A Long Short-Term Memory for AI Applications in Spike-based Neuromorphic Hardware

Arjun Rao[1,+], **Philipp Plank**[1,3,+], **Andreas Wild**[2], **and Wolfgang Maass***[1]

[1]Institute of Theoretical Computer Science, Graz University of Technology, Inffeldgasse 16b, Graz, Austria
[2]Intel Labs, Intel Corporation, 2111 NE 25th Ave, Hillsboro, OR 97124, USA
[3]Intel Labs, Intel Corporation, Lilienthalstr. 15, 85579 Neubiberg, Germany
*Corresponding author maass@igi.tugraz.at
[+]These authors contributed equally to this work

ABSTRACT

In spite of intensive efforts it has remained an open problem to what extent current Artificial Intelligence (AI) methods that employ Deep Neural Networks (DNNs) can be implemented more energy-efficiently on spike-based neuromorphic hardware. This holds in particular for AI methods that solve sequence processing tasks, a primary application target for spike-based neuromorphic hardware. One difficulty is that DNNs for such tasks typically employ Long Short-Term Memory (LSTM) units. Yet an efficient emulation of these units in spike-based hardware has been missing. We present a biologically inspired solution that solves this problem. This solution enables us to implement a major class of DNNs for sequence processing tasks such as time series classification and question answering with substantial energy savings on neuromorphic hardware. In fact, the Relational Network for reasoning about relations between objects that we use for question answering is the first example of a large DNN that carries out a sequence processing task with substantial energy-saving on neuromorphic hardware.

Energy consumption is a major impediment for more wide-spread applications of new AI-methods that use DNNs, especially in edge devices. Spike-based neuromorphic hardware is one direction that promises to alleviate this problem. This research direction is partially motivated by the method that brains use to run even more complex and larger neural networks than those DNNs that are used in current AI, with a total energy consumption of just 20W: Neurons in the brain only rarely emit spikes which mostly triggers energy consumption in neurons and synapses. But it has remained an open problem as to how DNNs that are needed for modern AI solutions could be implemented in neuromorphic hardware in such a sparse firing mode. Another open problem is how the LSTM units of such DNNs, that are needed for providing a working memory for sequence processing tasks, could be implemented in spike-based neuromorphic hardware. We present a biologically inspired solution to the second problem, that simultaneously provides a step towards also solving the first problem, since it reduces the firing activity of neurons that hold working memory content. We combine this method with a brain-inspired technique called membrane voltage regularization for enforcing sparse firing activity during the training of the DNN. We have tested the impact of these two innovations on computational performance and energy consumption for two benchmark tasks in an implementation on a representative spike-based chip: Intel's neuromorphic research chip Loihi [Davies et al., 2018]. We find significant reductions in the energy-delay product (EDP). In contrast to power, EDP accounts for the true energy and time cost per task/workload/computation. Simultaneously, these implementations demonstrate that two hallmarks of cognitive computations, both in brains and in machine intelligence, working memory and reasoning about relations between concepts or objects, can in fact be implemented more efficiently in spike-based neuromorphic hardware than in GPUs, the standard computing hardware for implementing DNNs.

Implementing a long short-term memory in spike-based neuromorphic hardware

Working memory is maintained in an LSTM unit in a special memory cell, to which read- and write-access is gated by trained neurons with sigmoidal activation function [Hochreiter and Schmidhuber, 1997]. Such an LSTM unit is difficult to realize efficiently in spike-based hardware. However, it turns out that by simply adding a standard feature of some biological neurons, slow after-hyperpolarizing (AHP) currents, a spiking neural network (SNN) acquires similar working memory capabilities as LSTM units over the time scale of the AHP currents. These AHP currents lower the membrane potential of a neuron after each of its spikes (see Fig. 1). Furthermore, these AHP currents can easily be implemented on Loihi with the desirable side benefit of reducing firing activity, and therefore
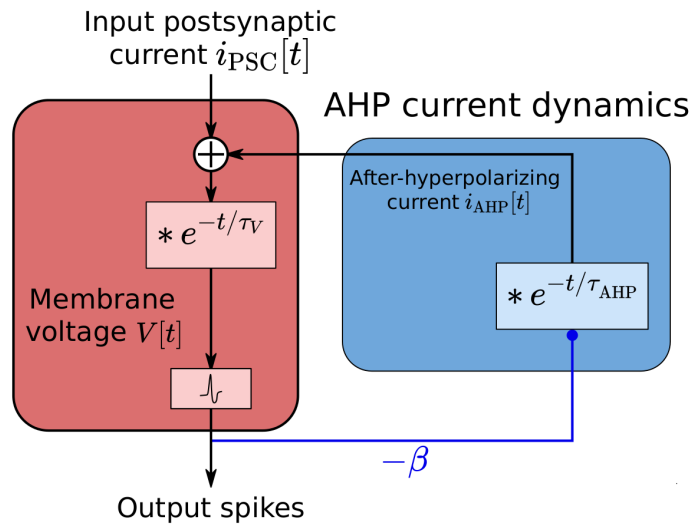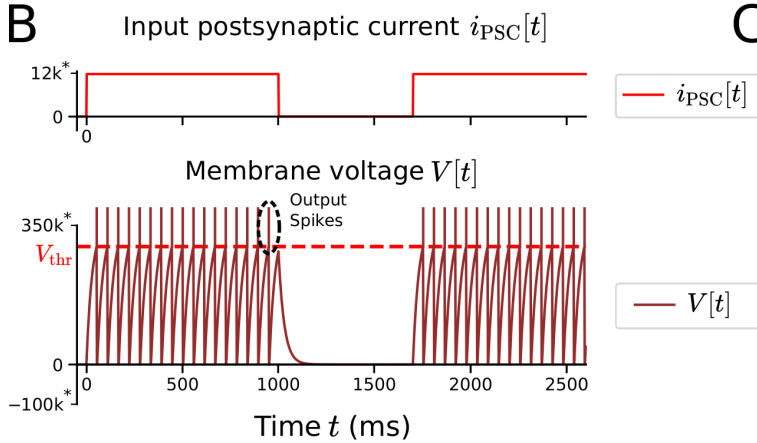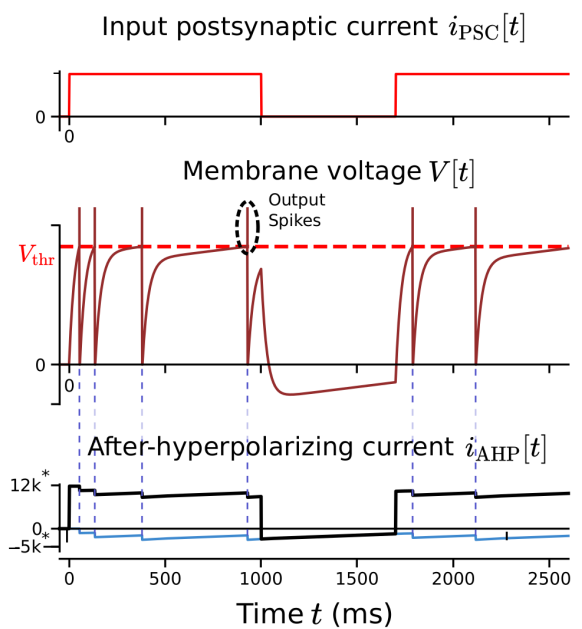
Figure 1. Schematics and dynamics of LIF neurons with and without AHP currents − **A)** Schematic for the implementation of spike frequency adaptation on LIF Neurons. **B)** Shows the response of the LIF neuron model without AHP currents (red compartment in panel A) to a synthetic constant input current. The input postsynaptic current (PSC) $i_{\text{PSC}}[t]$ is leaky-integrated into the membrane voltage $V[t]$. Spikes are emitted and $V[t]$ is reset each time the voltage crosses the threshold V_{thr}. In response to a piecewise constant input PSC, the neuron fires at a constant rate. **C)** The response to a piecewise constant input PSC, of a LIF neuron with AHP currents that shows spike frequency adaptation. The adaptation is implemented by means of an after-hyperpolarizing (AHP) current $i_{\text{AHP}}[t]$ triggered by the spiking of the neuron. Each output spike decreases (makes more negative) the AHP current thus reducing the total current that is integrated. This weakens subsequent spiking, and we see that even with a constant input PSC, the spike rate decreases over time i.e. we implement spike frequency adaptation (SFA). The decay of i_{AHP} is usually much slower than the decay of the membrane voltage ($\tau_{\text{AHP}} >> \tau_V$). Thus even after an extended gap of 700ms, the neuron retains memory of its previous input spikes and shows weaker spiking in response to the input PSC.

energy consumption. In addition, neurons with slow AHP currents capture another essential feature of LSTM units: Gradients can go iteratively through the content of a memory cell of an LSTM unit without being subject to exponential growth or decay because the content of the memory cell is effectively connected to itself with a weight of size 1. The current amplitude of AHP currents can be viewed as a replacement of the content of the memory cell of an LSTM unit, and because this amplitude decays slowly, gradient that go backwards in time through this hidden variable are also protected from exponential growth and decay. Therefore SNNs that contain neurons with slowly changing AHP currents can be trained very well with backpropagation through time (BPTT).

We refer to a SNN that contains these LIF neurons with slowly changing AHP currents, as a long short-term memory SNN (LSNN), borrowing the terminology of [Bellec et al., 2018b]. There, the slow dynamics of a different hidden variable of a spiking neuron model, a time-varying firing threshold, was used to provide a longer short-term memory. But this hidden variable cannot be readily implemented on Loihi. Note that neurons with AHP currents can participate also in the spike-based computations. Hence working memory function and computational processing need not be allocated to spatially separated units in the resulting LSNN. This is important because shuffling of information between processing and memory is commonly viewed as an important factor of the high energy consumption of standard computing hardware.

Fig. 1 B shows the dynamics of a LIF neuron without AHP currents, where the neuron performs leaky integration of the input postsynaptic current (PSC) $i_{\mathrm{PSC}}[t]$ to calculate the membrane voltage $V[t]$. The neuron emits spikes when this voltage exceeds a threshold and resets its value to zero. Fig. 1 C shows the dynamics of a LIF neuron with spike-induced AHP currents that hyperpolarize the membrane voltage. The AHP current (i_{AHP}) increases by an amount β whenever the neuron spikes, i.e., z(t) = 1, and decays between spikes with a large time constant of τ_{AHP}. This current, along with the input current i_{PSC} are leaky-integrated to calculate the membrane voltage. Upon each output spike, the increased negative value of the i_{AHP} reduces the total input current into the neuron, and thus inhibits subsequent spikes. The large value of τ_{AHP} (> 100ms) is what enables the recurrent network to retain memory over larger time spans. LIF neurons with AHP currents are precisely defined as follows:

$$i_{\mathrm{AHP}}[t + \Delta t] = \alpha_{AHP}\ i_{\mathrm{AHP}}[t] - \beta\ z[t] \tag{1}$$

$$V[t + \Delta t] = \begin{cases} \alpha_V V[t] + \frac{1}{g_V}\left(i_{\mathrm{PSC}}[t + \Delta t] + i_{\mathrm{AHP}}[t + \Delta t]\right) & \text{if neuron is not refractory} \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where $\alpha_V = e^{\frac{-\Delta t}{\tau_V}}$ and $\alpha_{AHP} = e^{\frac{-\Delta t}{\tau_{\mathrm{AHP}}}}$; τ_V and τ_{AHP} are the time constants of exponential decay of the membrane voltage and the AHP current respectively with $\tau_{\mathrm{AHP}} \gg \tau_V$. g_V is the membrane conductance. The definition of i_{PSC} as a function of input spikes, and a more detailed model is described in Methods. For the purpose of this paper, the membrane voltage and currents are unitless quantities and their values represent their values as seen in Loihi. The multi-compartment feature of Loihi allows the maintenance of the AHP current within the same neuron, hence LIF neurons with AHP currents can be implemented very efficiently on Loihi. Fig. 1 A shows the schematic of this multi-compartment neuron.

These networks of adaptive LIF neurons can be programmed onto the massively parallel Loihi multi-core architecture. Each of its neuro-core consists of multiple independent SRAMs holding neural and synaptic parameters. In addition, it computes the dynamics of up to 1024 single-compartment (no AHP) or up to 512 2-compartment (with AHP) neurons locally in-memory and thereby avoids expensive data movement between processing elements and external memories. 128 of such interconnected neuro-cores form a Loihi chip. Systems like the 32 chip Nahuku platform finally allow us to execute large scale models such as our Spiking RelNet.

It had been shown already in [Bellec et al., 2018b] that a related mechanism for spike-frequency adaptation, which is more difficult to realize in neuromorphic hardware, enables networks of spiking neurons to achieve a similar performance level as LSTM networks for many temporal processing tasks in current AI. We show that the previously discussed mechanism with AHP currents provides a similarly good performance of SNNs.

Comparing the energy consumption of spiking and non-spiking RNNs with Long Short-Term Memory for a standard time-series classification benchmark task

In order to test the energy efficiency of the proposed emulation of LSTM units with LIF neurons with AHP currents, we use a classical time series classification task: sequential MNIST (sMNIST). Here the pixel values of handwritten digits from the MNIST dataset [LeCun et al., 2010] are presented sequentially in a fixed order, pixel by pixel, and the task is to identify the underlying digit. The gray values of pixels are encoded by spikes through a population of

spiking input neurons that fire when the gray value crosses some threshold, where each neuron in the population has a different threshold (see Fig. 2 A). We trained a recurrent network consisting of 240 LIF neurons for this task, and implemented it on Loihi. A random subset of 100 of them were equipped with AHP currents. Using the technique of DEEP-R [Bellec et al., 2018a], we train the network to be sparsely connected with 20% of the recurrent connections enabled. Details on the network structure and parameters can be found in Methods.
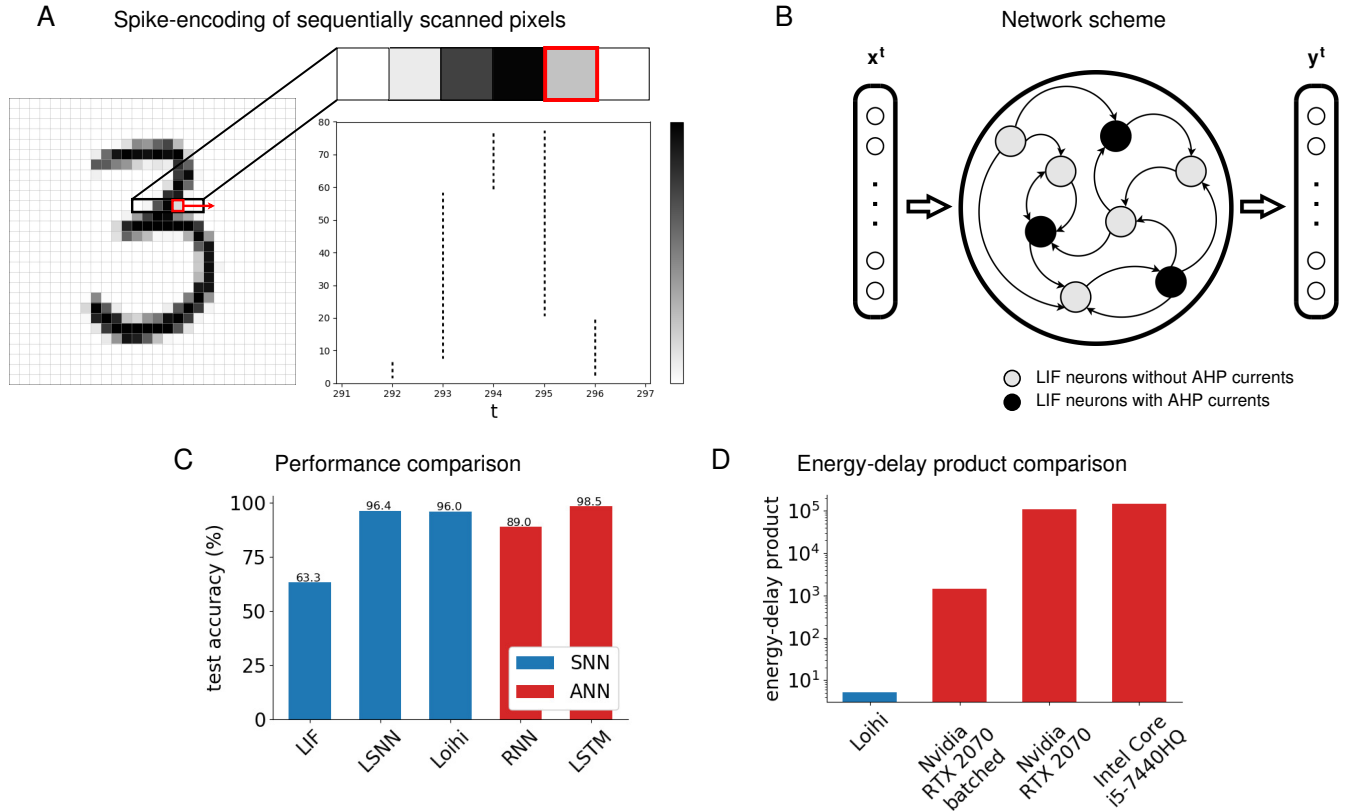

Figure 2. Illustration of the sMNIST task and comparison of performance and energy consumption on spiking and non-spiking hardware. A) The input pixels get encoded by spikes based on a threshold crossing method for a sequence of pixel values. 80 thresholds were used, represented by 80 input neurons, which send spikes depending on the change of the pixel value with respect to the previous pixel value. **B)** The network consist of an input layer sending spikes, a recurrently connected layer of LIF neurons with and without AHP currents, and a linear readout layer. **C)** The classification accuracy of the network running on Loihi was compared to the full precision LSNN, a network of LIF neurons without AHP currents, an artificial RNN, and an LSTM network, as in [Bellec et al., 2018b] **D)** The EDP was used to compare the time and energy performance of the spiking network running on Loihi and a corresponding LSTM network running on the GPU Nvidia RTX 2070 Super utilizing parallel evaluation of 100 samples at the same time (batched) and one sample at a time as well as the CPU Intel Core i5-7440HQ[1].

In order to compare accuracy, execution time and energy consumption with conventional hardware, we also implemented an LSTM network for solving the same task on CPUs and GPUs. The test accuracy of the spiking network on Loihi was 96.0% which is competitive against the full precision artificial networks as well as the best reported LSNN, using full precision from [Bellec et al., 2018b] (see Fig. 2 C). We focus on delay (execution time) and energy consumption as the main metrics in our benchmarks. Typically, there is a trade off between energy consumption and delay, e.g., increased energy (supply voltage) will decrease the delay, in integrated circuit chips build with CMOS technology, which is used in modern CPUs, GPUs and also Loihi. Therefore the product of the energy value and the measured delay, the EDP is ideal to compare applications between different hardware architectures, if these applications have a clear delay metric, e.g., time per classification. The EDP of the SNN running on Loihi is 4 orders of magnitude lower than the network on CPU or GPU (see Fig. 2 D) in batch size 1 regime, with Loihi outperforming over 2x on execution time and over a 1000x on energy consumption per inference.

Details regarding the benchmark procedure can be found in the Supplement.

A reason for this significant improvement in execution time and energy consumption compared to LSTMs on conventional hardware is the relatively small network of a few hundred neurons which is sufficient to solve this task. Thus, the network fits on a single chip of Loihi and uses only one or two neuro-cores. Being able to keep spike traffic within a neuro-core is the fastest and most efficient way to process spikes with Loihi. Another reason is the task itself, as processing the input pixel by pixel in a time series manner benefits the LSNN architecture of the network. The information processed is sparse over time, meaning firstly that the amount of information of a single pixel value is low and we require only a single time step to transfer this to the network, and secondly the large number of time steps (pixels) allows sufficient time for the neural states to evolve and assimilate information from the different pixels. Therefore this network architecture on neuromorphic hardware is most effective on processing time series data. Another aspect of performance in conventional neural networks is using parallel processing of batches of data to increase the throughput. Even with a batch size of 100 on the GPU the spiking network on Loihi operating in the batch size 1 regime is still more efficient. Furthermore, multiple instances of the spiking network for parallel computation on Loihi would also be possible, although there is a bottleneck for the input data transfer on current Loihi boards.

Energy-efficient implementation of a large DNN for relational reasoning in neuromorphic hardware

We wondered whether this implementation of working memory in spiking neurons could be used to also implement large DNNs for more demanding sequence processing tasks in an energy-efficient manner in spike-based neuromorphic hardware. Therefore we implemented and tested a spiking variant of the relational network (RelNet) of [Santoro et al., 2017] on Loihi, which we refer to as the Spiking RelNet. The question of whether this can be done in an energy-efficient manner is quite non-trivial since the Spiking RelNet consists primarily of feed-forward networks. The Spiking RelNet takes as input a set of K objects and a single question that are encoded respectively by input spike trains $o_1(t), \ldots, o_K(t)$ and $q(t)$, see Fig. 4 B. As indicated in Fig. 4 A it computes the function

$$RN\left([o_1(t), o_2(t), \ldots, o_K(t)], q(t)\right) = f_\phi\left(f_{agg}\left(\sum_{1 \leq i \leq j \leq N} g_\theta\left(o_i(t), o_j(t), q(t)\right)\right)\right), \tag{3}$$

with the output given through one-hot encoding of words by readout units. The only recurrent network modules are the ones indicated as module B of Fig. 4, that transform each input sequence (a sentences of words in natural language) into spiking activity of 200 neurons within a compressed time span of 37ms [2], see Fig. 4B. This input embedding of sentences was carried out by LSTM networks in the RelNet [Santoro et al., 2017], and is carried out by LSNNs in the Spiking RelNet. In the next processing step (panel C of Fig. 4), the resulting compressed spike codes for each pair of sentences in the story and for the question are processed in parallel by a copy of a feed-forward LIF network that implements the relational function $g_\theta\left(o_i(t), o_j(t), q(t)\right)$ that extracts salient relational information for question q from the two sentences. The outputs of these network modules are superimposed and connected to a LIF layer one-to-one, which implements the element-wise function f_{agg} (aggregation function in panel D). The readout network f_ϕ processes the output of f_{agg} through another feed-forward LIF network. The feed-forward networks don't use the AHP current. The answer to the question q is then given by an application of soft-max to one-hot readout neurons that each favor a particular word as the answer (see panel E in Fig. 4). For more details, see Methods and Supplement.

In the above description, we observe that the Spiking RelNet uses both recurrent networks (part B of Fig. 4 A) as well as feed-forward networks (Parts C, D, E of Fig. 4 A) to perform the calculation. When scaling up to tasks with a larger number K of objects, the fraction of these feed-forward components of RelNet increases (See Fig. 3 B). The reason is that the number of recurrent network modules scales linearly with K, whereas the number of feed-forward modules that compute the function g_θ increase quadratically with K, since we have an instance of g_θ for each pair of

---

[1]Loihi: Nahuku board (ncl-ghrd-01), CPU: Intel Core i9-7920X, RAM: 128GB, OS: Ubuntu 16.04.6 LTS, NxSDK: 0.95
Nvidia RTX 2070: Nvidia RTX 2070 Super, GPU-RAM: 8GB, CPU: Intel Core i7-9700K, RAM: 32GB, OS: Ubuntu 16.04.6 LTS, Python 3.6.5, TensorFlow-GPU: 1.14.0, CUDA: 10.0.
Intel Core i5-7440HQ: RAM: 16GB, OS: Windows 10 (build18362), Python 3.6.7, TensorFlow: 1.14.1
Performance results are based on testing as of July 9, 2021 and may not reflect all publicly available security updates. Results may vary.
[2]All time intervals and time constants are specified in terms of Loihi computation steps where we use the convention of one step corresponding to 1ms time (see Methods)
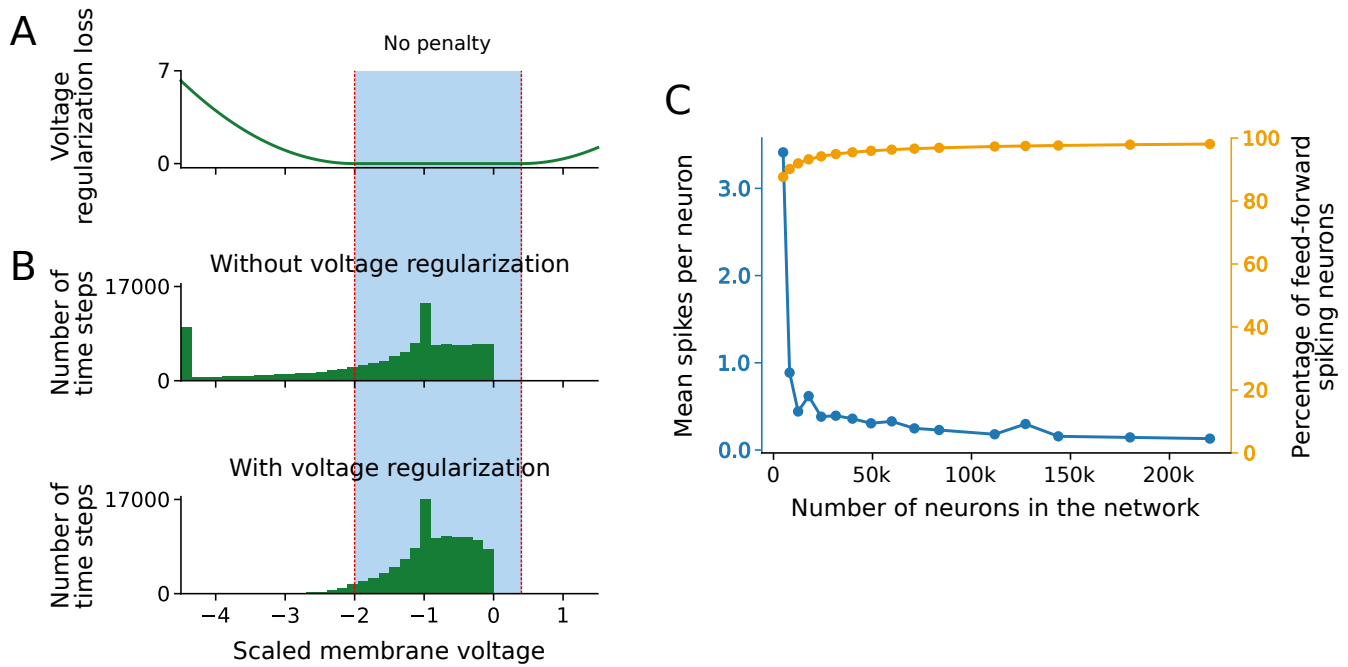
Figure 3. **Illustration of voltage regularization and its its capability to enforce –in conjunction with spike rate regularization– a sparse firing regime. A)** The voltage regularization penalty as function of the value taken by the scaled membrane voltage at a particular time step. The scaled membrane voltage is as defined in Eq. 13. A value of 0 corresponds to the spiking threshold, and a value of -1 corresponds to the value of the voltage corresponding to a zero input PSC. The membrane voltage is thus penalized if the scaled voltage is outside the range $[-2, 0.4]$. **B)** The distribution of the scaled voltage values across different batches, neurons, and time steps with and without regularization. **C)** The spikes used per neuron in relation to the network size (which varies for different story sizes). One observes that larger networks use fewer spikes per neuron as a result of spike rate regularization combined with voltage regularization, which results in savings in energy when run on hardware.

objects. Consequently, the numerous instances of the relational function (g_θ) occupy the majority of the hardware resources (see Fig. 5 C). This increasing fraction of feed-forward network modules is problematic from the perspective of energy efficiency, since prior emulations of feed-forward networks in spike-based hardware demonstrated that their advantage regarding energy-consumption gets lost for larger networks when high classification accuracy is required [Davies et al., 2021]. The reason is that these prior implementations had to use spike rate coding instead of event-based processing in order to achieve high classification accuracy. However, rate coding uses many spikes per neuron, thereby moving the network out of an energy-efficient working regime, and also increases the computation time of the network, i.e., reduces its throughput.

We show that this obstacle, which was largely based on experience with CNNs, can be overcome in the case of RelNet, since these networks can be implemented with high accuracy in a more event-based working regime. An important underlying difference to CNNs is that even for large problem instances, i.e., stories with many sentences, the number of relations that are relevant for answering a question tend to increase only linearly with the length of the story. Hence, with an aggressive spike-rate regularization during training (described in Methods and Supplement), one can force the network to focus its spiking activity on those events where potentially relevant relations are extracted from pairs of sentences. However, such strong spike rate regularization tends to affect the network performance in a substantial manner, since it drives many neurons into an ineffective state where their membrane potential is far away from the firing threshold, see the upper part of Fig. 3 B. We counter this this by adding an additional regularization term, called the voltage regularization loss to the loss function that penalizes the occurrence of these neuron states, see Fig. 3 A and Methods.

In addition, we introduced another method for encouraging the network to work in an event-based processing regime: We forced the network to encode its output in the membrane potential of readout neurons at a particular point in time (marked in Fig. 4 E). This compression of the time window for producing the network output induced upstream feed-forward parts C, D, E of RelNet to constrain their firing activity to rather short time windows, see
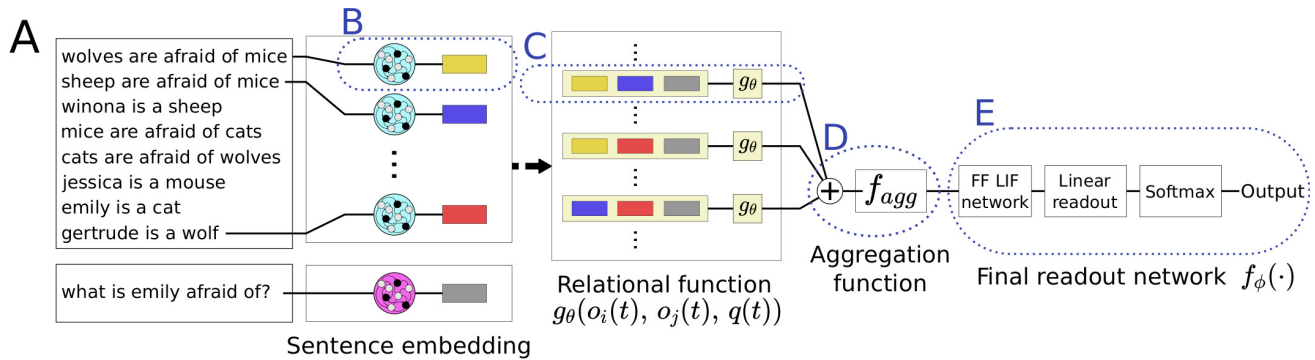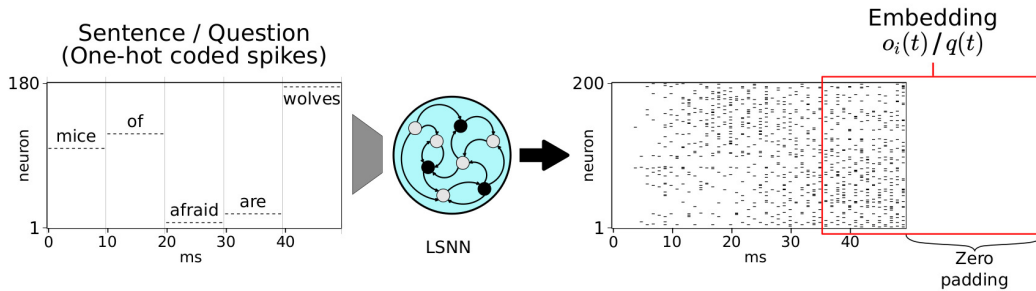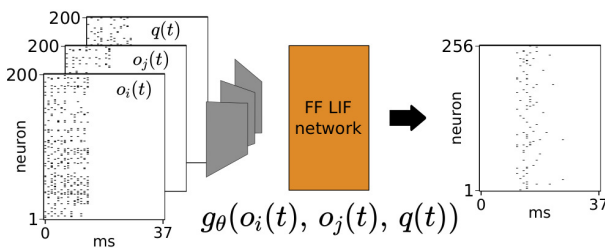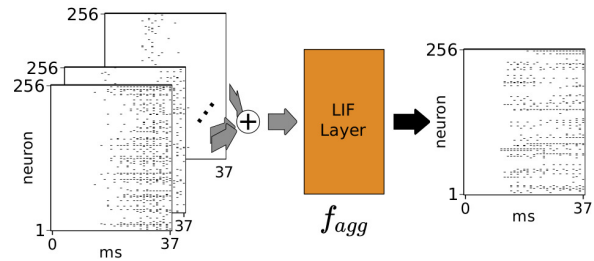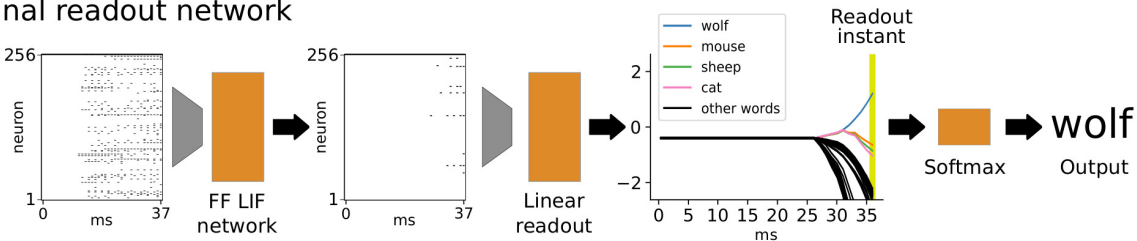
Figure 4. Spiking RelNet architecture and spike-coding schemes that it uses. A) The top-level Spiking RelNet architecture. We embed each sentence and the question into spike sequence objects $o_i(t)$ and $q(t)$ respectively via an LSNN. For each pair of sentence objects $o_i, o_j :: 1 \leq i \leq j \leq 20$, we apply the relational function g_θ to the triplet $(o_i(t), o_j(t), q(t))$. The outputs of the relational function are aggregated in a LIF Layer f_{agg} and then passed to the final readout function f_ϕ. **B)** The embedding scheme, where each word is provided for $T_{\text{word}} = 10$ms with one-hot coded spikes, aligned so that the first word is provided at the very end of the duration. The spikes in the last $T_{\text{inp}} = 14$ms are padded to a length of $T_{\text{sim}} = 37$ms (red box) to form a time-compressed sentence embedding $o_i(t)$ and $q(t)$. **C)** An instance of the spiking relational function g_θ operating on a sample triplet $(o_i(t), o_j(t), q(t))$. **D)** The aggregation layer is a layer of LIF neurons that receive one-to-one connections from each relational function instance. This aggregates the spike trains from across the relational function instances and outputs a spike sequence for the readout network. **E)** The final readout function consists of a three layer feed-forward LIF network followed by a linear readout (with one neuron per word in the dictionary), that integrates synaptic inputs only during the last 10ms (marked as yellow bar). The value of the readout at the final time step provides input to the softmax, whose output produces the the final answer through one-hot encoding of words.

Fig. 4. In addition, both the readout neurons and all network neurons used a rather short membrane time constant of 7ms, which makes it difficult to integrate information from firing rates of upstream neurons. As a result, the spike rate regularization managed to keep the average firing rates very low, in spite of the theoretically possible maximal firing rate of 1000 Hz caused by the absence of a refractory period in the neuron model, which was employed to enhance the backwards propagation of gradients in BPTT. As result, we see in Fig. 3 C that most neurons fired at most one spike during a network computation. Furthermore, the number of spikes per neuron decreased when RelNet was scaled up to larger instances that can answer questions about longer stories.

We tested the performance of the resulting spike-based RelNet implementation on Loihi for a standard benchmark dataset for question-answering; The bAbI dataset introduced by [Weston et al., 2015], that were also used for testing RelNet by [Santoro et al., 2017] This dataset consists of 20 different types of tasks, that each probe different challenges in reasoning about relational information contained in a set of sentences, i.e., a story. For example, tasks 4 and 5 require reasoning about a set of facts that are provided in the form of sentences with 2 arguments ("The office is north of the bedroom.") or 3 arguments ("Mary gave the cake to Bill."). Task 14 requires reasoning about temporal relationships between events, task 15 requires basic deduction, task 18 requires reasoning about relative sizes of objects, task 19 requires planning of a path, and task 20 requires reasoning about the likely motivations of an agent (see Supplement for an examples from tasks 15, 18, 19, 20). The questions are formulated in such a way that an answer can be given with a single word via one-hot encoding in the output (or with a sequence of two words in the case of path planning in Task 19; one has here an output line for each such possible sequence). According to the convention of [Weston et al., 2015] and [Santoro et al., 2017] a task is considered as being solved if the network has an error rate of at most 5% on instances of the task that had not been used for training. When applying a RelNet to solve this task, each sentence (question) forms an object o_i (q) that is embedded via LSNN's to a spiking representation $o_i(t)$ ($q(t)$). Thus, the difficulty of a particular instance of a bAbI task, and the required size of the RelNet grows quadratically with the number of sentences in the story, since the number of potential relations between the contents of sentences (in the context of the question) grows quadratically.

The whole SNN implementation of RelNet was largely trained end-to-end via BPTT for 17 of the bAbI tasks, with some extra measures to speed up training time (see Methods). We exluded 3 of the 20 bAbI tasks, 'Task 2: Two Supporting Facts', 'Task 3: Three Supporting Facts', and 'Task 18: Basic Induction', because also the ANN implementation of RelNet from [Santoro et al., 2017] was not able to solve these 3 tasks. The network is able to solve 16/17 tasks that it was trained on to errors under 5%. The performance of the network was unsatisfactory on task 17 "Positional Reasoning", as a result of the complex sentences needing more time steps to process (see Supplement).

Optimizing the performance of large RelNets in spike-based hardware.

For the longest stories that contain 20 sentences, the network contains 238604 neurons. When placing the densely connected recurrent and feed-forward layers in the Spiking RelNet onto Loihi, the hardware constraints on network connectivity (see Methods, Supplement) mean that we can place at-most 128 neurons per neuro-core (less than the maximum possible 1024). Our most resource efficient placement hence requires 2308 neuro-cores spread across 22 chips. Placing the network of this size onto Loihi brings with it the challenge of minimizing spike congestion. This happens primarily when we route spikes from the LSNNs that do the embedding (module B, Fig. 4), to the various instances of g_θ (module C, Fig. 4). A straightforward placement leads to excessive cross-chip spike transmission. This leads to significant delays in spike transmission, which slows down the computation. Therefore separate relay neuro-cores (marked green in Fig. 5 C and D), and an optimized allocation of instances of g_θ onto chips, were introduced for reducing across-chip spike transmission. This resulted in significant improvements of the EDP, see Fig. 5 E. The final optimized layout of the network over the chips can be see in Fig. 5 C and Supplement.

Another aspect of optimization concerns the number of time steps used, called the compute time, which affects not only the energy consumed and delay on Loihi, but also the training speed. We found that using spiking neurons without refractory period and membrane time constants of just 7 ms significantly reduced the required number of time steps, while causing only a mild decrease in accuracy.

Energy-efficiency of RelNet in spike-based neuromorphic hardware

We compared the energy consumption and delay of the spike-based implementation of RelNet on Loihi with GPU implementations of the ANN RelNet from [Santoro et al., 2017], see Table 1. One sees that the spike-based implementation consumes between 4 and 16 times less energy than the GPU implementation. The energy savings are lower for longer story sizes, apparently because these require the use of substantially more Loihi chips, and inter-chip communication appears to be less energy efficient in this spike-based hardware. One should also note that the average length of a story for the 16 datasets that we consider is just 6.5 sentences. The computation time on Loihi was slightly larger than on the GPU. But nevertheless, the resulting EDP remained lower for Loihi. For the

	sMNIST		Relational reasoning				
	GPU	CPU	GPU				
# cores on Loihi	1	1	124	332	700	1552	2320
# sentences (RR)	-	-	2	6	10	16	20
Energy ratio	7.467x	4.774x	16.49x	11.92x	7.78x	5.32x	4.36x
Latency ratio	2.82x	5.89x	0.73x	0.56x	0.44x	0.33x	0.38x
EDP ratio	21.026x	28.134x	12.10x	6.73x	3.41x	1.73x	1.67x

All ratios are shown against Loihi.

Table 1. Benchmarking results. Comparison and scaling analysis of the spiking relational network on Loihi against the corresponding ANN on CPU and GPU[3]. For the scaling analysis of the RelNet the data set was grouped by number of sentences per sample which in turn determines the number of configured LSNNs and therefore cores per sample. All measurements were done using 250 input samples, except for network size 16 where only 100 samples were used, as there are not enough test samples containing 16 sentences. The energy per inference was calculated using total power values. More detailed results can be seen in the Supplement.

longest and therefore slowest story size the average computation time per sample is 6.54 ms wall-clock time, which would still be sufficient for online applications like voice control or virtual assistants.

Discussion

We have shown that a key tool for sequence processing in recurrent neural networks in machine learning and AI, LSTM units, can be replaced in spike-based neuromorphic hardware by neurons with a biologically inspired mechanism for spike frequency adaptation (SFA). SFA was achieved -similarly as in the brain- through spike-triggered hyperpolarizing currents on the time scale of seconds. Since neurons with SFA can also be used for generic network computations, this solution does not require a separation of units for computing and working memory, hence it can be viewed as an in-memory computing solution for the case of working memory. Like other in-memory computing solutions it comes with the benefit of avoiding latencies and energy consumption that generally arises from traffic between computing and memory units. The resulting spike based solution for solving a benchmark time series classification tasks such as sMNIST turns out to be three orders of magnitudes more energy efficient than state-of-the-art implementations of LSTM networks on CPUs and GPUs, while achieving virtually the same performance. This property could be especially interesting to low latency processing of real-time workloads.

We have also shown that this method enables us to port large ANNs that involve LSTM units into spike-based hardware. We have demonstrated this for the example of relational networks, since these enable a qualitative jump in AI capabilities by supporting reasoning about relationships between objects in a story or image. An essential feature of our spike-based emulation of LSTM networks is that these networks can be trained very effectively through BPTT, like LSTM networks. In particular, the implementation of the RelNet on the neuromorphic chip Loihi achieved almost the same performance as the ANN counterpart. The resulting reduction of energy consumption for relational reasoning is less drastic as for the time series classification task sMNIST, because the relational network contains also a large fraction of feed-forward neural network modules. But we have shown that the feed-forward network modules can be organized through suitable output encoding and regularization mechanisms so that they not only interact seamlessly with the recurrent neural network modules, but also compute with very few spikes per neuron, thus in an event-based rather than rate coding regime. We believe that the energy efficiency of resulting spike-based feed-forward modules can be increased by more dedicated hardware. In that respect, RelNet appear to represent a more suitable target for implementing large AI networks in energy-efficient neuromorphic hardware than CNNs. Similar to [Santoro et al., 2017], we expect that relational networks in neuromorphic hardware can be used not only for solving question-answering tasks in natural language, but also for reasoning about relations between objects

in an image or in an auditory scene. This would provide a qualitative jump in AI capabilities of energy-efficient neuromorphic hardware.

Another interesting next step will be to enable on-chip training of these spike-based emulations of LSTM networks by using e-prop instead of BPTT, which has already been shown to work very well for networks of spiking neurons with SFA [Bellec et al., 2020]. Also one-shot learning capability has been demonstrated for these spiking networks [Scherr et al., 2020], and it is likely that the required method will also enable one-shot on-chip training of these networks.

Finally, spiking neurons with SFA are a first step in the direction of state-of-the-art point neuron models for neurons in the neocortex [Billeh et al., 2020]. Hence, if our emulation of neurons with SFA can be expanded towards these more general GLIF (generalized leaky integrate-and-fire) neuron models, it will become possible to emulate state-of-the-art models for parts of the neocortex in large energy efficient neuromorphic systems, thereby providing a new venue for simulating large neural networks of the brain at a substantially reduced energy cost. This would be an important breakthrough for the scientific analysis of these data-driven brain models that is currently starting. These perspectives point to a significant advantage of neuromorphic hardware such as Loihi or SpiNNaker [Furber et al., 2014] that supports the implementation of variations of the standard spiking neuron model as they arise in further work towards spike-based AI or neuromorphic implementations of large-scale data-driven models for neural networks of the brain.

Methods

LIF neuron model with after-hyperpolarizing (AHP) currents

The dynamical behavior of an LIF Neuron with after-hyperpolarizing (AHP) currents (indexed by j), as implemented in Loihi, is given by Eq. 1–8. Here we show the dynamic interaction between incoming spikes at time t, the resultant input postsynaptic current (PSC) $i_{\text{PSC},j}[t]$, the internal AHP current $i_{\text{AHP},j}[t]$, the membrane voltage $V_j[t]$, and the output spikes $z_j[t+1]$. The equations are explained subsequently

$$i_{\text{PSC},j}[t+\Delta t] = \alpha_I \, i_{\text{PSC},j}[t] + \sum_i w_{ij} \, z_i[t-d_{ij}] \tag{4}$$

$$i_{\text{AHP},j}[t+\Delta t] = \alpha_{AHP} \, i_{\text{AHP},j}[t] - \beta \, z_j[t] \tag{5}$$

$$V_j[t+\Delta t] = \begin{cases} \alpha_V V_j[t] + \frac{1}{g_V}(i_{\text{PSC},j}[t+\Delta t] + i_{\text{AHP},j}[t+\Delta t]) & \text{if neuron is not refractory} \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

$$z_j[t+\Delta t] = \begin{cases} 1 & \text{if } V_j[t+\Delta t] > b_0 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

$$V_j[t+\Delta t] \to 0 \text{ if } z_j[t+\Delta t] = 1 \tag{8}$$

Eq. 4–6 are represent temporal convolution with a exponentially decaying kernal. Here $\alpha_I = e^{-\frac{1}{\tau_I}}$, $\alpha_{AHP} = e^{-\frac{1}{\tau_{\text{AHP}}}}$, $\alpha_V = e^{-\frac{1}{\tau_V}}$, where τ_I, τ_{AHP}, and τ_V are the decay constants of the corresponding exponentials. β is update to the AHP current in response to an output spike. Since the LIF state transition is computed in Loihi, we associate a single compute step with 1ms biological time, correspondingly $\Delta t = 1$ms and $g_V = 1$.

Eq. 4 defines the PSC as a function of input spikes arriving through incoming synapses of weights w_{ij} and delays of d_{ij} steps. The LIF neuron without AHP currents corresponds to the case of $\beta = 0$, where the neuron performs a leaky-integration of $i_{\text{PSC},j}[t]$ to get the membrane voltage $V[t]$. When this voltage exceeds a threshold, it is reset to zeros and an output spike is generated. In this case, the memory of the neuron is limited by the voltage and PSC decay time constants τ_V and τ_I respectively, which are typically around 20ms. This means that even when connected in a recurrent fashion, the memory capacity for the network is typically at-most a 100ms.

Eq. 5 defines the AHP current. With $\beta > 0$, each output spike i.e. $z_j[t] > 0$ will cause $i_{\text{AHP},j}[t]$ to become more negative by a value of β. When leaky-integrated into the membrane voltage $V[t]$ (Eq. 6), this increased negative value of $i_{\text{AHP}}[t]$ lowers the rate of subsequent spikes, leading to *spike frequency adaptation*. The decay time constant of the AHP current τ_{AHP} is much longer than τ_V, τ_I, typically > 100ms. The slow decay means that this inhibition persists over a much longer duration thus functioning as a longer-term memory cell. This longer lasting memory proves invaluable to solve the complex tasks demonstrated in this work.
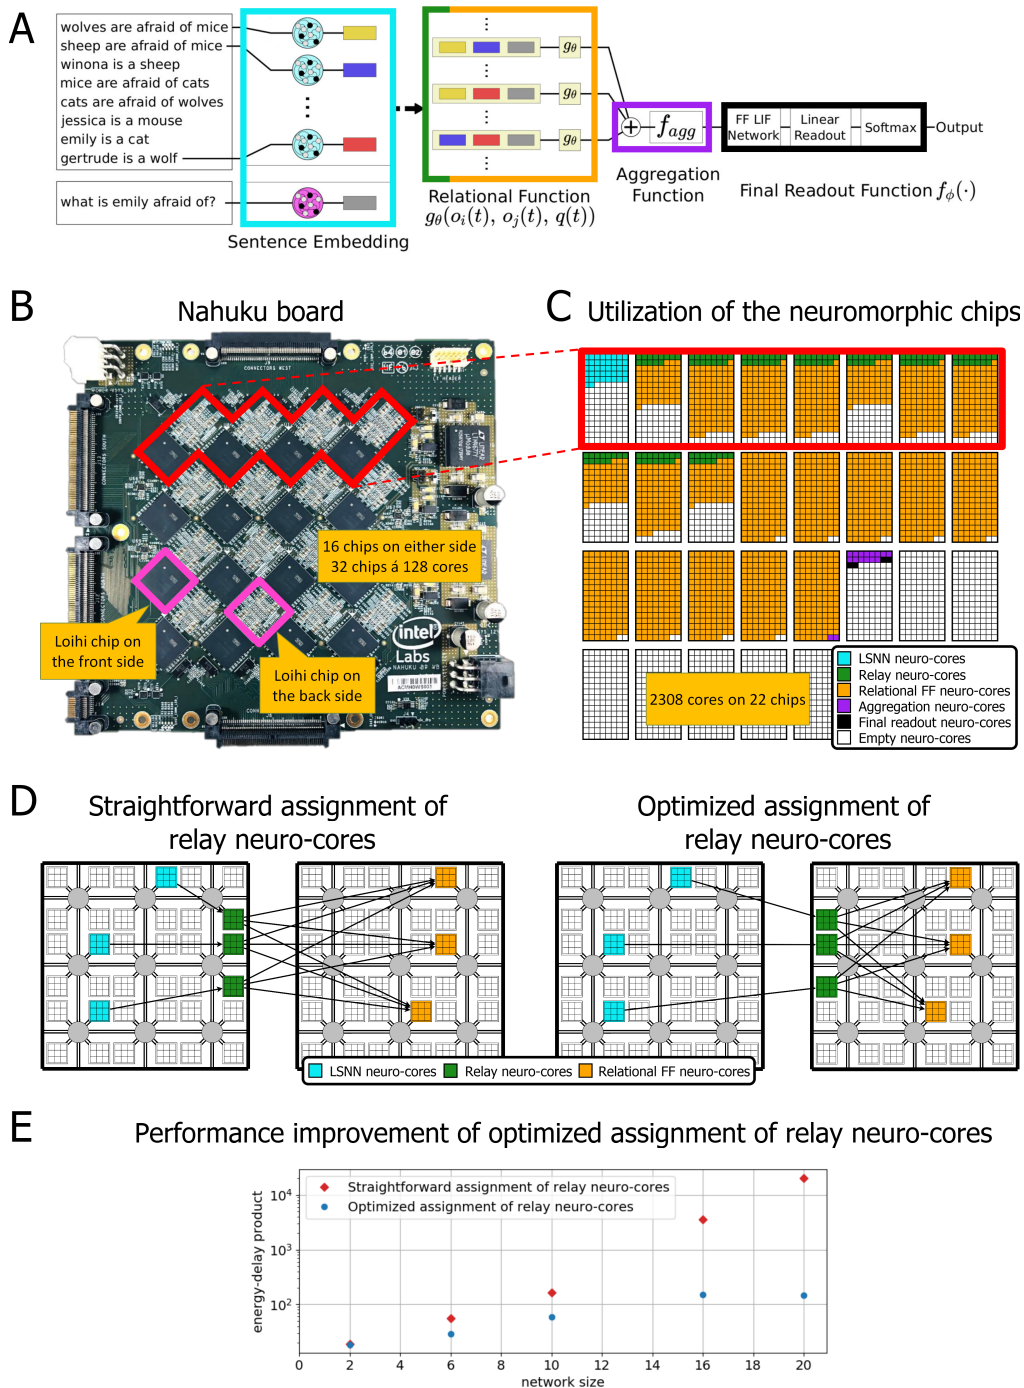
Figure 5. Spiking RelNet placement and optimization on Loihi A) Highlighting the different parts of the Spiking RelNet with the color code used in C and D. **B)** The Spiking RelNet was configured on a Nahuku board with 32 Loihi chips. On each side of the board 16 chips are placed in a checkerboard pattern. Each chip has 128 neuromorphic computing units or neuro-cores. **C)** The full scale Spiking RelNet which can solve tasks up to 20 sentences utilizes 2308 neuro-cores on 22 chips. The detailed mapping of the different layers is shown. In order to minimize cross-chip spikes some chips were not fully utilized. **D)** Straightforward assignment of relay neuro-cores on the same chip as the source neuro-cores on the left side and optimized assignment of the relay neuro-cores on the chip of the target neuro-cores, to minimize inter-chip spike traffic, on the right side. Note that the target neuro-cores must be assigned carefully to do this efficiently. **E)** Shows the benefit of the optimized assignment in energy-delay product as a function of network size as measured by number of bAbI sentences which roughly corresponds to the number of Loihi chips for the RelNet solving bAbI tasks on Loihi. Thus we see that the placement of the Spiking RelNet is optimized in terms of resource utilization as well as network performance on the hardware.

Details of LIF network training

In this section we describe important details pertaining to the training of networks of LIF neurons with and without AHP currents. In all equations below, we drop the neuron index j for brevity.

The scaled voltage

For subsequent details regarding LIF network training, we find it useful to define a normalized version of the membrane voltage i.e. a *scaled voltage* v_s.

We first notice that the membrane voltage $V[t]$ is a sum of two voltage components $V_{\mathrm{PSC}}[t]$ and $V_{\mathrm{AHP}}[t]$ which are a result of leaky-integrating i_{PSC} and i_{AHP} respectively. Correspondingly, we can rewrite Eq. 6,8 describing the voltage evolution as follows:

$$V_{\mathrm{PSC}}[t + \Delta t] = \begin{cases} \alpha_V\, V_{\mathrm{PSC}}[t] + \frac{1}{g_V} i_{\mathrm{PSC}}[t + \Delta t] & \text{if neuron is not refractory} \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

$$V_{\mathrm{AHP}}[t + \Delta t] = \begin{cases} \alpha_V\, V_{\mathrm{AHP}}[t] + \frac{1}{g_V} i_{\mathrm{AHP}}[t + \Delta t] & \text{if neuron is not refractory} \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

$$V[t + \Delta t] = V_{\mathrm{PSC}}[t + \Delta t] + V_{\mathrm{AHP}}[t + \Delta t] \tag{11}$$

$$\begin{aligned} V_{\mathrm{PSC},j}[t + \Delta t] &\to 0 \\ V_{\mathrm{AHP},j}[t + \Delta t] &\to 0 \end{aligned} \quad \text{if } z_j[t + \Delta t] = 1 \tag{12}$$

The scaled voltage $v_s[t]$ is defined below:

$$v_s[t] = \frac{V[t] - b_0}{b_0 - V_{\mathrm{AHP}}[t]} \tag{13}$$

$v_s[t]$ takes the value of 0 when $V[t] = b_0$ and a value of -1 when $V[t] = V_{\mathrm{AHP}}[t]$. This is motivated by the fact that $V_{\mathrm{AHP}}[t]$ is the value that $V[t]$ would take if there was no input PSC.

The surrogate gradient

The generation of spikes from the membrane voltage (Eq. 7), involves the use of a step function centered at the neuron threshold. This function is non-differentiable at the neuron threshold and provides a non-informative gradient of zero at all other points. Thus in order to use gradient back-propagation to train networks of LIF Neurons, we consider a surrogate gradient for the step function similar to methods used in previous works [Bellec et al., 2018b, Zenke and Vogels, 2021, Esser et al., 2016, Shrestha and Orchard, 2018, Neftci et al., 2019, Zenke and Ganguli, 2018, Zhu et al., 2021].

We rewrite the thresholding equation Eq. 7 in terms of the scaled voltage $v_s[t]$:

$$z[t + \Delta t] = h\left(v_s\right) \equiv h\left(\frac{V[t] - b_0}{b_0 - V_{\mathrm{AHP}}[t]}\right) \tag{14}$$

where h is the unit step function and $v_j[t]$ is referred to as the *scaled voltage*.

We then use the following piece-wise linear surrogate gradient function to serve as a pseudo-derivative of the step function $h(\cdot)$.

$$\frac{dh}{dv_s} \triangleq \begin{cases} \gamma\left(1 + \frac{v_s}{v_-}\right) & \text{if } -v_- \leq v_j < 0 \\ \gamma\left(1 - \frac{v_s}{v_+}\right) & \text{if } 0 \leq v_j \leq v_+ \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

where v_- and v_+ define the support of the surrogate gradient and γ is a dampening factor that affects the magnitude of the derivative.

Thus $h'(v_j)$ peaks at a value of γ for $v_j = 0$ and linearly decays to zero at the values of $-v_-$ and v_+.

Spike rate regularization

For each neuron k, we calculate the mean rate $\bar{\rho}_k$ across all batches. we then add the following regularization loss

$$L\rho = \lambda_\rho \left(\sum_j (\bar{\rho}_k - \rho_{target})^2\right)^2 \tag{16}$$

where ρ_{target} is a target rate and λ_ρ is the parameter that controls the strength of the regularization. This loss encourages the mean spike rate of each neuron across a random batch to be as close to the target rate ρ_{target}. This ensures that the network activity does not die out and that the spike rate stays sparse owing to the low value of ρ_{target}. The outermost square is in order to dynamically reduce the regularization strength as the loss becomes smaller.

When training the Spiking RelNet, we use a more aggressive spike rate regularization to limit the total spike rate across all across the instances of the relational function g_θ. This is described below in the section on the training of the Spiking RelNet.

Voltage regularization

The spike rate regularization has a tendency to push the synaptic weights low enough that the membrane voltages become very negative. This leads to a large number of time steps where the voltage values fall outside the support of the surrogate gradient and thus no gradient information can be propagated through them, which impedes gradient back propagation. Thus we are motivated to add a loss that penalizes voltages that fall significantly outside the support of the surrogate gradient function defined in Eq. 15. Since the surrogate gradient is defined in terms of the scaled voltage v_s (Eq. 13), we define the voltage regularization loss in terms of it as well.

For each neuron j and time step n, we calculate the loss component

$$L_v^{(j,n)} = (\text{relu}(v_{s,j}[n] - 0.4))^2 + (\text{relu}(-v_{s,j}[n] - 2.0))^2 \tag{17}$$

The total voltage regularization loss is given by

$$L_v = \lambda_v \left(\text{mean}_{i,n} L_v^{(i,n)} \right)^2 \tag{18}$$

The above penalizes all neurons at all time instants that the scaled voltage v_s goes outside the range $[-2.0, 0.4]$. This prevents the network from using voltages that are excessively negative and increases the proportion of voltage values that lie within the support of the surrogate-gradient. Moreover, limiting the range of the voltage values is also crucial in order to be able to fit the voltage values onto the range offered by the fixed precision registers on Loihi.

Use of PSC kernels

For the LIF neuron model, the membrane voltage resets to zero upon spiking, and stays zero for the duration of the refractory period. This means that the gradient cannot propagate through the membrane voltage beyond the last spike. For a LIF Neuron with AHP currents this issue is alleviated by the slow decay of the AHP current through which gradients can be propagated much further into the past. However, for the feed-forward layers used in the relational network, which don't use the AHP current, we need to make use of the PSC to propagate gradients, as it is unaffected by the spiking of the neuron. We thus find that the use of a non-zero PSC decay time constant τ_I, i.e. an exponentially decaying PSC, offers improved performance upon training compared to using a delta PSC. The use of a non-delta PSC means that a change in the weight of an input synapse changes the rate at which the membrane potential rises and therefore has the capacity to smoothly modify the spike time.

Details for the application to sMNIST

Input encoding

The gray values of the pixels from an MNIST image were encoded in spikes. 80 input neurons were used and each pixel was associated with a particular threshold for the gray value. So there were 79 linear spaced thresholds between 0 and 256. Every second threshold refereed to an increasing gray value, while the others refereed to a decreasing gray value. If the gray value increases when transitioning from one pixel to the next, every second input neuron from the last threshold to the next threshold generates a spike. The pseudo code for the input encoding can be seen in the Supplement. The last input neuron becomes active after the presentation of all 784 pixels for 56ms, thus the presentation of one sample takes 840ms. This last input neuron which generates a spike at every time step after the image presentation indicates the end of an sample. The classification happened at the last time step i.e. time step 840 of a sample. Each of the 10 output neuron denoted a digit and the neuron with the highest membrane potential on the last time step defined the predicted class.The network was implemented on the Intel Loihi chip using NxNet API from the NxSDK v0.95.

Network structure

An LSNN was used consisting of 240 neurons, 180 excitatory and 60 inhibitory. A random subset of 100 of the excitatory neurons were equipped with AHP currents. Additionally 80 input neurons were used to perform an input spike encoding of the images, and 10 output neurons were used corresponding to the 10 classes of the MNIST dataset. The overall connectivity of the network, including the input and output connectivity, was kept at 20%, meaning that only 20% of the possible synapses between the neurons were used. This was achieved by using a rewiring technique named DEEP-R [Bellec et al., 2018a] during training. The hyper-parameters which were used to train the network for Loihi were $\beta = 96$, baseline threshold $b_0 = 127$, $\tau_V = 20$, $\tau_{adap} = 700$ as well as a refractory period and delay of 1ms.

Details for the Spiking RelNet

In this section, we describe in detail the structure of the Spiking RelNet as applied to the bAbI tasks.

High-level network outline

Building on the general architecture proposed in [Santoro et al., 2017], the Spiking RelNet takes as input K objects $o_i(t)$, and a question object $q(t)$ and implements the following function to compute its output.

$$RN\left([o_1(t), o_2(t), \ldots, o_K(t)], q(t)\right) = f_\phi\left(f_{agg}\left(\sum_{1 \leq i \leq j \leq K} g_\theta\left(o_i(t), o_j(t), q(t)\right)\right)\right) \tag{19}$$

Fig. 4 A shows the basic outline of this network. When applied to the bAbI task, the sentences of the story and the question are embedded into the spike sequences $o_i(t)$ and $q(t)$ respectively by means of LSNNs, which are recurrent networks consisting of LIF Neurons both with and without AHP currents. To provide the input to the LSNN, we assign an input neuron corresponding to each distinct word used in the bAbI dataset. The words in a sentence/question are then presented in sequence, with each word being presented for a duration of $T_{\text{word}} = 10$ms during which only the corresponding input neuron fires continuously. We then take the spike activity of the LSNN over the last $T_{\text{inp}} = 14$ms, and pad it to a length of $T_{\text{sim}} = 37$ms to form the embedding spike sequences $o_i(t)$ and $q(t)$ (see Fig. 4 B).

The function g_θ is the relational function. It receives as input a triplet of spike sequences $(o_i(t), o_j(t), q(t))$ corresponding to a pair of sentences and questions, and produces a spike sequence output It is implemented as a four layer feed-forward spiking neural network with LIF neurons. We have an instance of g_θ for each pair of sentences $i, j :: i \leq j$, so that the ordering of the sentences in the stories is made available to the network.

The function f_{agg} is an element-wise function. This is implemented by means of a LIF layer to which each instance of g_θ is connected one-to-one, where the set of input weights from an instance of g_θ to this layer is shared across all instances. This is an addition to the architecture proposed in [Santoro et al., 2017] and plays an essential role in enabling the implementation of Spiking RelNets onto neuromorphic hardware.

The function f_ϕ is the readout function. It is implemented as a 3-Layered feed-forward LIF network followed by a linear readout (see section below) and a softmax layer. f_ϕ outputs, for each unique word present in the bAbI dataset, the probability of that word being the answer to the question. This probability is used to compute a cross-entropy loss that is used to train the network using gradient back-propagation.

The LIF neurons used in g_θ, f_{agg} and f_ϕ don't use AHP currents. For more detailed parameters pertaining to the layers, see Supplement.

The linear readout

The design of the linear readout is crucial to the performance of the relational network. The linear readout consists of a network of specialized readout neurons, with one neuron for each word in the database of words used in the bAbI task.

The readout neuron is a variant of the LIF Neuron without AHP currents, where $\tau_I = \tau_{\text{readout}} = 7.0$ms and $\tau_V = \infty$, and the threshold $b_0 = \infty$. This corresponds to a neuron which does not spike, but where the PSC decays with the readout time constant τ_{readout} and the neuron performs (non-leaky) integration of the PSC to calculate the membrane potential. However, we chose to enable the integration of the PSC into the voltage only $T_{\text{readout}} = 10$ms prior to the final step. The value of the membrane voltage at the final step is scaled by a fixed scalar and forms the input to the softmax (see Fig. 4E). This design incentivizes the spike activity of the final layers to occur in a confined time window close to the final time step, while allowing the precise timing of the spike to influence the final output, leading to a high information capacity in a short time window.

Training the relational network

The Spiking RelNet requires many time steps of compute time compared to the non-spiking RelNet, making the loss computation and gradient back-propagation through time many times more expensive in the spiking case. Simply training the network end-to-end with the cross-entropy loss requires an impractically long time for the network to converge, as well as leading to pathological spike rates and low performance. The solutions to these issues are:

In order to speed convergence, We first train a non-spiking relational network to solve the bAbI tasks, where LSTM's are used to embed the questions and words. We then train the LSNN to reproduce the outputs of the LSTMs for the various input sentences in the dataset. The weights of these pre-trained LSNNs are fixed, and they are used to perform the embedding while we train the relational function g_θ and readout function f_ϕ. This helps the network converge in much fewer training epochs than the end-to-end trained non-spiking relational network. This makes the training feasible for a Spiking RelNet.

The emergence of pathological spike rates and membrane voltage values is solved by the use of spike rate regularization and membrane voltage regularization described above. We use a more aggressive regularization for the spike rates in the instances of the relational function g_θ, where the regularization forces the total spike rate summed over all instances of g_θ towards a low target rate. This minizes the number of spikes transmitted to the aggregation layer f_{agg}, thus reducing cross chip transfer of spikes. It also forces the network to only generate spikes corresponding to those sentence pairs relevant to the question. The resultant low spike rate seen in Fig. 3B results in a very power and delay efficient implementation of feed-forward spiking networks onto neuromorphic hardware.

Placement of the Spiking RelNet onto Loihi

The Loihi Nahuku board consists of 32 interconnected Loihi chips, each of which contains 128 neuro-cores. The neuro-core is the fundamental computational unit that computes the dynamics of the LIF neurons with and without AHP currents. Loihi allows one to connect any neuro-core on any chip, to any other neuro-core on any other chip thus enabling large networks to be placed on the board. However due to hardware limitations, the number of connections and the connectivity is constrained as described below. Additionally, transporting a large number of spikes across different chips incurs significant latency. We discuss here the strategies to place the Spiking RelNet within these constraints.

The LSNN network that solves the sMNIST task contains 240 neurons connected with 20% of the recurrent connections enabled. This network is small enough to fit in a single chip and occupy only one neuro-core. The Spiking RelNet is a much larger network. Considering a maximum of $M = 20$ sentences in a story, the Spiking RelNet has M instances of the LSNN's that embed sentences, plus one for the questions. Additionally, there exists an instance of the relational function g_θ for each pair of sentences $o_i(t), o_j(t) :: i \leq j$, making a total of $\frac{M(M+1)}{2} = 210$ instances. Each of these instances is implemented as a separate network on Loihi, leading to a total network size of $238,604$ neurons. The placement of this network needs to take into consideration many constraints regarding connectivity, memory, and the latency of spike transport. The associated challenges and solutions are outlined in this section.

Synaptic memory limit

Each neuro-core has a limited amount of SRAM memory which can be used to store synaptic parameters. This limits the number of incoming synapses to a particular neuro-core. The precise number is dependent on synaptic parameters and we have found an empirical limit of around 40000 synapses per neuro-core. Except the aggregation layer, all layers in the network have dense input and recurrent synaptic connections. Thus each layer needs to be placed over multiple neuro-cores in order to store the input and recurrent connections.

Fanout limits – LSNN relay layer

The total number of neuro-cores to which the neurons of a neuro-core connect to is limited to 2048, and 4096 for intra-chip connections. This plays a role when connecting the LSNNs to the large number of instances of the relational function g_θ.

Thus, one can split the neurons across multiple neuro-cores to reduce the number of output connections per neuro-core. However splitting a recurrent LSNN network across too many neuro-cores increases latency. Instead we use relay layers. A *relay layer*, as the name suggests, simply reproduces the spiking activity of the layer that forms its input. Each LSNN is thus connected to multiple relays which then each fanout to a smaller number of instances of g_θ.

Limits pertaining to fanin – The aggregation layer

For any neuro-core C, Loihi limits the number of neurons that can be connected to that neuro-core to 4096. Unlike the two constraints above, this constraint on the fanin to a neuro-core introduces a fundamental restriction to the

network architectures that can be implemented on Loihi.

The layer that receives the output from the instances of g_θ receives input from $\frac{M(M+1)}{2} = 210$ instances. For this layer to not violate the fanin constraint, the connection from the output of g_θ to this layer must be sparse. Thus, we introduce an *aggregation layer* to which each instance of g_θ is connected in a sparse one-to-one manner, with shared weights across instances. The sparse connection enables the aggregation layer to be implemented within the fanin constraints

For a more detailed treatment of the constraints, as well as the number of neuro-cores required to place each layer, see Supplement.

Optimizing network placement to minimize congestion in cross chip spike transport

Placing the LSNN, relay, the relational networks taking into consideration only the connectivity constraints, we notice significant delays that occur due to transporting spikes from the LSNN and relay networks to the instances of g_θ. This is oweing to the large number of spikes that need to be transferred across different Loihi chips. Thus, we need in addition to optimize the placement of the instances of g_θ, and the relay networks in a manner that minimizes cross-chip spike transport. We break down this general objective into the following constraints.

- All relay networks must be connected only to relational function instances that are placed on the same Loihi chip. We thus choose to place the initial layer of the relational function instances in the same chip as the relay networks that give them their input.

- We aim to minimize the number of relay networks required. Each chip has a limit of 128 neuro-cores and thus a limit on the number of g_θ instances that can be placed. This means that for each chip, we must choose the set of g_θ instances in such a manner that the number of distinct sentences needed as input is minimized.

- For each g_θ instance, all layers after the first one are to be placed on the same chip.

The layout that we arrive at with the above principles is described in the Supplement. The resultant improvement in delay and the corresponding energy delay product is shown in Fig. 5E.

Data availability

The MNIST dataset [LeCun et al., 2010] is freely available at http://yann.lecun.com/exdb/mnist/. The bAbI dataset [Weston et al., 2015] is freely available at https://research.fb.com/downloads/babi/.

Code availability

The Loihi source code is freely available from Github (https://github.com/intel-nrc-ecosystem/models/tree/master/nxsdk_modules_ncl/lsnn/apps/smnist and https://github.com/intel-nrc-ecosystem/models/tree/master/nxsdk_modules_ncl/lsnn/apps/relnet)

References

Bellec et al., 2018a. Bellec, G., Kappel, D., Maass, W., and Legenstein, R. (2018a). Deep rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*.

Bellec et al., 2018b. Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., and Maass, W. (2018b). Long short-term memory and learning-to-learn in networks of spiking neurons. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 795–805. Curran Associates, Inc.

Bellec et al., 2020. Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., and Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(1):3625.

Billeh et al., 2020. Billeh, Y. N., Cai, B., Gratiy, S. L., Dai, K., Iyer, R., Gouwens, N. W., Abbasi-Asl, R., Jia, X., Siegle, J. H., Olsen, S. R., Koch, C., Mihalas, S., and Arkhipov, A. (2020). Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron*, 106(3):388–403.e18.

Davies et al., 2018. Davies, M., Srinivasa, N., Lin, T., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., Liao, Y., Lin, C., Lines, A., Liu, R., Mathaikutty, D., McCoy, S., Paul, A., Tse, J., Venkataramanan, G., Weng, Y., Wild, A., Yang, Y., and Wang, H. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *micro*, 38(1):82–99.

Davies et al., 2021. Davies, M., Wild, A., Orchard, G., Sandamirskaya, Y., Guerra, G. A. F., Joshi, P., Plank, P., and Risbud, S. R. (2021). Advancing neuromorphic computing with loihi: A survey of results and outlook. *Proceedings of the IEEE*, 109(5):911–934.

Esser et al., 2016. Esser, S. K., Merolla, P. A., Arthur, J. V., Cassidy, A. S., Appuswamy, R., Andreopoulos, A., Berg, D. J., McKinstry, J. L., Melano, T., Barch, D. R., et al. (2016). Convolutional networks for fast, energy-efficient neuromorphic computing. *Proceedings of the national academy of sciences*, 113(41):11441–11446.

Furber et al., 2014. Furber, S. B., Galluppi, F., Temple, S., and Plana, L. A. (2014). The spinnaker project. *Proceedings of the IEEE*, 102(5):652–665.

Hochreiter and Schmidhuber, 1997. Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

LeCun et al., 2010. LeCun, Y., Cortes, C., and Burges, C. (2010). Mnist handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, 2.

Neftci et al., 2019. Neftci, E. O., Mostafa, H., and Zenke, F. (2019). Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63.

Santoro et al., 2017. Santoro, A., Raposo, D., Barrett, D. G., Malinowski, M., Pascanu, R., Battaglia, P., and Lillicrap, T. (2017). A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976.

Scherr et al., 2020. Scherr, F., Stöckl, C., and Maass, W. (2020). One-shot learning with spiking neural networks. *bioRxiv*.

Shrestha and Orchard, 2018. Shrestha, S. B. and Orchard, G. (2018). Slayer: Spike layer error reassignment in time. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

Weston et al., 2015. Weston, J., Bordes, A., Chopra, S., Rush, A. M., van Merriënboer, B., Joulin, A., and Mikolov, T. (2015). Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.

Zenke and Ganguli, 2018. Zenke, F. and Ganguli, S. (2018). Superspike: Supervised learning in multilayer spiking neural networks. *Neural computation*, 30(6):1514–1541.

Zenke and Vogels, 2021. Zenke, F. and Vogels, T. P. (2021). The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks. *Neural Computation*, 33(4):899–925.

Zhu et al., 2021. Zhu, X., Zhao, B., Ma, D., and Tang, H. (2021). An efficient learning algorithm for direct training deep spiking neural networks. *IEEE Transactions on Cognitive and Developmental Systems*.

Acknowledgements

Author contributions statement

A.R., P.P. and W.M. contributed to the design and planning of the experiments. A.R. and P.P. carried out the experiments. A.R., P.P., A.W. and W.M. participated in the analysis of the experimental data. A.R., P.P., A.W. and W.M. wrote the manuscript.

Competing interests

The authors declare competing interests as follows. P.P. and A.W. are currently employed by Intel Labs, developers of the Loihi neuromorphic system. W.M. and A.R. are members of the Intel Neuromorphic Research Community and W.M. has received research funding from Intel for related work.

Additional information

Supplementary information is available.

Supplementary information

A Long Short-Term Memory for AI Applications in Spike-based Neuromorphic Hardware

Arjun Rao, Philipp Plank, Andreas Wild, Wolfgang Maass

July 9, 2021

S1 Supplementary results

S1.1 Energy and time benchmarking

In order to evaluate the energy efficiency of our spiking networks implemented on the neuromorphic chip Loihi from Intel [Davies et al., 2018] we measured the energy and execution time of a task and compared it with the artificial neuronal network implementation on conventional hardware, i.e., CPUs and GPUs.

For the network executed on a Loihi system the performance is measured by reading out sensors of the hardware. Regarding the energy measurement, Loihi chips of a system are powered by on-board voltage regulators that support power telemetry over an I²C interface. These voltage regulators are used to collect power usage information. In particular, the SDK of Loihi allows to split the contributions of static and dynamic power consumption as well as estimate the contribution of neuro-cores and on-chip synchronous x86 cores to the overall power consumption.

For Energy measurement on CPU the Intel Power Gadget 3.5, a software based power estimation tool, was used. For GPUs we used the nvidia-smi tool to measure the power, which is also a software based estimation tool. The nvidia-smi tool does not give a detailed breakdown of where power is consumed, but rather report the power draw of the whole board. Therefore we measured a baseline idle power draw, which we considered for the static energy. Afterwards we measure the power consumption during the workload, which denotes to the total energy and then we calculated the dynamic energy by subtracting the static energy from the total energy.

For both Loihi and CPU/GPU the measurements were performed for a workload running long enough to get in a steady state for power draw. Therefore, the batch size 50 and 100 examples on the GPU required us to run the test set several times to achieve a steady state. The execution time for networks executed on Loihi and networks running on CPU or GPU was measured on python level using the timeit module and can be considered a wall-clock time. This wall-clock time is then divided by the number of samples used for inference to calculate the latency. The execution time was measured independent of the power measurements.

In Table S1 and Table S2 are given the detailed results of our benchmarking efforts for the sMNIST task and the Spiking RelNet respectively.

Hardware	# cores		Power (mW)			Time per time step (μs)	Latency (ms)	Latency ratio	Energy per Inference (mJ)	Energy ratio	Energy Delay Product (μJs)	EDP ratio
			Static	Dynamic	Total							
		x86 cores	0.08	24.33	24.41				0.34			
Loihi	1	neuron cores	0.51	0.91	1.42	16.79	14.11	1.00x	0.02	1.00x	5.14	1.00x
		total	0.59	25.24	25.83				0.36			
Nvidia RTX 2070	-	batch size 1	33898.00	34602.00	68500.00	-	39.73	2.82x	2721.51	7,467.20x	108125.39	21,025.64x
		batch size 50	33898.00	38171.00	72069.00	-	37.44	2.65x	53.97	148.07x	2020.46	392.89x
		batch size 100	33898.00	53287.00	87185.00	-	40.41	2.86x	35.23	96.67x	1423.70	276.85x
Intel Core i5-7440HQ	-	batch size 1	2040.00	18886.00	20926.00	-	83.15	5.89x	1740.00	4,774.16x	144680.74	28,134.05x

Table S1: **Benchmark results for sMNIST.** Comparison of energy and time measurements of the spiking LSNN network on Loihi against the corresponding LSTM on GPU and CPU solving the sMNIST task[1].

[1]Loihi: Nahuku board (ncl-ghrd-01), CPU: Intel Core i9-7920X, RAM: 128GB, OS: Ubuntu 16.04.6 LTS, NxSDK: 0.95
Nvidia RTX 2070: Nvidia RTX 2070 Super, GPU-RAM: 8GB, CPU: Intel Core i7-9700K, RAM: 32GB, OS: Ubuntu 16.04.6 LTS, Python 3.6.5, TensorFlow-GPU: 1.14.0, CUDA: 10.0.
Intel Core i5-7440HQ: RAM: 16GB, OS: Windows 10 (build18362), Python 3.6.7, TensorFlow: 1.14.1
Performance results are based on testing as of July 9, 2021 and may not reflect all publicly available security updates. Results may vary.
[2]Loihi: Nahuku board (ncl-ghrd-01), CPU: Intel Core i9-7920X, RAM: 128GB, OS: Ubuntu 16.04.6 LTS, NxSDK: 0.95
GPU: Nvidia RTX 2070 Super, GPU-RAM: 8GB, CPU: Intel Core i7-9700K, RAM: 32GB, OS: Ubuntu 16.04.6 LTS, Python 3.6.5, TensorFlow-GPU: 1.14.0, CUDA: 10.0.
Performance results are based on testing as of July 9, 2021 and may not reflect all publicly available security updates. Results may vary.

Hardware	# sentences # cores		Power (W)			Time per time step (μs)	Latency (ms)	Latency ratio	Energy per Inference (mJ)	Energy ratio	Energy Delay Product (μJs)	EDP ratio
			Static	Dynamic	Total							
Loihi	20 2320 cores	x86 cores	0.01	0.44	0.44	45.73	6.54	1.00x	2.89	1.00x	148.47	**1.00x**
		neuron cores	1.89	1.14	3.03				19.82			
		total	1.90	1.57	**3.47**				**22.70**			
Nvidia RTX 2070	20	batch size 1	33.50	5.89	39.39	-	2.51	0.38x	98.88	4.36x	248.18	1.67x
		batch size 50	33.50	73.86	107.36	-	4.43	0.68x	9.51	0.42x	42.14	0.28x
		batch size 100	33.50	82.38	115.88	-	8.26	1.26x	9.57	0.42x	79.06	0.53x
Loihi	16* 1552 cores	x86 cores	0.00	0.43	0.43	55.89	7.99	1.00x	3.47	1.00x	151.77	**1.00x**
		neuron cores	1.16	0.78	1.94				15.52			
		total	1.17	1.21	**2.38**				**18.99**			
Nvidia RTX 2070	16	batch size 1	33.36	5.47	38.82	-	2.6	0.33x	100.94	5.32x	262.45	1.73x
		batch size 50	33.36	51.47	84.83	-	4.82	0.60x	8.18	0.43x	39.42	0.26x
		batch size 100	33.36	76.36	109.71	-	5.43	0.68x	5.96	0.31x	32.35	0.21x
Loihi	10 700 cores	x86 cores	0.01	0.44	0.45	36.36	5.20	1.00x	2.33	1.00x	58.73	**1.00x**
		neuron cores	0.86	0.86	1.72				8.96			
		total	0.87	1.30	**2.17**				**11.30**			
Nvidia RTX 2070	10	batch size 1	33.90	4.63	38.53	-	2.28	0.44x	87.86	7.78x	200.31	3.41x
		batch size 50	33.90	54.37	88.27	-	3.47	0.67x	6.13	0.54x	21.26	0.36x
		batch size 100	33.90	65.15	99.05	-	3.97	0.76x	3.93	0.35x	15.61	0.27x
Loihi	6 332 cores	x86 cores	0.01	0.45	0.46	27.64	3.95	1.00x	1.81	1.00x	29.11	**1.00x**
		neuron cores	0.42	0.99	1.41				5.56			
		total	0.43	1.44	**1.86**				**7.37**			
Nvidia RTX 2070	6	batch size 1	33.80	5.58	39.38	-	2.23	0.56x	87.82	11.92x	195.84	6.73x
		batch size 50	33.80	44.76	78.56	-	3.2	0.81x	5.03	0.68x	16.09	0.55x
		batch size 100	33.80	52.82	86.62	-	3.76	0.95x	3.26	0.44x	12.25	0.42x
Loihi	2 124 cores	x86 cores	0.01	0.46	0.47	22.96	3.28	1.00x	1.53	1.00x	18.36	**1.00x**
		neuron cores	0.17	1.07	1.24				4.06			
		total	0.18	1.53	**1.70**				**5.59**			
Nvidia RTX 2070	2	batch size 1	33.27	4.99	38.26	-	2.41	0.73x	92.20	16.49x	222.21	12.10x
		batch size 50	33.27	41.62	74.89	-	2.92	0.89x	4.37	0.78x	12.77	0.70x
		batch size 100	33.27	46.38	79.64	-	3.48	1.06x	2.77	0.50x	9.64	0.53x

Table S2: **Benchmark results for question-answering task using RelNet** Benchmarking comparison and scaling analysis of the Spiking RelNet on Loihi against the corresponding ANN on GPU[2]. The data set was grouped by number sentences per sample which in turn determines the number of LSNNs and therefore cores per sample. Measurements were done using 250 input samples. The energy per inference was calculated using total power values.

*For network size 16 only 100 input samples were used, as there are not enough test samples containing 16 sentences.

Output accuracy of the Spiking RelNet

The Spiking RelNet is trained on the combined data from 17 out of 20 bAbI tasks and it's performance is compared to an implementation of a non-spiking RelNet in Table S3. We have excluded the 3 tasks "Task 2: Two Supporting Facts", "Task 3: Three Supporting Facts", and "Task 16: Basic Induction". For an example of some of the tasks on which the network was trained, see Fig. S1. This is because the non-spiking RelNet did not converge on these tasks. The network we trained was able to solve $16/17$ tasks to within a 5% Error, which is the threshold at which a task is considered solved (used in [Santoro et al., 2017], and [Weston et al., 2015]). The task "Task 17: Positional Reasoning" has a rather high error, which we think is because the comparatively complex sentences in this task require a longer compute time to process. In Table S3, we show two additional columns. The first one corresponds to using $T_{sim} = 45$ ms. For this case, we pad the embeddings upto a longer T_{sim}, and increase the amount of time for which we run the feed-forward part of the relational network (i.e. modules C–E in main text Fig. 4). We observe here that while Task 17 is not yet under 5% error, the error has dropped significantly. The second one shows a simulation where all the time constants, refractory periods, T_{inp}, T_{sim}, and time per word in embedding (i.e. modules B–E in main text Fig. 4), are tripled. We see here that the additional temporal resolution allows all 17 tasks to be solved within 5% classification error.

S2 Supplementary Methods

S2.1 Input encoding for sMNIST

In Listing S1 the pseudo code for the input encoding used in the sMNIST task is shown. We assume that the current pixel value and next pixel value of the input image are presented, the number of thresholds were chosen to be half of the input neurons and thresholds are linearly spaced between 0 and 255 (number of threshold times).

```
while threshold_counter <= num_thresholds:

    thr = thresholds[threshold_counter]

    # transition from a lower pixel value to a higher pixel value
    if current_pixel_value <= thr and next_pixel_value >= thr:
        input neuron with the id 2*threshold_counter spikes

    # transition from a higher pixel value to a lower pixel value
    if current_pixel_value >= thr and next_pixel_value <= thr:
        input neuron with the id (2*threshold_counter + 1) spikes

    threshold_counter += 1
```

Listing S1: **Input encoding sMNIST.** Pseudo code for the input spike encoding of MNIST images used in the sMNIST task.

Task Name	Spiking RelNet	Non-spiking RelNet	Spiking RelNet incr. compute time	
			$T_{sim} = 45$ steps	All times tripled
Task 1: Single Supporting Fact	1.0	0.4	0.6	1.0
Task 2: Two Supporting Facts	–	20.8	–	–
Task 3: Three Supporting Facts	–	25.0	–	–
Task 4: Two Argument Relations	0.1	0.0	0.0	0.1
Task 5: Three Argument Relations	2.3	0.6	1.2	2.3
Task 6: Yes/No Questions	0.2	0.0	0.3	0.4
Task 7: Counting	0.7	0.6	0.6	1.4
Task 8: Lists/Sets	0.5	0.1	0.3	0.9
Task 9: Simple Negation	0.1	0.1	0.1	0.7
Task 10: Indefinite Knowledge	2.3	1.8	1.5	1.7
Task 11: Basic Coreference	0.9	1.5	1.6	2.1
Task 12: Conjunction	4.8	3.6	4.3	4.2
Task 13: Compound Coreference	3.9	2.5	3.6	3.6
Task 14: Time Reasoning	0.9	0.7	0.5	0.0
Task 15: Basic Deduction	0.1	0.0	0.0	0.0
Task 16: Basic Induction	–	52.6	–	–
Task 17: Positional Reasoning	18.4	4.6	6.5	2.3
Task 18: Size Reasoning	1.5	0.6	0.8	0.2
Task 19: Path Finding	0.8	7.9	1.5	3.7
Task 20: Agent Motivations	0.0	0.0	0.0	0.4

Table S3: The above table lists the percentage error of different architectures on the different tasks of the bAbI dataset. According to [Weston et al., 2015], tasks with errors under 5% are considered solved. The first two columns compare the performance of the Spiking RelNet to the non-spiking RelNet. The next two list the performances for Spiking RelNets for which the compute steps are increased. The column on the left corresponds to the case where the compute steps of the feed-forward networks $T_{sim} = 45$ time steps (compared to 37 time steps in the original network). For the column on the right, the network has all time lengths and time constants tripled compared to the original network. When allowing longer compute time, the Spiking RelNet is able to successfully solve all 17 tasks to within a few percent error of the non-spiking RelNet.

Task 16: Basic Deduction

```
1   Cats are afraid of wolves.
2   Mice are afraid of cats.
3   Sheep are afraid of mice.
4   Gertrude is a cat.
5   Wolves are afraid of sheep.
6   Jessica is a mouse.
7   Emily is a wolf.
8   Winona is a cat.

Q   What is jessica afraid of?
A   Cat
```

Task 18: Size Reasoning

```
1    The container fits inside the suitcase.
2    The chocolate fits inside the chest.
3    The box of chocolates fits inside the suitcase.
4    The chocolate fits inside the box.
5    The chocolate fits inside the container.
6    The container fits inside the suitcase.
7    The chocolate fits inside the box.
8    The suitcase is bigger than the chocolate.
9    The chocolate fits inside the chest.
10   The container is bigger than the box.
11   The box of chocolates fits inside the chest.
12   The chest fits inside the container.
13   The box fits inside the container.

Q    Is the chest bigger than the suitcase?
A    no
```

Task 19: Path Finding

```
1   The garden is west of the bathroom.
2   The bathroom is north of the bedroom.
3   The bathroom is south of the office.
4   The hallway is south of the bedroom.
5   The kitchen is east of the hallway.

Q   How do you go from the hallway to the
    bathroom?
A   n,n
```

Task 20: Agents Motivation

```
1   Antoine is tired.
2   Sumit is thirsty.
3   Jason is thirsty.
4   Jason moved to the kitchen.
5   Yann is hungry.

Q   Where will yann go?
A   Kitchen
```

Figure S1: **Examples for four types of tasks in the bAbI dataset.** Each example consists of a story and a question which allows a single word answer (except for Task 19, where two directions can be given as answer). The dataset consists of 20 types of tasks, each targeting a different aspect of reasoning. In this figure, examples from four tasks are shown. Each one demonstrates the requirement for relational reasoning in order to successfully answer the question. "Basic Deduction", "Path Finding", and "Size Reasoning" require only reasoning across multiple sentences. Task 20 "Agents Motivation" requires in addition association of concepts with information that is external to the current story.

S2.2 Parameters

The parameters pertaining to the LIF Neuron with AHP currents are listed in Table S4

S2.3 Details of the Spiking RelNet architecture

The embedding networks

In order to embed sentences and questions into spike sequence, we use LSNN Networks, which are networks of recurrently connected LIF neurons both with and without AHP currents. The LSNN network for sentences uses a different set of weights than those used by the LSNN network for questions, however all other parameters are identical and are described below.

The LSNNs each contain 200 LIF neurons of which a random subset of 100 neurons are LIF neurons with AHP currents. The synaptic delays take an integer value uniformly from 1 to 3 ms. Note here that all time lengths and time constants are specified in terms of computation steps on Loihi, with 1 computation step corresponding to 1 ms.

The input to the LSNN is described here. We assign each distinct word from the bAbI

| | | Spiking RelNet | |
Parameter	sMNIST	Original	triple length
Neuron Parameters:			
PSC decay τ_I (steps)	0.0	7.0	20.0
Voltage decay τ_V (steps)	20.0	7.0	20.0
AHP current decay τ_{AHP} (steps)	700.0	40.3	120.0
AHP current decrement β / V_{thr}	0.756	0.176	0.062
Surrogate gradient Parameters:			
Dampening factor γ	0.3	0.0	0.5
Scaled voltage support $[v_-, v_+]$	$[-1.0, 1.0]$	$[-1.0, 0.5]$	$[-1.0, 0.5]$

Table S4: **LIF Neuron Parameters:** Here we detail the parameters for the LIF neurons used in the different examples.

dataset an index, and associate an input neuron. When a word is provided as input to the LSNN, the corresponding input neuron, and only that neuron, emits spikes continuously for $T_{\mathrm{word}} = 10$ ms. To present a sentence or question to the LSNN, each word in it is encoded as above, and input to the LSNN in sequence so that the first word is aligned to the final time step. Thus the input to the LSNN takes at most 10 ms $* N_{\mathrm{words}} = 110$ ms where $N_{\mathrm{words}} = 11$ is the maximum number of words in a sentence or question of the bAbI task. These input neurons are connected to the LSNN in an all-to-all manner.

The final embedded spike sequences $o_i(t)/q(t)$ are formed by taking The spike activity of the LSNN over the last $T_{inp} = 14$ ms, and padding them with zeros up to a total compute time of $T_{sim} = 37$ ms. These embeddings are the input to the instances of the relational function $g_\theta(o_i(t), o_j(t), q(t))$, and consequently all the feed-forward LIF-networks from this point on are run for a length of T_{sim}. The zero padding and the longer compute time $T_{sim} > T_{inp}$ are necessary so that the spikes from the embedding have enough time steps to propagate through the several layers of the feed-forward LIF networks.

The relational function

Currently, the relational function g_θ is implemented as a 4-layer feed-forward network of LIF neurons without AHP currents, with 256 neurons in each layer. Each layer is connected all-to-all to the next. The synaptic delays take up values from 1 to 3 ms. The relational function takes as input a triplet containing a pair of embedded sentences and the embedded question $(o_i(t), o_j(t), q(t))$, and returns an output spike sequence. The spike sequences of the sentence pair and question are weighted and fed via all-to-all connections to the first layer.

The above relational function is applied once for each pair of sentences i, j in the story where $i \leq j$. Note that there are instances of g_θ that receive the same sentence twice i.e. $i = j$ The output of each instance of the relational function is a spike matrix,

The aggregation layer

The aggregation layer is an addition to the original architecture proposed in [Santoro et al., 2017], which is necessary due to constraints of neuromorphic hardware. These constraints are discussed in section S2.5. The aggregation is a single layer of LIF neurons without AHP currents, such that the output of each instance of g_θ is connected one-to-one to this layer. This implements an element-wise function f_{agg} on the outputs of g_θ instances summed. Each neuron of the aggregation layer receives as input the sum of the output spikes across all g_θ instances, and outputs a spike sequence.

The readout function

The readout function $f_\phi(\cdot)$ is implemented using a 3 layer feed-forward LIF network without AHP currents, with layer sizes 256, 512, 160, followed by a specially designed linear readout that is described in the methods section of the main text. Each layer, including the readout, is connected all-to-all to the next. The synaptic delays take up values from 1 to 3 ms.

S2.4   Details of Spiking RelNet training

Pretraining the LSNN networks

In order to reduce the number of epochs required to train the Spiking RelNet, we choose to pre-train the LSNN's that embed the sentences and questions as below. We first train a non-spiking implementation of the RelNet end-to-end until we reach optimal performance. The non-spiking LSNN uses LSTM units to embed the sentences and questions into representations. We then train the LSNN to reproduce the output of these pre-trained LSTM networks for all the sentences used in the database.

In order to readout a value from the LSNN, which can be compared to the LSTM output, we use a linear readout similar to the one used to train the entire RelNet (described in the main Methods). The number of readout units matches the dimension of the LSTM that we wish to approximate i.e. 32. We then compare the value of this readout to the output of the LSTM, using mean squared error. This mean squared error is used as the loss function to train the LSNN weights using back-propagation in time (BPTT).

The weights of the LSNN are then frozen, and the resultant spike embeddings learnt here are used as the input when training the weights of the feed-forward part of the Spiking RelNet. Note that when we use the spikes of these pre-trained LSNN's as input to the g_θ instances, the readout weights that were used to pre-train the LSNN are not used in any way. We directly connect the LSNN spikes to g_θ using randomly initialized weights and train these weights. We find that, one pre-trained, all the information pertaining to the sentence is encoded within the spikes of the LSNN, which can then be processed by the feed-forward part.

Rate regularization for g_θ instances

For all other layers, the rate regularization pushes the mean rate of each neuron across the batch toward a specific target rate. However we use a more aggressive regularization for the spike rates of the instances of g_θ.

Corresponding to each layer of g_θ, we calculate the following loss. Consider the b'th story in the batch. For this story, we denote spike rate of the neuron k, in the (i,j)'th instance of g_θ as $\rho_{k,ij}^b$ Now we define $R_k^b = \sum_{1 \leq i \leq j \leq M} \rho_{k,ij}^b$ as the total spike rate of neuron k across all g_θ instances for the b'th story. We then define the rate regularization loss as below:

$$L_R = \lambda_R(\text{mean}_k \ (\text{mean}_b(R_k^b) - R_{\text{target}})^2)^2$$

We calculate a similar loss for each layer of g_θ and sum them up as a part of the final loss. For each neuron, this regularization loss pushes the sum of its spike rate across the g_θ instances to a target value of R_{target}. For our networks, $R_{\text{target}} = 300\,\text{Hz}$. For a task with two sentences, this translates to a target of 3.7 spikes per neuron, and for a task with 20 sentences, this corresponds to 0.05 spikes per neuron. The corresponding spike rates achieved when trained can be seen in the main text Fig. 3B.

S2.5   Constraints on connectivity on Loihi

In this section we discuss the restrictions on network connectivity and the strategies used to place a Spiking RelNet that processes stories up to $M = 20$ sentences in length. In order to understand the following section, it is useful to define some of the relevant terminology:

Neuro-Core A neuro-core (short for neuromorphic core) is a fundamental computational element in the Loihi chip. Each neuro-core can compute the dynamics of up to 1024 Neurons, and contains a shared SRAM which contains data pertaining to the weights of the incoming synapses as well as shared configuration and state variables.

Chip A Loihi chip is a block of 128 interconnected neuro-cores integrated within the same silicon substrate. It is possible for multiple chips to be connected together to allow for a larger number of interconnected neuro-cores and thereby larger networks.

Axon The axon is a structure that is part of the infrastructure that implements connectivity and spike transport in Loihi. Loihi implements the connectivity between different neuro-cores via inter-core connections called axons. Each axon indexed by (i, C) is a connection between a presynaptic neuron i and a postsynaptic neuro-core C. If an axon (i, C) is connected, all spikes generated by the presynaptic neuron i, are routed through this axon to neuro-core C, where it is weighted by the relevant weights and delivered to the postsynaptic neurons. Each axon (i, C) is considered to be an *input axon* of neuro-core C and an *output axon* of the neuro-core that contains neuron i

We now discuss here the various constraints and their impact on the network architecture and placement.

Synaptic memory limit

As mentioned earlier, the SRAM in a neuro-core is used to store the parameters for any incoming connection to that neuro-core. The limited per-neuro-core SRAM memory for synaptic parameters puts a limit on the number of incoming synapses to a particular neuro-core. Except the aggregation layer, all layers in the network have densely connected input synapses which translates to a larger number of incoming connections than can fit into a single neuro-core. Thus, need to appropriately split the postsynaptic neurons across several

neuro-cores so that the total number of synapses coming into each neuro-core can fit into the memory.

Limits pertaining to fanouts – LSNN relay layer

Loihi has two limits that pertain to the fanout of a layer.

- *output axon limit* – The number of outgoing axons from a neuro-core is limited in general to 2048, and 4096 if all the outgoing axons are connected to neuro-cores within the same chip.

- *neuro-core fanout limit* – The number of different neuro-cores to which a single neuron can be connected to is at-most 512.

This constraint plays a role only in the case of the connections from the LSNN's to the first layer of the instances of g_θ. To see this, we have $\frac{M(M+1)}{2} = 210$ instances of g_θ. For a particular sentence k, there are exactly M pairs $(i,j) :: 1 \le i \le j \le M$ that contain k. There is also the pair (k,k) which contains k twice. This means that the output of the corresponding LSNN-k ($o_k(t)$) is connected to a g_θ instance $M + 1$ times in an all-to-all manner. It takes 4 neuro-cores (Table S5) to place the first layer of an instance of g_θ. This implies that each neuron of LSNN-k is connected to $4 \times (M+1) = 84$ neuro-cores, implying 84 output axons per LSNN neuron. Similarly for the question-LSNN, since it forms an input to all the instances of g_θ, we get $4 \times \frac{M(M+1)}{2} = 840$ output axons per LSNN neuron.

Given the above number of output axons per neuron, the output axon limit puts a stronger limit on the number of neurons per neuro-core than the synaptic memory limit. However, the LSNN is a recurrent network and inter-core communication will lead to significant latency if the neurons of the LSNN are spread across too many neuro-cores. Moreover, for the question-LSNN, that fact that each neuron fans out to 840 neuro-cores means that we violate the neuro-core fanout limit.

This motivates the use of relay networks. A *relay layer*, as the name suggests relays spikes from the layer that forms the input. It has the same number of neurons as the input layer and the input layer is connected in a one-to-one manner to it. Each time a neuron of the relay layer receives a spike, it generates a spike, thereby reproducing the input spike train as the output spike train.

Each LSNN is connected to multiple relay layers, each of which fans-out to a subset of the g_θ instances that take the corresponding sentence/question as input. Since the connection from the LSNN to the relay neurons is one-to-one, and we fanout to fewer relay networks than the original fanout to the g_θ instances, we have a much reduced fanout from the LSNN. So much so that this is no longer the a constraint in placing the LSNN. Also each relay neuron fans out to fewer neuro-cores than the original LSNN, and can be split across neuro-cores without additional latency cost, thus satisfying all fanout constraints.

The actual choice of relays and the fanouts from the relays is determined in a manner that minimize congestion in spike transport (described below in section S2.6).

Each sentence-LSNN fans out to 4 relay networks and the question-LSNN fans out to 10. Each LSNN can be placed on 2 neuro-cores (see Table S5), meaning each neuro-core has 100 neurons. The number of output axons per neuro-core for the LSNN's are then $4 * 100 = 400$ and $10 * 100 = 1000$ for the sentence-LSNN and question-LSNN respectively, both well within the output axon limits.

Input axon limit – the aggregation layer

Loihi limits each neuro-core to have a maximum of 4096 input axons. Note here that if a neuron i is connected to even a single neuron in neuro-core C, the axon (i, C) is an input axon of neuro-core C. In this case, we denote neuron i as being presynaptic to the neuro-core C. The input axon limit means that for any neuro-core, a maximum of 4096 neurons can be presynaptic to that neuro-core. Unlike the above two constraints which can be worked around, the input axon limit introduces a fundamental restriction to the network architectures that can be implemented on Loihi.

In the original formulation of relational networks in [Santoro et al., 2017], the outputs from all the relational function instances are summed together to form the input to the final readout function f_ϕ. In an ideal scenario, it is possible to implement this using spiking networks using the fact that incoming spike inputs can be summed into the PSC's of the postsynaptic neurons. However, when the summed spike train forms an input that is connected all-to-all into the f_ϕ network, it becomes a problem. To implement this, we would need to connect each relational function instance in an all-to-all manner to the final readout function, with shared weights across the different instances. However, consider that the output dimension of the relational function is 256. This connectivity would mean that if we consider a single neuron of the first layer of f_ϕ, this neuron would receive input from all $256 \cdot \frac{M(M+1)}{2} = 53760$ output neurons across the instances of g_θ. This means that a neuro-core with even a single such neuron would have 53760 presynaptic neurons, and thus input axons, which is completely beyond the input axon limit.

In order to deal with this constraint, we have modified the original architecture of the Relational Network by adding an *aggregation layer*. The aggregation layer is a layer of spiking LIF neurons without AHP currents, that has the same number of neurons as the output layer of g_θ, and to which each instance of g_θ is connected one-to-one with weights shared across the g_θ instances. This leads to each neuron from the aggregation layer having 210 presynaptic neurons. The input axon limit thus allows up to 19 neurons per neuro-core, and the 256 neurons of the aggregation layer can be placed onto 14 neuro-cores.

In Table S5, we give the number of neuro-cores required to place an instance of a layer is given for the different layers of the Spiking RelNet. Also mentioned are the constraints that lead to the layers being split across as many neuro-cores.

S2.6 Optimized network placement to minimize spike congestion

In the methods section of the main text, we discuss how it is necessary to place the relay layers and the the several instances of the relational function g_θ in a careful manner so that the amount of cross-chip communication is minimized.

In order to perform this optimization, it is quite difficult to actually optimize the cross-chip communication as that is not easy to calculate directly as a function of the network placement. Instead, we first notice that a majority of the spike congestion occurs when transferring the spikes from the LSNN networks, via the relay layers to the input of the various instances of g_θ. Thus it makes sense to optimize the placement of the relay layers and first layer of the instances g_θ in such a manner that the number of connections across chips is minimized. This objective is translated into the following constraints on the network placement.

Relational Network Layer	Layer Size	Number of cores per instance	Connectivity limit	Number of Instances	Total Cores
LSNN (Sentences)	200	2	Synaptic Memory	$M = 20$	40
LSNN (Question)	200	2	Synaptic Memory	1	2
Relay networks (sentences)	200	1-2*	Output Axon	$4M = 80$*	100
Relay networks (questions)	200	3-5*	Output Axon	10*	42
g_θ Layer 1	256	4	Synaptic Memory	$\frac{M(M+1)}{2} = 210$	840
g_θ Layer 2	256	2	Synaptic Memory	210	420
g_θ Layer 3	256	2	Synaptic Memory	210	420
g_θ Layer 4	256	2	Synaptic Memory	210	420
Aggregation Layer	256	14	Input Axon	1	14
f_ϕ Layer 1	256	2	Synaptic Memory	1	2
f_ϕ Layer 2	512	4	Synaptic Memory	1	4
f_ϕ Layer 3	160	3	Synaptic Memory	1	3
Linear Readout Layer	180	1	Synaptic Memory	1	1
				Total Cores:	2308

* Determined by placement scheme that minimizes spike congestion.

Table S5: **Parameters for the placement of the Spiking RelNet onto Loihi:** This table gives, for each layer of the relational network, the number of cores required to place all the neurons in a single instance of that layer. Among the three limits on connectivity, i.e. limits on synaptic memory, input axons, and output axons, the limit that ultimately results in the layer needing to be divided among multiple cores is specified as the corresponding connectivity limit. Also mentioned are the number of instances of each network corresponding to a maximum story size of $M = 20$ sentences, and the total number of cores needed

- Firstly, a relay layer that forms an input to an instance of g_θ is located on the same chip as the first layer of that instance. That is, we place initial layer of the g_θ instances on the same chip as the relay layers that give them their corresponding input sentence and question embeddings.

- With this constraint, the connections from the LSNN's to the relay layers become the cross-chip communication that needs to be optimized, meaning that we need to minimize the number of relay layers used. Each chip has a limit of 128 neuro-cores and thus a limit on the number of g_θ instances whose initial layers can be placed on a single chip. Thus, when choosing the g_θ instances to be placed on the same chip, we select them such that the number of distinct sentences needed as input for this set of instances is minimized, thus minimizing the number of relay layers.

The resultant placement of the relay networks and the initial layer of the instances of g_θ is detailed in Fig. S2.

For the subsequent layers of the g_θ instances, we simply place each instance in sequence with only the constraint that all the remaining 3 Layers should lie on the same chip. We find that the cross-chip communication from the first layer to the subsequent layers is very minimal and does not cause a significant delay.

The advantage of such this layout is two-fold. Firstly, since a majority of the connections are within the chip, this significantly reduces the congestion that happens when

transporting the spikes across chips. Secondly, since each relay network fans out only to neuro-cores that are placed on the same chip, the output axon limit is now 4096 rather than 2048 thus requiring fewer relay networks.
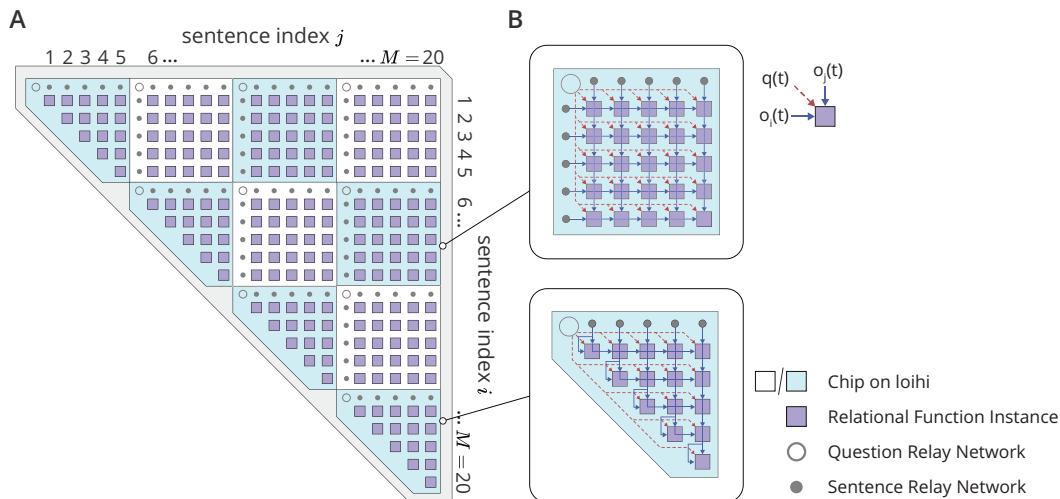

Figure S2: **Placement of relay networks and the initial layer of g_θ instances that minimizes spike congestion** − **A)** Here we show how the first layer of the several instances of the relational function g_θ are placed across the several chips. The instances of g_θ are arranged according to the indices of the input sentences $o_i(t), o_j(t) :: i \leq j$. Each blue/cyan block corresponds to a single Loihi chip, within which we show the g_θ instances whose first layer is placed on that chip, along with the relay layers that give them the input. Each cell also has a relay layers for the question embedding $q(t)$. The instances are grouped into squares because this minimizes the number of distinct inputs (i.e. relay layers) needed. **B)** A blowup of a chip showing the connections between the relay layers and the contained instances of g_θ.

References

[Davies et al., 2018] Davies, M., Srinivasa, N., Lin, T., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., Liao, Y., Lin, C., Lines, A., Liu, R., Mathaikutty, D., McCoy, S., Paul, A., Tse, J., Venkataramanan, G., Weng, Y., Wild, A., Yang, Y., and Wang, H. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *micro*, 38(1):82–99.

[Santoro et al., 2017] Santoro, A., Raposo, D., Barrett, D. G., Malinowski, M., Pascanu, R., Battaglia, P., and Lillicrap, T. (2017). A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976.

[Weston et al., 2015] Weston, J., Bordes, A., Chopra, S., Rush, A. M., van Merriënboer, B., Joulin, A., and Mikolov, T. (2015). Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.