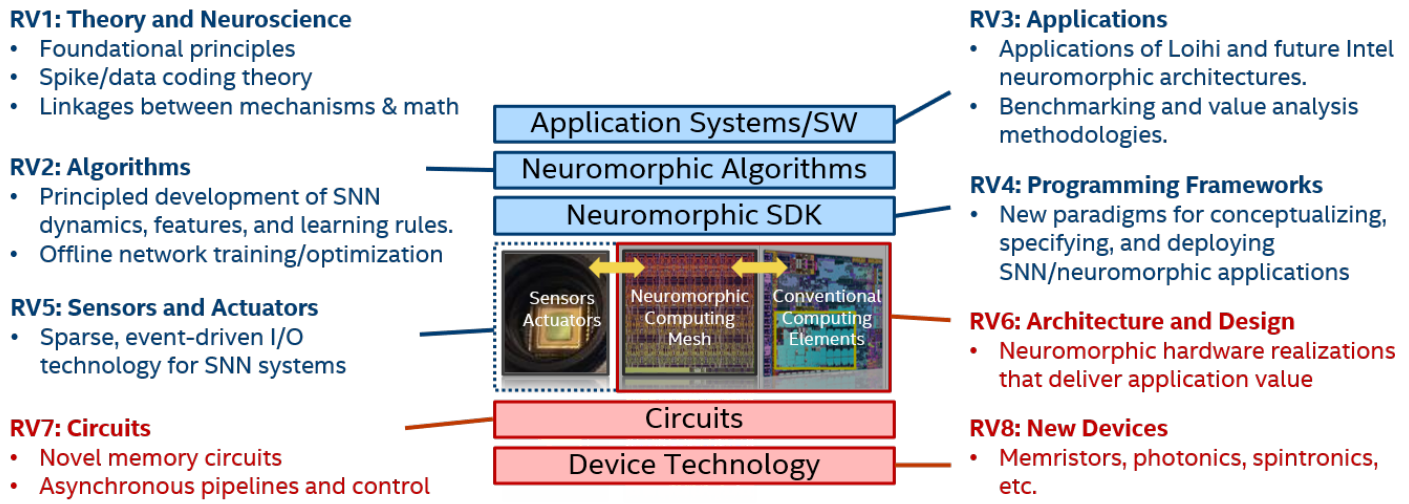

Figure 3. NEUROMORPHIC RESEARCH VECTORS. BLUE VECTORS FALL WITHIN THE DOMAIN OF INRC-FUNDED RESEARCH. RED ARE OUT OF SCOPE.

Intel support and resources will be directed to projects that align with these priorities, as they offer the greatest potential for success in the neuromorphic research field.

Proposals seeking funding in the INRC should use the Intel Loihi 2 hardware platforms and the Lava open-source framework. Proposals may take advantage of specific Loihi 2 features, as described in Box 2, to achieve the greatest benefits of the neuromorphic architecture. A general overview of Intel's first and second generation Loihi chips are available in [2] and [3], respectively. The use of common neuromorphic technologies will allow results to be rigorously compared across all relevant metrics: correctness, precision/accuracy, speed of execution, power consumption, and resource utilization. Technical objectives should be defined in terms of those metrics such that results can be quantitatively assessed.

Groups seeking funding must articulate an informed, forward-looking orientation. Proposals should reflect a strong understanding of results obtained to date by the INRC or with other neuromorphic platforms as documented by [1]. Projects should not assume software programming or training capabilities that do not yet exist, or else they should present a credible plan for developing those tools in Lava and making them available to other INRC members.

Furthermore, projects should accommodate the evolution of neuromorphic hardware architectures. Current Loihi and Loihi 2 chips offer a very flexible foundation for neuromorphic research, but changes in future systems and software are likely and project plans should be reasonably able to accommodate these.

Our broader objective is to accelerate progress in a collaborative fashion, so we wish to see functional, open-source code contributed to Lava (See Box 3) to allow others to replicate and advance on progress.

RV1: Theory and Neuroscience

The long-term success of neuromorphic technology depends on sound theoretical foundations that support robust algorithms and applications. As a new computational paradigm differing from conventional computer architectures in fundamental ways, neuromorphic computing currently lacks unifying theoretical frameworks, such as the Turing Model, which leads to a highly fragmented exploration space.

Past INRC RFPs have supported theoretical work in the areas of computational complexity frameworks, characterization of neural dynamics, and unified engineering figures of merit. While further work in these areas could be important for progress in the field, below are a few example areas where neuromorphic theory is currently lacking:

- Theoretical analysis of different information coding strategies based on application objectives. For example, when to use temporal coding, phase coding, rate coding, population coding, plus generalizations to other coding strategies such as sigma-delta with graded spikes.
- Holistic models of fault tolerance and recoverability for SNNs, including the effects of noise and stochastic synapses.
- Characterizing the computational tradeoffs and complexity analysis associated with promising neuromorphic frameworks such as vector symbolic architectures [4] and dynamic neural fields [5].
- Theory-driven investigation of the "accuracy gap" for deep SNNs trained with surrogate gradient backpropagation.

RV2: Algorithms

Central to the advancement of neuromorphic computing is the development of algorithms that leverage the novel features of neuromorphic architectures and satisfy their hardware constraints. These are algorithms that utilize *sparsity* of connectivity, communication, and activity. They should include *dynamically evolving state* within each neuron that is excited by inputs and inhibited through feedback loops. Learning algorithms have access to local state variables at the neuron and the synapse as well as synaptic traces for encoding granular reward, error, or neuromodulatory signals. Novel neuro-inspired features such as stochasticity, structural plasticity, event-driven computation, and temporal information coding can also provide unique advantages on neuromorphic architectures.

Box 2. Features of Intel's Loihi 2 neuromorphic hardware architecture.

- **Programmable neuron models**. Intel's first-generation Loihi processor supported a leaky-integrate-and-fire (LIF) spiking neuron model with the ability to aggregate neural units into multi-compartment dendritic trees to communicate graded neuron state between compartments. Intel's Loihi 2 neuromorphic processor generalizes this capability, with fully programmable neuron models that support a broad range of differential equations for state variable dynamics, configurable spike conditions, and state machines. Supported neuron models include adaptive threshold LIF, Resonate-and-Fire, Hopf resonators, sigma-delta coding, and many others. See for examples.
- **Graded spikes**. Many SNN chips, including Loihi 1, can only send binary-valued spike messages between neurons. While binary spikes can perform a remarkable amount of computation -- as best demonstrated by the brain -- in digital hardware, spikes can be generalized to carry integer-valued payloads with little cost in performance or energy. Loihi 2 supports such graded spikes, enabling more complex event-based messaging while preserving the sparse and time-coded communication of SNNs and providing greater numerical precision.
- **Three-factor learning rules**. Loihi 1 primarily supported two-factor learning rules (involving pre- and post-synaptic traces), with a third modulatory term set in a diffuse manner from graded "reward" broadcasts. Loihi 2 leads the next generation of neuromorphic chips supporting more targeted and localized third factors in learning rules. As an example, these may be error signals mapped to specific neurons, available as third factors in synaptic learning rules.

Algorithms research proposals should fully consider all recent learnings [1] and should include a plan for rigorous benchmarking to current state-of-the-art conventional solutions. The value of the proposed algorithms should be motivated in the context of a specific application and associated real world constraints, informed by the challenges and opportunities facing neuromorphic technology deployment.

RV2 projects should advance beyond theory and modeled examples to provide generally usable software modules in Lava targeting Loihi and other future neuromorphic platforms. Others should be able to easily apply results to their own problems, preferably over a range of different application domains.

The following areas are of particular interest.

Novel neuromorphic neuron models

Compared to the stateless neuron models of deep learning (e.g., ReLU), biologically inspired neuron models include *time-varying state variables*. In neuromorphic hardware, these neurons offer several computational advantages. They introduce time-varying behavior into a network, allowing the network to efficiently encode time-varying input signals, to make predictions, and to produce complex output sequences.

Loihi 2 implements neurons using programmable neuron cores that support a broad class of neuron models (see Box 2). These models can have nearly arbitrary internal dynamics and support both spike-based communication and continuous transmission of graded information. The former sparsifies long-range communication with event-based messages triggered by some spike condition; the latter provides high-precision computation within a local cluster of neurons, where the cost of communication and fanout are low.

Generally, we see significant promise of this algorithmic approach for demonstrating more compact, intelligent, and efficient nonlinear signal processing solutions, e.g., for audio and radio frequency processing.

In addition, significant potential has been demonstrated by networks composed of stochastic neurons. These models have been used to represent Bayesian or more general graphical models [6] and can solve a range of hard problems such as the computation of marginal probabilities, or maximum likelihood [7]. These models have for instance been applied to solving constraint satisfaction problems with SNNs [8] and could offer one possible realization of causal graphical models for efficient continual learning. Such stochastic spiking networks are now demonstrating significant outperformance compared to classical approaches on conventional hardware architectures [1].

Opportunities for progress could include advancing stochastic SNN theory, combining stochastic networks with other components such as offline-trained DNNs, and extending stochastic networks to areas such as probabilistic inference and model predictive control.

Computing with neurodynamics

To date, the best results using Loihi have been achieved in networks that solve problems that identify optimization and search solutions using network-level spiking neuron dynamics. In these networks, spikes have the effect of dimension-wise prioritizing asynchronous gradient descent steps, leading to orders-of-magnitude gains in speed and efficiency compared to traditional solutions for same problems. Examples include Lasso regression, constraint satisfaction, quadratic unconstrained binary optimization, shortest-path graph search, and similarity search, with other promising applications on the horizon to problems such as quadratic programming, integer linear programming, Subspace Locally Competitive Algorithm [9], Minimax optimization [10], and probabilistic inference [6] [7] [12] [13].

One important open challenge involves hierarchically composing such optimizing networks to solve larger problems or to solve the problems with higher precision than offered by the state variables of single neurons. Conventional multilevel solver techniques (e.g., [14]) may be of value here, in addition to insights from neuroscience.

Online learning

In the neuromorphic research field, much attention has been directed to learning algorithms that approximate backpropagation (or gradient descent in parameter space) with online learning rules that respect locality and other neuromorphic architectural constraints . Despite encouraging progress, major hurdles remain to be resolved before these approaches can yield practical value. While these algorithms can operate continuously, they rely on unrealistic assumptions about the statistics of real-world data, such as independent identically distributed (iid) samples or highly controlled training scenarios. Furthermore, operating online doesn't improve the data efficiency of backprop, which is a fundamental challenge for many edge applications.

Several example approaches of interest are described include on-chip few-shot transfer learning , surprise-driven learning such as CLAPP or Contrastive Predictive Coding , and stochastic learning with bounded resources such as complex synapses to extend memory lifetime, or *replay* for pinning significant memories.

Vector Symbolic Architectures

Vector Symbolic Architectures, also known as Hyperdimensional Computing algorithms, have shown great promise for neuromorphic systems with algorithms such as the resonator network . Recent applications to scene understanding and visual odometry have highlighted the potential for achieving state-of-the-art results with far more compact and explainable networks than those trained with deep learning. However VSA demonstrations to date tackle small-scale problems with dense vector encodings that do not easily scale to hyperdimensional sizes.

To realize the potential of VSAs as powerful symbolic framework for programming powerful algorithms on neuromorphic hardware, Intel Labs is developing a lava-vsa library that will support a variety of different vector encodings, including those exploiting sparsity for efficient and scalable mapping to spiking hardware. We are interested in funding credible proposals for productive collaboration in this domain.

Offline training and network optimization

A number of backprop-style offline training tools for spiking neural networks have emerged over the past few years. However, in most cases, technical limitations have limited the networks trained with these tools to relatively shallow and simple architectures, and small dataset sizes, compared to modern deep learning. Model training convergence and performance both suffer significant disadvantages that limit the broader applicability of these methods. Innovations in offline optimization and significant improvements to training algorithms are needed to reach performance that is competitive with standard ANNs. Promising approaches include Hessian-based methods, spike time-based gradients, use of novel neuron models and features to ease training, among others.

Meanwhile, we see strong evidence of a barrier to the performance of deep SNNs with backprop-based offline training. Spike-based neural network models are fundamentally not differentiable, so the surrogate gradient methods deployed in lava-dl and other SNN deep learning toolchains fundamentally involve approximations that introduce errors compared to training conventional ANNs. To overcome this, we are very interested in supporting new ideas and approaches to the problem of offline supervised training, including the Forward-Forward algorithm [28], evolutionary methods, and neural architecture search.

Neuromorphic transformers

In recent years, the advent of transformers has led to breakthroughs in many areas of artificial intelligence, from natural language processing to computer vision and multi-modal understanding. For sequential data processing applications, transformers have surpassed the performance of recurrent neural network models, especially in their ability to scale and train successfully with very large data sets.

Today's transformers are ill-suited for directly mapping to neuromorphic architectures. As feed-forward models operating on vectorized sequence data, they are designed to run best on conventional data parallel architectures with batched input and vast amounts of memory. However, at the heart of transformers are constructs that evoke neuro-inspired mechanisms: an attention mechanism closely related to associative memories and Hopfield networks [29], and positional encoders compatible with resonate-and-fire spectral transforms [30]. These properties suggest that neuro-inspired transformer models may be developed that operate on input data temporally, projecting recurrent state to attractor-based memory mechanisms achieving state-of-the-art speed and efficiency by computing with neurodynamics.

Development of such transformer-inspired neuromorphic models that are of great interest to Intel for funding, taking recent efforts [31] [32] into consideration.

RV3: System Applications

We seek real-world application demonstrations at the intersection of research and today's best engineering solutions. Over the long term we see a vast domain of applications for neuromorphic devices, but Intel is primarily focused on supporting applications that are commercially relevant and viable using today's neuromorphic hardware and algorithms.

To substantiate commercial relevance, we encourage application projects that include participation, support, or co-investment from industry or government organizations. For example, this could be a corporate advisor, use of data from an end customer, or integration into a commercial system platform. For highly

Box 3. Open-Source Lava Framework for Neuromorphic Computing

Lava is an open-source software framework to develop applications for neuromorphic hardware architectures. It provides developers with the abstractions and tools to develop distributed and massively parallel applications. These applications can be deployed to heterogeneous system architectures containing conventional processors as well as neuromorphic chips that exploit event-based message passing for communication. The Lava framework comprises high-level libraries for deep learning, constrained optimization, and others for productive algorithm development. It also includes tools to map those algorithms to different types of hardware architectures.

- **Composable Processes with Efficient Event-Based Message Passing**. Lava applications consist of a connected graph of processes mapped to a heterogenous execution platform including both conventional and neuromorphic components. Messages between Lava processes vary from asynchronous single-bit spikes to buffered packets with arbitrary payloads.
- **Extensible Libraries through Multi-Backend and Multi-Abstracton Compilers.** Lava supports a wide range of application programming paradigms such as offline deep learning , attractor networks and hyperdimensional computing, and highly efficient mathematical optimization.
- **Easy to use and extend Python code**. For the broadest possible adoption among neuromorphic developers, all libraries and features in Lava are built in Python and flexibly interface with external tool such as ROS, PyTorch, or Brian. Lava is free and open-source on GitHub.

The *Lava Extension for Loihi* (lava-loihi) is a proprietary package developed by Intel which offers high-performance Lava process execution on Loihi 2 systems. The extension is offered free of charge to all INRC members, and contributions to the extension are encouraged for research teams with sufficient neuromorphic software experience.

Visit https://github.com/lava-nc for more information and to get started.

compelling proposals, Intel can partner and co-develop technical assets to achieve a successful outcome.

Application proposals should clearly describe the value of the neuromorphic solution in relation to limitations of current solutions and how specifically neuromorphic technology will be used. Prior evidence that one or more project members possess significant technical expertise in the application domain is expected for credible proposals in this RV. Proposals should articulate the commercial impact of a successful outcome, based on measurable and significant advances the project will demonstrate over the current state-of-the-art.

Promising categories of System Applications research include, but are not limited to:

Audio processing, especially applications that need to operate continuously in an always-on fashion at low power levels and where a fast response and online adaptation is needed, e.g. wake-on-voice, dynamic noise suppression, automatic speech recognition, sound detection and localization, speaker identification, and blind source separation.

Signal processing for security, failure detection, and sensor networks. Examples range from radar, sonar, biometric, and turbine monitoring to cybersecurity intrusion detection to sensor network processing for earthquake prediction and oil field analysis.

Video processing, especially detection and classification of events and activities with a dynamic signature within the input video, such as activity recognition, target tracking, or novel object acquisition.

Human-machine interfacing: gesture recognition for cursor control or sign language interpretation, gaze tracking, speech processing, tactile/haptic sensor processing. Brain-computer interfaces (EEG, EMG, direct nerve/neural probes) that demonstrate real-world advantages for gaming and people with disabilities. Additional value may come from applying neuromorphic compute to wireless interfaces in this domain.

Routing and scheduling, such as NP-Hard vehicle routing problems for logistics, task and job scheduling in data centers, and traffic routing for networks. These types of problems effectively leverage neuromorphic optimization capabilities.

Real-world robot deployments, including perception and control of delivery robots, warehouse robotic systems, healthcare robots, and collaborative robot applications.

Aerospace and satellite-based applications, including radio-frequency signal processing and telecommunications, trajectory control, low SWaP multispectral image processing.

RV4: Programming Models and Frameworks

The maturity of software in the neuromorphic field remains low, and this creates a formidable barrier to industry adoption of neuromorphic technology. Code sharing between groups is minimal, and published examples generally are difficult if not impossible to replicate by others. There are very few examples of composability, abstraction, and modularity in the algorithms studied and published. While some promising frameworks have open-source code, prohibitive licensing terms limit widespread adoption and community-wide contribution. Intel launched the

open-source Lava framework as a solution (See Box 3 "What is Lava?").

With this RFP we hope to encourage developers with diverse backgrounds and interests to improve and extend the Lava framework. By bringing new ideas and perspectives to the software challenges, we see opportunities for great gains in areas that are bottlenecks to progress today: developer productivity, training efficiency, system composability, and libraries for powerful features like structural plasticity and evolutionary optimization.

Beyond the immediate priority of building out and optimizing the central capabilities of the Lava framework, we see several long-term compelling directions for Lava development. A few examples are listed below. We welcome the research community to take the lead in these areas with support from Intel.

Development of Domain-Specific Languages (DSL) spanning the levels of the neuromorphic computing stack. The goal might be to unify the various levels of abstraction, providing developers a consistent experience when working from low-level hardware configuration up to behavioral specification, or it may take a narrower focus on the end-user/application layer and maximize the efficiency of specifying useful applications.

Porting Lava to other neuromorphic architectures. The current Lava compiler and runtime support execution across several architectures, including Intel CPUs, RISC-V processors, and Loihi neuro cores. To encourage convergence on a common framework, transparent benchmarking, and the greatest possible breadth, Intel welcomes efforts to extend Lava's support to other architectures, both neuromorphic and conventional. Support for GPU and FPGA backends is particularly relevant for this RFP, e.g. by implementing support for OpenCL.

Developer Productivity and Experience Improvements including developer tools, graphical user interfaces, performance profilers, and interfaces to established packages. This type of project should provide value to the existing community and/or extend the usefulness of Lava and Loihi to new user communities. In addition, projects which make significant contributions to improving the neuromorphic developer experience by investigating and solving for limitations of the Lava software platform are encouraged.

RV5: Event-Based Interfaces

Over the past several years, event-based vision sensing technology has seemingly matured with the advent of commercially available sensors and large investments from numerous industry and government organizations. Neuromorphic processing of event-based sensor output promises many advantages over conventional architectures, yet algorithmic and hardware scaling challenges limit the near-term commercial viability of this combination of technologies. We view some of these challenges as fundamental, exacerbated by the pixel-level granularity of features produced by today's event-based sensors.

We are interested in supporting fundamental research that addresses these pain points: (1) application-driven modeling of future sensor architectures that tightly integrate novel photodiode sampling circuits with neuromorphic processing, both near and far; (2) novel spatiotemporal filtering techniques prototyped on Loihi and/or FPGAs that extract meaningful features with a minimum of parameters and compute cost; (3) feedback-driven attention and active sensing mechanisms that improve the speed, efficiency, and resource needs of visual inference and learning.

Beyond vision sensors, Intel may consider funding and offering in-kind support for interface and hardware engineering projects demonstrating the value of novel event-based sensor and actuator technologies, such as electronic skins, cochlea-inspired audio processing, muscle-like actuators, and wireless interfaces.

Research Proposal Process

The process for responding to the RFP and joining the Intel Neuromorphic Research Community is outlined below.

Note: only academic groups will be considered for funding under this RFP. However, commercial and government organizations can apply to Join the INRC and use Loihi hardware.

1) Submit your proposal online using the application here Join the INRC. To be considered for funding, you must select the "Research (PI)" membership category.
2) Intel will review submitted proposals and contact you with any requests for additional information, and to share our funding determinations.
3) If your proposal is accepted, an INRC Participation Agreement must be executed between Intel and all institutions participating in the proposed research. See the section below for more info on this agreement. Once this is executed, you will be granted:
    - Remote access to Loihi via the Neuromorphic Research Cloud (NRC).
    - Access to the *Lava extension for Loihi*, Intel's proprietary plugin for running Lava models on Loihi hardware platforms.
    - Physical access to loaned Loihi hardware systems **as needed** for the research project.
4) Conduct your research and participate in the community.
5) Share your progress, results, and feedback in INRC events.

Further guidance will be provided upon receipt of a project proposal. Feel free to send questions at any time to inrc_interest@intel.com.

Eligibility

We welcome groups of all types and locations to submit research proposals and engage in the community, subject only to U.S. export control laws.

Intel research grants are only offered to academic research groups.

Proposal Submission

Along with this RFP, Intel provides an *INRC Project Proposal Template* that should be filled in with the information and documentation required from respondents seeking funding. Refer to that document for detailed guidance on what information to include in your INRC project proposal, while preserving the template structure. A proposal may be rejected if it does not include the required information and documents.

For projects not seeking funding, respondents may submit project plans in any format as long as all required information is covered.

Some recommendations:

- Delete all commentary and guidance text from the template document.
- Strive for brevity.
- Feel free to submit more than one proposal.
- Keep the scope of each project narrowly defined and limited to a single research vector.

Please note that we are unable to receive proposals that are provided under an obligation of confidentiality. Proposals should therefore include only public information. If you represent a corporate entity with proprietary IP considerations, please contact us prior to submitting a proposal.

INRC Participation Agreement

To provide INRC members access to pre-production neuromorphic research hardware, Intel requires each organization participating in a proposal to execute a legal agreement covering the technology license, research results, confidentiality, and liability. Below is a plain-language summary of the intention of this legal document. The agreement:

1. Is executed between Intel and the research member's organization, not the individual. It will automatically cover new project members and can be updated to cover new projects. It must be signed by *an authorized legal representative*. This is often not the project PI.
2. Provides a license to use Intel's pre-production hardware and associated software ("Lava extension for Loihi") for the proposed research.

   Note that all general Lava software on GitHub is provided under BSD-3 and LGPL-2.1 licensing, so you are free to use that code without a signed participation agreement.
3. Provides Intel a license to use project results. This does not transfer ownership or restrict your ability to use the research, except that you agree to report back to Intel on the research and allow Intel to use those reported results.

   For example, in the course of your project, you might measure the performance of the Loihi chip while running a model for sorting jelly beans. Intel asks that you report these results, and Intel may use them to promote Loihi or improve future devices.
4. Requires researchers to provide Intel an opportunity to review publications prior to release and request the removal of any Intel Confidential information.
5. Covers Intel and the research member under a simple non-disclosure agreement. This allows either party to share confidential information with written notice and ensure the other party won't share that information further. This helps Intel provide early access to technology roadmaps and benchmarks and helps researchers to get feedback on preliminary results or project ideas.
6. Confirms the rights of members and Intel to publish research findings and sets out a simple process to avoid accidentally publishing confidential information.
7. Includes terms and conditions for remote access to Intel's neuromorphic research cloud, such as account security and tech support access.

Intel appreciates the diverse nature of member organizations and has designed the participation agreement to be consistent with industry standards and very permissive, while still protecting both the member and Intel. Modifications to the agreement are generally not permitted.

Progress Updates and Results

Once your INRC research project is approved and underway, Intel collects quarterly updates on your progress and results. These updates enable us to promote the research of the community and share exciting results with the press and analysts to encourage greater external support for your work. Intel also invites members to present their results at regular INRC Forums and semi-annual INRC workshops.

Intel supports and encourages publishing results in public, peer-reviewed forums. We will do our best to support live demonstrations and independently hosted hands-on workshops using Loihi hardware systems, subject to U.S. export control, security, availability, and other constraints.

Loihi Hardware Access

Intel provides two forms of access to Loihi-based neuromorphic systems and software to use in your research:

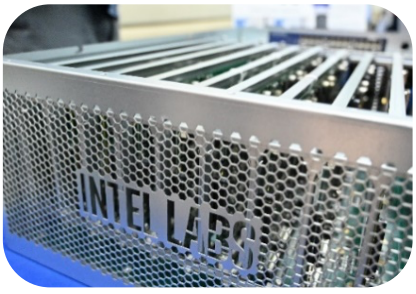

Neuromorphic Research Cloud

Most members of the INRC are provided remote access to a pool of Loihi systems through the Neuromorphic Research Cloud. Research members can remotely login to virtual machines to program and submit jobs to the cloud systems. Members have access to Loihi and Loihi 2 systems, including 8-chip Kapoho Point systems, 32-chip Nahuku systems, and the large-scale Pohoiki Springs system supporting up to 100 million neurons.

The neuromorphic research cloud offers the best platform for prototyping, controlled benchmarking, and developing neuromorphic systems. Most jobs run immediately, hardware is

configured and maintained by the Intel team, and the software and environment are easy to set up.

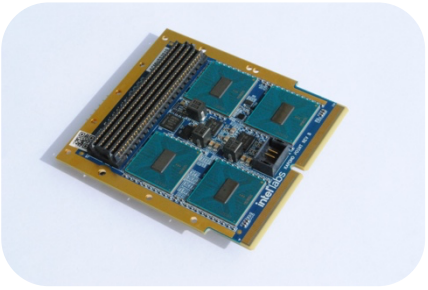

Kapoho Point
Intel can provide some research groups a Kapoho Point system for on-site use. This system includes 8 Loihi 2 chips. Due to the small form-factor and low power consumption, it can be embedded into edge devices and connected directly to sensors or robots and supports a wide variety of research use cases.

Kapoho Point is best suited for groups that have neuromorphic research experience and have developed applications on the Neuromorphic Research Cloud.

Evaluation Criteria for Proposals

In order of importance, the evaluation criteria for this solicitation are as follows:

1. **Potential contribution and relevance to Intel and the broader industry**: The proposed research should directly support a technology solution that addresses the RVs outlined above, leading to technological advances with the potential for ongoing technology transfer in collaboration with Intel and the broader industry.

2. **Technical innovation**: Proposed solutions of interest should clearly push the boundaries of technical innovation and advancement. Research that is not of interest in this program include incremental advancements to state-of-the-art and current design practices. Feasibility of new algorithms/techniques should be demonstrated through SW/HW implementations. Projects seeking funding should target Loihi hardware platforms and the Lava software framework to enable algorithmic capabilities and application proof of concept demonstrators that others can build on. Technical objectives should be defined in terms of quantitative target metrics (precision/accuracy, speed of execution, power consumption, and resource utilization) as detailed in "Technical Objectives of Research." Funded projects will be enabled with remote access to Loihi and future neuromorphic hardware platforms via our Neuromorphic Research Cloud (NRC) system and limited access to physical loaned systems as needed/approved. See sections on Engagement Process and Loihi Hardware and SDK for more information.

3. **Clarity of overall objectives, intermediate milestones and success criteria**: The proposed Research Plan should clearly convey that the PIs have the knowledge and capability to achieve the stated research goals. It is understood that any research program will have uncertainties and unanswered questions at the proposal stage, but a clear path forward in key challenge areas must be identified and justified. Teams are expected to demonstrate progress toward project goals at quarterly milestones and monthly project status updates. The proposal should explicitly point out which RV is being addressed, the synergy among them if more than one RV, the plan and milestones towards building research prototypes, plan for ongoing technology transfers, and the anticipated proof of concept outcome. The technical suitability of proposals to RV2: Algorithms and RV3: Systems Applications will be evaluated according to the criteria included in the INRC project proposal template, as included for reference in Appendix 1 and Appendix 2, respectively. Strength of project management will also be considered.

4. **Qualification of participating researchers**: The extent to which expertise and prior experience bear on the problem at hand. Please elaborate on track records of building research prototypes (e.g., open-source research code and collateral on GitHub) and resulting publications from past relevant projects.

5. **Cost effectiveness and cost realism**: The extent to which the proposed work is both feasible and impactful within the proposed resource levels will be examined.

6. **Potential for co-funding**: Opportunity for matching grants and co-funding with other funding entities, such as SRC, NSF, DARPA, NSERC, etc. will be given significant consideration.

7. **Potential for broader impact**: Intel supports the advancement of computing education and diverse participation in STEM. Significant consideration will be given to proposals in which the outcome of the research can influence the development of new curriculum initiatives impacting undergraduate or graduate education at the respective universities (e.g., exposure to latest industry technologies or tools in a classroom). Proposals are encouraged to elaborate on how the proposed work is anticipated to impact student education on campus and/or the broader academic community.

Intel Note:

As an industry leader, Intel pushes the boundaries of technology to make amazing experiences possible for every person on earth. From powering the latest devices and the cloud you depend on to driving policy, diversity, sustainability, and education, we create value for our stockholders, customers, and society. Intel expects suppliers in our supply chain to be strong partners in making Intel successful through support of Intel's goals and commitments to diversity, sustainability, and education.

In light of Intel's strong commitment to diversity and creating an inclusive environment, in your proposal please address: (a) your organization's commitment to diversity and inclusion with respect to race, national origin, gender, veterans, individuals with diverse abilities and LGBTQ, (b) a summary of your performance in this area and any initiatives you are pursuing, and (c) the diverse team you propose for this project, including leadership, support, and any subcontracting you propose (such as to minority- or women-owned businesses).

Intellectual Property

This solicitation affords proposers the option of submitting proposals for the award of a grant or gift, a sponsored research agreement, or other agreement as appropriate. Intel reserves the right to negotiate the final choice of agreement.

The final award terms are expected to follow one or the other of two high-level intellectual property (IP) approaches. Either: (1) Intel and the university will jointly agree that IP developed under a grant or gift will be placed in the public domain, including offering software under an open-source license, or (2) Intel and the university will negotiate a sponsored research agreement with more specific IP terms, which, at a minimum, will require the university to grant Intel and other sponsors (if any) a broad non-exclusive royalty free license to foreground IP.

It is a requirement to follow approach (1) if a project's software development directly enhances the Lava software framework or builds on pre-existing INRC shared software libraries.

Please note that Intel is unable to receive proposals under an obligation of confidentiality. All proposals submitted should therefore include only public information. Accepted proposals may be published to the INRC member site for community reference (*i.e.* visible to all other members engaged in INRC research), specifically sections 1-7. Groups will have control over all such content on the INRC website and may request for their project details not to be shared at all in this manner with other members, if so desired.

References

[1] M. Davies, A. Wild, G. Orchard, Y. Sandamirskaya, G. A. Fonseca Guerra, P. Joshi, P. Plank and S. Risbud, "Advancing Neuromorphic Computing with Loihi: A Survey of Results and Outlook," *Proceedings of the IEEE,* pp. 911-934, 2021. Results may vary.

[2] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C.-K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataramanan, Y.-H. Weng, A. Wild, Y. Yang and H. Wang, "Loihi: a Neuromorphic Manycore Processor with On-Chip Learning," *IEEE Micro,* 2018.

[3] Intel Labs, "Taking Neuromorphic Computing to the Next Level with Loihi 2 Technology Brief," 2021.

[4] D. Kleyko, M. Davies, E. P. Frady, P. Kanerva, S. J. Kent, B. A. Olshausen, E. Osipov, J. M. Rabaey, D. A. Rachkovskij, A. Rahimi and F. T. Sommer, "Vector Symbolic Architectures as a Computing Framework for Nanoscale Hardware," *arXiv 2106.05268,* 2021.

[5] Y. Sandamirskaya, "Dynamic neural fields as a step toward cognitive neuromorphic architectures," *Frontiers in Neuroscience,* vol. 7, no. 1662-453X, p. 276, 2014.

[6] D. Pecevski, L. Buesing and W. Maass, "Probabilistic Inference in General Graphical Models through Sampling in Stochastic Networks of Spiking Neurons," *PLOS Computational Biology,* vol. 7, pp. 1-25, 2011.

[7] S. Habenschuss, Z. Jonke and W. Maass, "Stochastic Computations in Cortical Microcircuit Models," *PLOS Computational Biology,* vol. 9, pp. 1-28, 2013.

[8] Z. Jonke, S. Habenschuss and W. Maass, "Solving Constraint Satisfaction Problems with Networks of Spiking Neurons," *Frontiers in Neuroscience,* vol. 10, p. 118, 2016.

[9] D. M. Paiton, S. Shepard, K. H. R. Chan and B. A. Olshausen, "Subspace Locally Competitive Algorithms," in *NICE '20*, Heidelberg, Germany, 2020.

[10] Q. Li and C. Pehlevan, "Minimax Dynamics of Optimally Balanced Spiking Networks of Excitatory and Inhibitory Neurons," *arXiv 2006.08115,* 2021.

[11] S. Ludwig, J. Hartjes, B. Pol, G. Rivas and J. Kwisthout, "A spiking neuron implementation of genetic algorithms for optimization," in *BNAIC/Benelearn Communications in Computer and Information Science*, Leiden, 2020.

[12] H. Jang, O. Simeone, B. Gardner and G. Andre, "An Introduction to Probabilistic Spiking Neural Networks: Probabilistic Models, Learning Rules, and Applications,," *IEEE Signal Processing Magazine,* vol. 36, no. 6, pp. 64-77, November 2019.

[13] L. Buesing, J. Bill, B. Nessler and W. Maass, "Neural Dynamics as Sampling: A Model for Stochastic Computation in Recurrent Networks of Spiking Neurons," *PLOS Computational Biology,* vol. 7, no. 11, 2011.

[14] E. Treister and I. Yavneh, "A Multilevel Iterated-Shrinkage Approach to l1 Penalized Least Square Minimization," *IEEE Trans. Signal Proc,* vol. 60, no. 12, pp. 6319-6329, 2012.

[15] G. Bellec, F. Scherr, A. Subramoney, E. Hajek, D. Salaj, R. Legenstein and W. Maass, "A solution to the learning dilemma for recurrent networks of spiking neurons," *Nature Communications,* 2020.

[16] F. Zenke and E. O. Neftci, "Brain-inspired learning on neuromorphic substrates," *Proceedings of the IEEE,* vol. 109, no. 5, pp. 935-950, 2021.

[17] J. Sacramento, R. P. Costa, Y. Bengio and W. Senn, "Dendritic error backpropagation in deep cortical microcircuits," *arXiv:1801.00062,* Dec 2017.

[18] B. Scellier and Y. Bengio, "Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation," *Frontiers in Computational Neuroscience,* vol. 11, no. 1662-5188, p. 24, 2017.

[19] K. Stewart, G. Orchard, S. B. Shrestha and E. Neftci, "Online Few-Shot Gesture Learning on a Neuromorphic Processor," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems,* vol. 10, no. 4, pp. 512-521, 2020.

[20] B. Illing, J. Ventura, G. Bellec and W. Gerstner, "Local plasticity rules can learn deep representations using self-supervised contrastive predictions," *arXiv 2010.08262,* 2021.

[21] A. van den Oord, Y. Li and O. Vinyals, "Representation Learning with Contrastive Predictive Coding," *arXiv 1807.03748,* 2019.

[22] M. K. Benna and S. . Fusi, "Computational principles of synaptic memory consolidation," *Nature Neuroscience,* vol. 19, no. 12, pp. 1697-1706, 2016.

[23] A. . Roxin and S. . Fusi, "Efficient Partitioning of Memory Systems and Its Importance for Memory Consolidation," *PLOS Computational Biology,* vol. 9, no. 7, pp. 1-13, 2013.

[24] N. . Chenkov, H. . Sprekeler and R. . Kempter, "Memory replay in balanced recurrent networks," *PLOS*

Computational Biology, vol. 13, no. 1, p. , 2017.

[25] E. P. Frady, S. J. Kent, B. A. Olshausen and F. T. Sommer, "Resonator Networks, 1: An Efficient Solution for Factoring High-Dimensional, Distributed Representations of Data Structures," *Neural Computation,* vol. 32, no. 12, pp. 2311-2331, 2020.

[26] A. Renner, L. Supic, A. Danielescu, G. Indiveri, B. Olshause, Y. Sandamirskaya, F. T. Sommer and E. F. Paxon, "Neuromorphic Visual Scene Understanding with Resonator Networks," *arXiv,* p. 12880, 2022.

[27] A. Renner, L. Supic, A. Danielescu, G. Indiveri, E. P. Frady, F. T. Sommer and Y. Sandamirskaya, "Neuromorphic Visual Odometry with Resonator Networks," *arXiv,* p. 02000, 2022.

[28] G. Hinton, "The Forward-Forward Algorithm: Some Preliminary Investigations," *arXiv:2212.13345,* 2022.

[29] H. Ramsauer, B. Schäfl, J. Lehner and S. Hochreiter, "Hopfield Networks is All You Need," *arXiv:2008.02217,* 2021.

[30] G. Orchard, E. P. Frady, D. Ben-Dayan Rubin, S. Sanborn, S. B. Shrestha, F. T. Sommer and M. Davies, "Efficient Neuromorphic Signal Processing with Loihi 2," *arXiv:2111.03746,* 2021.

[31] D. Krotov and J. Hopfield, "Large Associative Memory Problem in Neurobiology and Machine Learning," *arXiv:2008.06996v3,* 2021.

[32] L. Kozachkov, K. V. Kastanenka and D. Krotov, "Building Transformers from Neurons and Astrocytes," *bioRxiv,* 2022.

[33] S. B. Shrestha and G. Orchard, "SLAYER: Spike Layer Error Reassignment in Time," in *32nd Conference on Neural Information Processing Systems*, Montréal, Canada, 2018.

Legal Disclaimers

The issuance of this RFP and the submission of a response by a respondent or the acceptance of such a response by Intel Corporation ("Intel") does not obligate Intel in any manner. The RFP is not an offer or a contract. Intel is not obligated to contract for any of the products/services described in the RFP. Intel reserves the right to:

1) amend, modify or withdraw this RFP;
2) revise any requirement of this RFP;
3) waive any requirements of this RFP that are not material;
4) seek clarifications and revisions of responses to this RFP;
5) require supplemental statements or information from any responsible party;
6) accept or reject any or all responses to this RFP;
7) extend the deadline for submission of responses to this RFP or otherwise modify the schedule set forth in this RFP;
8) negotiate potential terms with any respondent to this RFP;
9) engage in discussions with any respondent to this RFP to correct and/or clarify responses;
10) require clarification at any time during the procurement process and/or require correction of responses for the purpose of assuring a full and complete understanding of a respondent's proposal and/or determine a respondent's compliance with the requirements of the solicitation; and
11) cancel, or reissue in whole or in part, this RFP, if Intel determines in its sole discretion that it is its best interest to do so.

Intel may exercise the foregoing rights at any time without notice and without liability to any respondent or any other party for its expenses incurred in preparation of responses hereto or otherwise. All costs associated with responding to this RFP will be at the sole cost and expense of the respondent. Intel makes no representation or warranty and shall incur no liability under any law, statute, rules or regulations as to the accuracy, reliability or completeness of this RFP.