Intel Labs Announcement

Neuromorphic Computing Lab

Intel Neuromorphic Research Community

REQUEST FOR PROPOSALS (RFP)

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

INRC Workshops:

The INRC hosts two annual workshops (Spring and Fall) dedicated to promoting the latest research from Intel Labs and members and providing in-depth technical training on Intel Loihi and Lava. Visit neuromorphic.intel.com for more info and to plan your submission.

Overview

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 chips and software to INRC members so the capabilities and advantages of neuromorphic approaches can be evaluated 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 in Proceedings of the IEEE, "Advancing Neuromorphic Computing With Loihi: A Survey of Results and Outlook," (Davies, et al., 2021) summarizes the findings of the first three years of research with Loihi. We expect new projects funded for ongoing research to incorporate these learnings and focus on the most promising near-term directions for demonstrating applied value.

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, Loihi-2, CPU, and GPU, 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.

What's New?

Starting in 2022, Intel is updating the INRC to help the community grow more quickly, simplify our processes, and provide more useful resources to members. Here's a brief list of changes:

- The INRC Participation Agreement has been simplified to help new members join the community quickly and easily. See the section **INRC Participation Agreement** below for a summary.
- New INRC member categories have been created. While this RFP contains useful information for anyone interested in neuromorphic computing, it is primarily aimed at helping new Research Members submit project proposals. See the section on Affiliate membership and Industry membership to find out if those categories will support your needs.
- A new online form allows you to Join the INRC faster and easier, guiding you through the most essential components of an INRC project proposal.

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. Since 2018, Intel has organized three broad funding calls and will consider sufficiently compelling proposals from PIs at any time.

To request funding outside of a broad call, Intel recommends PIs submit a project proposal using the attached template addressing one of the research vectors defined by this RFP (RV1 through RV5, see below). Out-of-cycle projects should aim for a duration of one year or less. Typical grants from Intel and INRC sponsors support one student or postdoc per project. All proposals should 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 intentions to pursue such opportunities will be considered favorably. Intel will provide letters of support for external funding applications.

Background

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

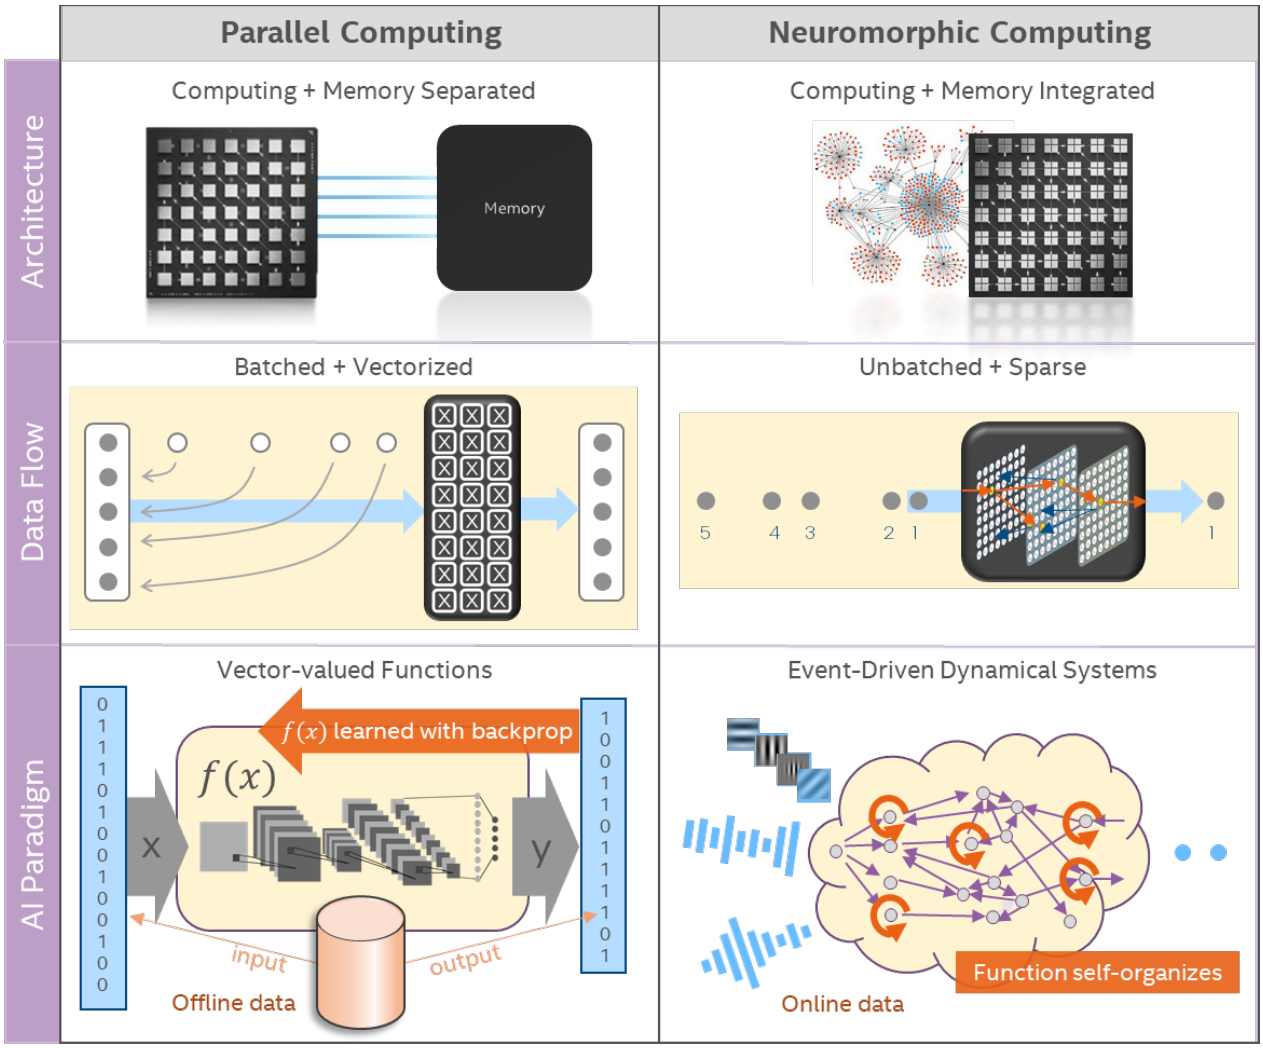

Figure 1. SIMPLIFIED VIEW CONTRASTING TODAY'S PARALLEL COMPUTING AND NEUROMORPHIC COMPUTING, WITH AI PARADIGMS CORRESPONDING TO EACH ONE.

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 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 in order 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.

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

- **Streaming input data** with temporal information structure (e.g. audio, video, or any signals changing on microsecond-to-second time scales), especially when the data events of interest arrive infrequently and unpredictably.
- **A need for fast pattern matching, search, and optimization**. Neuromorphic architectures fundamentally implement a neural network computational model, which have a large body of literature supporting sub-symbolic processing and pattern matching in high-dimensional spaces, as well as optimizing network-defined energy functions through emergent neural dynamics.
- **A need for adaptation, fine-tuning, or associative learning** in response to arriving information.
- **A need for low latency responses**, e.g. as in closed-loop control applications. The time and resource cost of *batching* and *vectorization*, necessary for efficient use of conventional architectures, may be unacceptable.
- **Power constrained**. 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 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 today by chips such as Intel's Loihi series (Davies, et al., 2018), neuromorphic technology can provide value for applications characterized by specific properties, shown in Box 1. These loosely correspond to the same application needs that shaped brain evolution in nature. Rapid response to arriving information allows mobile organisms to evade threats, capture prey, while rapid online learning in response to minimal information 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.

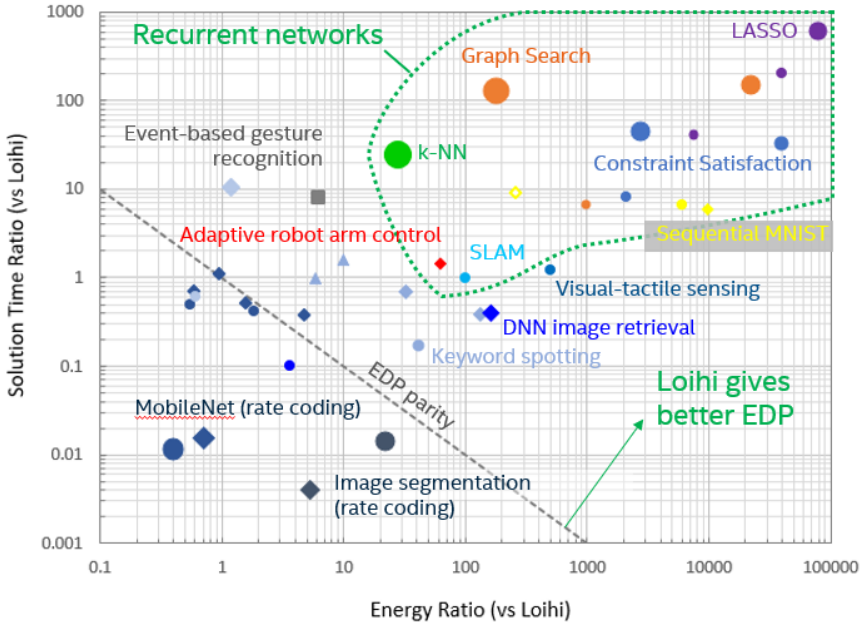

Figure 2. LOIHI RESULTS SHOWING RELATIVE GAINS IN SOLUTION ENERGY AND LATENCY VERSUS REFERENCE ARCHITECTURES. SEE (DAVIES ET AL 2021) FOR DETAILS.

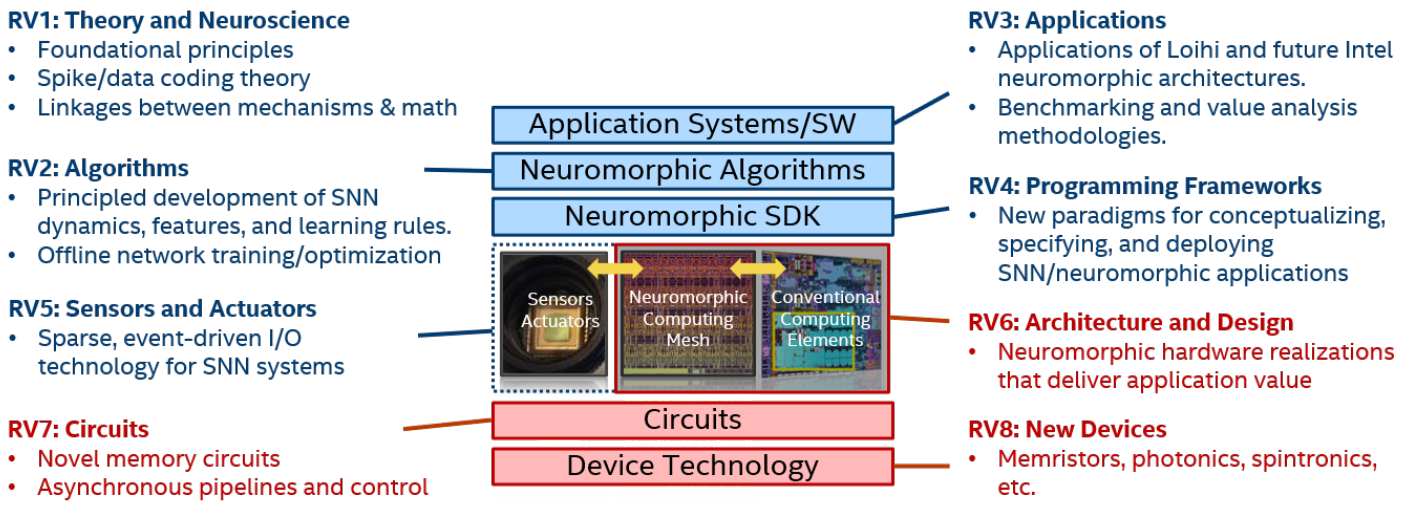

Figure 3. NEUROMORPHIC RESEARCH VECTORS. BLUE VECTORS FALL WITHIN THE DOMAIN OF INRC-FUNDED RESEARCH. RED ARE OUT OF SCOPE.

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 (Davies, et al., 2021) 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 compared to conventional ones, at least not over 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 a wide range of real-world applications are still a severe limitation to progress. To enable commercially relevant applications and attract increased investment to the field, more attention needs to be directed to the most pressing of these near-term challenges. This 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 of 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 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 a 2018 IEEE Micro publication (Davies, et al., 2018) and an Intel Tech Brief (Intel Labs, 2021), 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, e.g. as documented in (Davies et al. 2021). 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 code contributed to the Lava GitHub with permissive licensing terms. This allows 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

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

- **Graded spikes**. Many traditional SNN chips, including Loihi, only support binary-valued spike messages. While binary spikes can perform a remarkable amount of computation (as best demonstrated by the brain), in digital hardware spikes can be easily generalized to carry integer-valued payloads with little extra cost in either performance or energy. Several recent neuromorphic chips support such graded spikes, and future Intel chips will too. These generalized spike messages support event-based messaging, preserving the desirable sparse and time-coded communication properties of SNNs, while also providing greater numerical precision.
- **More general neuron models**. The first gen Loihi supported a generalized leaky-integrate-and-fire (LIF) spiking neuron model with the ability to aggregate neural units into dendritic trees communicating graded dynamic state variables towards the soma (root) compartment. Intel's Loihi 2 neuromorphic processor generalizes this further, with fully programmable neuron models that allow each compartment to state variable changes as nearly arbitrary difference equations with configurable spike conditions and state machines. A greatly expanded set of neuron models are now supported, including adaptive threshold LIF, Resonate-and-Fire, Hopf resonators, sigma-delta coding, and many others (See (Orchard, et al., 2021) for some examples).
- **Three-factor learning rules**. Loihi primarily supports two-factor learning rules (involving pre- and post-synaptic traces), with a third modulatory term set in a diffuse manner from graded "reward" broadcasts. Future chips will support more targeted and localized third factors in learning rules, for example those mapped to specific postsynaptic neurons as projected error signals. Neuron thresholds and other parameters will also support more flexible programmed plasticity.

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.
- Mapping the conceptual frameworks of vector symbolic architectures (Kleyko, et al., 2021) and dynamic neural fields (Sandamirskaya, 2014) to the neuroscientific state-of-the-art.
- 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 can only use state variables that are *locally available* at each synapse and neuron. Novel neuro-inspired features such as stochasticity, structural plasticity, event-driven computation, and temporal information coding can also provide unique gains on neuromorphic architectures.

Algorithms proposals should fully consider all recent learnings (Davies, et al., 2021) 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 sub-vectors are of particular interest.

RV2.1 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 near-term promise of this algorithmic approach for demonstrating more compact, intelligent, and efficient nonlinear signal processing solutions, e.g., for audio processing.

RV2.2 Optimization with neural dynamics

The best quantified Loihi results come from networks that optimize well-defined objectives using attractor and other network-level dynamics. Examples include Lasso regression, constraint satisfaction, and similarity search, with other promising generalizations on the horizon such as Subspace Locally Competitive Algorithm (Paiton, Shepard, Chan, & Olshausen, 2020), Minimax optimization (Li & Pehlevan, 2021), genetic algorithms (Ludwig, Hartjes, Pol, Rivas, & Kwisthout, 2020) and probabilistic inference (Pecevski, Buesing, & Maass, 2011) (Habenschuss, Jonke, & Maass, 2013) (Jang, Simeone, Gardner, & Andre, 2019).

An important challenge involves hierarchically composing such optimizing networks to solve larger problems or to solve the same problems with higher precision. Multiscale modeling techniques from conventional applied math may be of value here, in addition to novel insights from neuroscience.

RV2.3 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 (Bellec, et al., 2020) (Zenke & Neftci, 2021) (Sacramento, Costa, Bengio, & Senn, 2017) (Scellier & Bengio, 2017). 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 (Stewart, Orchard, Shrestha, & Neftci, 2020), surprise-driven learning such as CLAPP (Illing, Ventura, Bellec, & Gerstner, 2021) or Contrastive Predictive Coding (van den Oord, Li, & Vinyals, 2019), and stochastic learning with bounded resources such as complex synapses to extend memory lifetime (Benna & Fusi, 2016), or *replay* for pinning significant memories (Roxin & Fusi, 2013) (Chenkov, Sprekeler, & Kempter, 2017).

RV2.4 Stochastic spiking neural networks

Stochastic spiking neural networks have been used to represent Bayesian or more general graphical models (Pecevski, Buesing, & Maass, 2011) and are known to solve a wide range of hard problems such as the computation of marginal probabilities, or maximum likelihood (Habenschuss, Jonke, & Maass, 2013). These models have for instance been applied to solving constraint satisfaction problems with SNNs (Jonke, Habenschuss, & Maass, 2016) 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 (Davies, et al., 2021).

Opportunities for progress might include (1) advancing stochastic SNN theory and demonstrating these frameworks on neuromorphic hardware, (2) combining stochastic networks with standard components such as offline-trained DNNs and associative memories, and (3) algorithmically extending these frameworks to areas of high value such as probabilistic inference and model predictive control.

RV2.5 Vector Symbolic Architectures

Vector Symbolic Architectures, also known as Hyperdimensional Computing algorithms, have shown significant promise for neuromorphic systems with algorithms such as the resonator network (Frady, Kent, Olshausen, & Sommer, 2020). VSA demonstrations to date tackle small scale problems and scaling these algorithms could prove challenging due to their use of hardware resources.

Research into VSA frameworks that provide sparsity in connectivity and activity may enable this class of algorithms to be efficiently scaled on neuromorphic hardware and far outperform conventional dense-vector approaches. Insofar as the brain implements such hyperdimensional vector symbolic computations, we can be confident that a highly efficient, neuro-inspired sparse implementation exists. We encourage greater attention to this promising new direction for neuromorphic computing.

RV2.6 Offline training and network optimization

While a wealth of backprop-style offline training tools for spiking neural networks have been demonstrated over the past few years, computational constraints limit them to smaller and simpler network architectures and dataset sizes than conventional deep learning. Scaling up computational resources quickly faces diminishing returns. Innovations in offline optimization methods are needed to train larger and more complex neuromorphic networks. Examples could include hybrid training approaches (e.g. ANN conversion-based pre-training followed by direct fine-tuning), Hessian-based methods, 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 involve approximations that introduce errors compared to training conventional ANNs. To overcome this, we see promise in evolutionary methods for training, Bayesian optimization of network structures, hardware-aware training, and neural architecture search to realize better performance with fewer layers.

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

Box 3. What is Lava?

- **Magma**. At its foundation is Intel's next generation SDK that provides a low-level interface for mapping and executing neural network models and applications to neuromorphic hardware. This layer includes cross-platform execution support so applications can be developed in simulation before being deployed to Loihi or other neuromorphic platforms. This layer also includes a profiler that can measure or estimate performance and energy consumption across the targeted back-end platforms.
- **Channel-based asynchronous message passing**. Lava specifies, compiles, and executes a collection of processes mapped to a heterogenous execution platform including both conventional and neuromorphic components. Unifying all inter-process communication is an event-based message passing framework sometimes referred to as Communicating Sequential Processes or the Actor model. Messages in Lava vary in granularity from single-bit spikes to buffered packets with arbitrary payloads.
- **Offline training.** Lava supports tools such as SLAYER (Shrestha & Orchard, 2018) so a range of different event-driven neural networks (LIF SNNs, RF SNNs, SDNNs, and others) can be trained offline with backpropagation and integrated with other modules specified in Lava.
- **Integration with third party frameworks**. Lava is fully extensible, supporting interfaces to third party frameworks such as ROS, YARP, TensorFlow, and hopefully many more in the future, providing a truly heterogeneous execution environment.
- **Python interfaces**. For ease of adoption, all libraries and features in Lava are exposed through Python, with optimized libraries and C/C++/CUDA/OpenCL code under the hood where necessary to provide excellent performance.
- **An open-source framework with permissive licensing**. Lava is hosted publicly on GitHub and runs on CPU/GPU platforms without requiring any legal agreement with Intel. The software is available for free use under BSD-3 and LGPL-2.1 licensing.

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

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 sufficiently compelling proposals, Intel is open to partnering and co-developing technical assets if doing so will support 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. Domain expertise is essential for credible proposals in this RV. Proposals should articulate a strong case for the commercial viability of a successful outcome, based on measurable and significant advances the project will demonstrate over the current state-of-the-art.

Promising example categories of application demonstrations include:

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>

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, .

RV4: Programming Models and Frameworks

The overall software maturity in the neuromorphic field remains low. 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 supports the open-source Lava framework as a solution (See Box 3 "What is Lava?").

We aim to focus greater attention across all levels of a single, unified software framework. 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.

Performance optimization and extension of Lava Compiler to other neuromorphic architectures. The current Lava Compiler and Runtime support heterogeneous execution on different compute resources such as x86/ARM CPUs and Loihi neuro cores. To foster greater convergence and converging developer efforts on a shared platform, the Lava Compiler is open to extension to other neuromorphic backends besides Loihi. In addition, the development of highly efficient compilation algorithms to optimize algorithm performance and resource utilization on memory-constrained neuromorphic chips is still in its early stages. Thus, it leaves many opportunities for community members to improve overall compiler efficiency and speed.

Formal specification and verification of Lava processes spanning structural, functional, and behavioral levels of abstraction. A promising approach for improving the explainability and composability of certain classes of neuromorphic networks is to rigorously derive emergent properties of their dynamics, e.g. attractor states and stability conditions, given a specification of their structure and parameters. Such capabilities would abstract away low-level details of the networks' behavior, such as precise trajectories in phase space, while providing practical insights to the application developer. Such neuromorphic network verification capabilities would require new mathematical analysis tools going beyond the discrete math and logical theorem provers of conventional formal software verification.

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.

Overall Research Goals

Intel's goals for the INRC relate to accelerating neuromorphic research progress and enabling the commercial adoption of future neuromorphic technology:

- Identify, develop, and characterize algorithms that exploit the novel properties of event-based/spiking neuromorphic hardware architectures to deliver orders of magnitude gains in latency, energy efficiency, and data efficiency compared to conventional solutions.
- Guide the iterative development of neuromorphic, non-Von Neumann architectures with the algorithmic and architectural insights from the above.
- Develop an open, cross-platform software framework that aligns and grows the neuromorphic research ecosystem, laying the groundwork for a future commercial developer ecosystem.
- Prototype real-world applications with Lava and Intel's neuromorphic research silicon to identify the areas where neuromorphic technology might deliver the greatest commercial value.

Whenever possible, Intel will support project proposals that align well with these goals.

Member Categories

Intel offers multiple categories of community membership to best match the engagement profiles across a wide range of interested organizations. Intel seeks to support the broadest, most inclusive community possible given legal and other constraints. The first step in engaging in the INRC is for an organization to identify its most appropriate member category:

- **Affiliate Member:** Participate in community events, learn about neuromorphic computing, share research. This category is open to any members of research organizations who want to follow Intel's neuromorphic research and general community activities, but who don't plan to use Intel's hardware technology in a research project right away. This membership is well-suited to theoretical researchers, neuroscientists, industry analysts, and others who don't need direct access to hardware.
- **Research Member:** Join the community to conduct a research project using Intel's Loihi hardware technology. This category includes most of the current members of the INRC. This is the best option for applied scientists and algorithms researchers who do need to access neuromorphic hardware to conduct their research. A formal research proposal is required, either by completing the *INRC Project Proposal Template* if requesting funding, or by submitting supporting documentation via the online INRC engagement system otherwise. Research members must belong to an established research organization, such as a university lab or company.
- **Industry Member:** Evaluate neuromorphic applications for your business or organization. This category is tailored for corporate, government, or large non-profit members seeking to test or evaluate neuromorphic technology for real-world use cases. Potential industry members should be prepared to submit a detailed description of the intended application and the results that will be shared with Intel.

Engagement Process

The process for joining the Intel Neuromorphic Research Community is outlined below.

1) Submit your application online to Join the Intel Neuromorphic Research Community. Research and Industry members should have their project plan ready. Intel will review the submission and contact you with any questions.
2) **Affiliate members:** Intel will contact you when approved and add you to the INRC member database. You'll receive notifications for INRC news and events and access to a range of community benefits.
3) **Research and Industry members:** An **INRC Participation Agreement** must be executed between Intel and all organization(s) affiliated with the request. 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.

Project Proposal Submission

Along with this RFP, we are providing an *INRC Project Proposal Template* document that lists all information and documentation required of each respondent seeking funding. Please 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 over the past four years to be consistent with industry standards and very permissive, while still protecting both the member and Intel. To ensure consistency, 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 offers two forms of access to Loihi-based neuromorphic systems and software for you to use in your research.

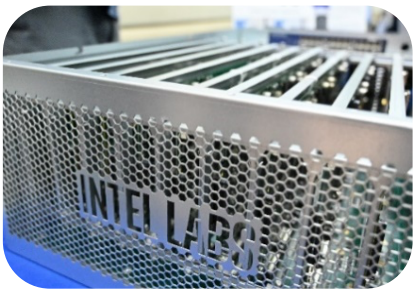

Neuromorphic Research Cloud

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, an even 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.

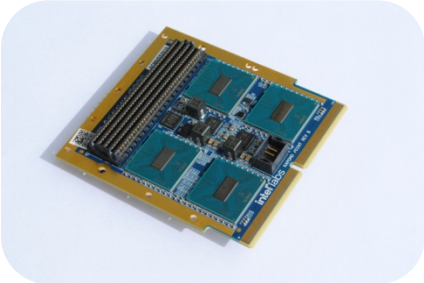

Kapoho Point

Beginning in 2022, Intel can provide some groups a Kapoho Point system for on-site use. This system supports 8 Loihi 2 chips and due to the small form factor and low power requirements, it can be connected directly to sensors, including the Prophesee event-based vision sensor, or robots at the edge. Kapoho Point is designed to support 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 Seeking Funding

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/collaterals 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 closely synergistic 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/tools in classroom setting). 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.

Point of Contact for Inquiries and Submissions

Current proposal submission instructions can be found at the Join the INRC page hosted on the community web site. (https://intel-ncl.atlassian.net/wiki/spaces/INRC/pages/1784807425/Join+the+INRC)

All inquiries should be sent to inrc_interest@intel.com

References

Andrew G. Barto, S. M. (2003). Recent Advances in Hierarchical Reinforcement Learning. *Discrete Event Dynamic Systems, 13*, 341–379.

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*.

Ben Dayan Rubin, D. D., & Fusi, S. (2007). Long memory lifetimes require complex synapses and limited sparseness. *Frontiers in Computational Neuroscience, 1*, 7. Retrieved 8 16, 2021, from http://ini.uzh.ch/admin/extras/doc_get.php?id=42366

Benna, M. K., & Fusi, S. (2016). Computational principles of synaptic memory consolidation. *Nature Neuroscience, 19*(12), 1697-1706. Retrieved 8 16, 2021, from https://nature.com/articles/nn.4401

Botvinick, M., Niv, Y., & Barto, A. C. (2009). Hierarchically Organized Behavior and Its Neural Foundations: A Reinforcement Learning Perspective. *Cognition, 113*(3), 262-280. Retrieved 8 17, 2021, from https://sciencedirect.com/science/article/pii/s0010027708002059

Buesing, L., Bill, J., Nessler, B., & Maass, W. (2011). Neural Dynamics as Sampling: A Model for Stochastic Computation in Recurrent Networks of Spiking Neurons. *PLOS Computational Biology, 7*(11).

Chen, X., Fan, H., Girshick, R., & He, K. (2020). Improved Baselines with Momentum Contrastive Learning. *arXiv 2003.04297*.

Chenkov, N., Sprekeler, H., & Kempter, R. (2017). Memory replay in balanced recurrent networks. *PLOS Computational Biology, 13*(1). Retrieved 8 16, 2021, from http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1005359

Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., . . . Wang, H. (2018). Loihi: a Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro*.

Davies, M., Wild, A., Orchard, G., Sandamirskaya, Y., Fonseca Guerra, G. A., Joshi, P., . . . Risbud, S. (2021). Advancing Neuromorphic Computing with Loihi: A Survey of Results and Outlook. *Proceedings of the IEEE*, 911-934.

Frady, E. P., Kent, S. J., Olshausen, B. A., & Sommer, F. T. (2020). Resonator Networks, 1: An Efficient Solution for Factoring High-Dimensional, Distributed Representations of Data Structures. *Neural Computation, 32*(12), 2311-2331.

Frady, E. P., Kleyko, D., & Sommer, F. T. (2020). Variable Binding for Sparse Distributed Representations: Theory and Applications. *arXiv preprint arXiv:2009.06734*.

Fusi, S., & Abbott, L. F. (2007). Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience, 10*(4), 485-493. Retrieved 8 16, 2021, from https://nature.com/articles/nn1859

Fusi, S., Drew, P. J., & Abbott, L. F. (2005). Cascade models of synaptically stored memories. *Neuron, 45*(4), 599-611. Retrieved 8 16, 2021, from https://sciencedirect.com/science/article/pii/s0896627305001170

Habenschuss, S., Jonke, Z., & Maass, W. (2013). Stochastic Computations in Cortical Microcircuit Models. *PLOS Computational Biology, 9*, 1-28.

Hoare, C. A. (1978). Communicating sequential processes. *Communications of the ACM, 21*(8), 666-677.

Illing, B., Ventura, J., Bellec, G., & Gerstner, W. (2021). Local plasticity rules can learn deep representations using self-supervised contrastive predictions. *arXiv 2010.08262*.

Imam, N., & Cleland, T. A. (n.d.). Rapid learning and robust recall in a neuromorphic olfactory circuit. *(in review)*.

Intel Labs. (2021). Taking Neuromorphic Computing to the Next Level with Loihi 2 Technology Brief. Retrieved from https://www.intel.com/content/www/us/en/research/neuromorphic-computing-loihi-2-technology-brief.html

Jang, H., Simeone, O., Gardner, B., & Andre, G. (2019, November). An Introduction to Probabilistic Spiking Neural Networks: Probabilistic Models, Learning Rules, and Applications,. *IEEE Signal Processing Magazine, 36*(6), 64-77. doi:10.1109/MSP.2019.2935234

Jonke, Z., Habenschuss, S., & Maass, W. (2016). Solving Constraint Satisfaction Problems with Networks of Spiking Neurons. *Frontiers in Neuroscience, 10*, 118.

Katkov, M., Romani, S., Tsodyks, M., & Tsodyks, M. (2017). Memory Retrieval from First Principles. *Neuron, 94*(5), 1027-1032. Retrieved 8 16, 2021, from https://cell.com/neuron/fulltext/s0896-6273(17)30292-1

Kleyko, D., Davies, M., Frady, E. P., Kanerva, P., Kent, S. J., Olshausen, B. A., . . . Sommer, F. T. (2021). Vector Symbolic Architectures as a Computing Framework for Nanoscale Hardware. *arXiv 2106.05268*. Retrieved from https://arxiv.org/abs/2106.05268

Kreiser, R., Renner, A., Leite, V. R., Serhan, B., Bartolozzi, C., Glover, A., & Sandamirskaya, Y. (2020). An On-chip Spiking Neural Network for Estimation of the Head Pose of the iCub Robot. *Frontiers in Neuroscience, 14*(1662-453X), 551.

Li, Q., & Pehlevan, C. (2021). Minimax Dynamics of Optimally Balanced Spiking Networks of Excitatory and Inhibitory Neurons. *arXiv 2006.08115*.

Lohstroh, M. (2020). Reactors: A Deterministic Model of Concurrent Computation for Reactive Systems. EECS Department, University of California, Berkeley.

Ludwig, S., Hartjes, J., Pol, B., Rivas, G., & Kwisthout, J. (2020). A spiking neuron implementation of genetic algorithms for optimization. *BNAIC/Benelearn Communications in Computer and Information Science. 1398*, pp. 91--105. Leiden: Springer.

Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (2015). Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *The Journal of Neuroscience, 35*(21), 8145-8157. Retrieved 8 17, 2021, from https://ncbi.nlm.nih.gov/pubmed/26019331

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

Paiton, D. M., Shepard, S., Chan, K. H., & Olshausen, B. A. (2020). Subspace Locally Competitive Algorithms. *NICE '20.* Heidelberg, Germany.

Parisi, G. (1986). A nenory which forgets. *J Phys A, 10*(19), 617.

Pecevski, D., Buesing, L., & Maass, W. (2011). Probabilistic Inference in General Graphical Models through Sampling in Stochastic Networks of Spiking Neurons. *PLOS Computational Biology, 7*, 1-25.

Pehlevan, C. (2019). A Spiking Neural Network with Local Learning Rules Derived From Nonnegative Similarity Matching. *arXiv*. Retrieved from https://arxiv.org/abs/1902.01429

Ravindran, B., & Barto, A. G. (2003). *Relativized options: choosing the right transformation*. Retrieved 8 17, 2021, from https://aaai.org/papers/icml/2003/icml03-080.pdf

Romani, S., Amit, D. J., & Amit, Y. (2008). Optimizing one-shot learning with binary synapses. *Neural Computation, 20*(8), 1928-1950. Retrieved 8 16, 2021, from https://mitpressjournals.org/doi/10.1162/neco.2008.10-07-618

Roxin, A., & Fusi, S. (2013). Efficient Partitioning of Memory Systems and Its Importance for Memory Consolidation. *PLOS Computational Biology, 9*(7), 1-13. Retrieved 8 16, 2021, from https://ncbi.nlm.nih.gov/pmc/articles/pmc3723499

Sacramento, J., Costa, R. P., Bengio, Y., & Senn, W. (2017, Dec). Dendritic error backpropagation in deep cortical microcircuits. *arXiv:1801.00062*.

Sandamirskaya, Y. (2014). Dynamic neural fields as a step toward cognitive neuromorphic architectures. *Frontiers in Neuroscience, 7*(1662-453X), 276.

Scellier, B., & Bengio, Y. (2017). Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation. *Frontiers in Computational Neuroscience, 11*(1662-5188), 24.

Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward Causal Representation Learning. *IEEE Advances in Machine Learning and Deep Neural Networks, 109*(5).

Shrestha, S. B., & Orchard, G. (2018). SLAYER: Spike Layer Error Reassignment in Time. *32nd Conference on Neural Information Processing Systems.* Montréal, Canada.

Stewart, K., Orchard, G., Shrestha, S. B., & Neftci, E. (2020). Online Few-Shot Gesture Learning on a Neuromorphic Processor. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems, 10*(4), 512-521.

Susi, G., Garcés, P., Paracone, E., Cristini, A., Salerno, M., Maestú, F., & Pereda, E. (2021). FNS allows efficient event-driven spiking neural network simulations based on a neuron model supporting spike latency. *Scientific Reports, 11*(1), 12160.

Tang, G., Shah, A., & Michmizos, K. P. (2019). Spiking Neural Network on Neuromorphic Hardware for Energy-Efficient Unidimensional SLAM. *IROS.* Retrieved from https://arxiv.org/abs/1903.02504

van den Oord, A., Li, Y., & Vinyals, O. (2019). Representation Learning with Contrastive Predictive Coding. *arXiv 1807.03748*.

Zenke, F., & Neftci, E. O. (2021). Brain-inspired learning on neuromorphic substrates. *Proceedings of the IEEE, 109*(5), 935-950.

Zhang, J., Yu, H., & Xu, W. (2021). Hierarchical Reinforcement Learning By Discovering Intrinsic Options. *international conference on Learning Representations.*

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.