

What fundamental issue in the neuromorphic research field does the NeuroBench framework aim to address?	The lack of standardized benchmarks, making it difficult to measure advancements and compare performance.
NeuroBench is a benchmark framework for neuromorphic algorithms and systems, collaboratively designed by researchers from ____ and ____.	Industry, academia
What are the two main tracks of the NeuroBench framework?	The algorithm track and the system track.
What is the primary purpose of the NeuroBench algorithm track?	To enable hardware-independent evaluation of neuromorphic algorithms using complexity metrics.
What is the primary purpose of the NeuroBench system track?	To measure the real-world speed and efficiency of fully deployed neuromorphic hardware solutions.
How does the interplay between the algorithm and system tracks in NeuroBench create a 'virtuous cycle'?	Algorithm innovations guide system design, while system-level insights accelerate further algorithmic progress.
Which NeuroBench track would be used to evaluate an SNN algorithm simulated on a conventional GPU?	The algorithm track, as it is designed for hardware-independent evaluation.
According to the paper, what are the three outstanding challenges that NeuroBench aims to solve for neuromorphic benchmarking?	Lack of a formal definition, implementation diversity, and rapid research evolution.
What is the term for the memory size, in bytes, required to represent a model's parameters and buffers in the NeuroBench algorithm track?	Footprint.
In the NeuroBench algorithm track, what metric is defined as the number of zero weights divided by the total number of weights in a model?	Connection Sparsity.
What NeuroBench algorithm metric measures the average sparsity of neuron outputs over all layers, timesteps, and samples during execution?	Activation Sparsity.
What are 'static metrics' in the NeuroBench algorithm track?	Metrics that can be determined from the model alone without execution, such as Footprint and Connection Sparsity.
What are 'workload metrics' in the NeuroBench algorithm track?	Metrics that depend on the execution or simulation of the model with benchmark data, such as correctness and Activation Sparsity.
What does the 'Synaptic Operations' metric in NeuroBench measure?	The average number of synaptic operations per model execution, based on neuron activations and their fanout synapses.
Synaptic operations with non-binary activations are considered ____, while those with binary activations are considered ____.	multiply-accumulates (MACs), accumulates (ACs)
What do 'Eff. MACs' and 'Eff. ACs' metrics represent in NeuroBench?	Effective synaptic operations that disregard zero activations and zero connections, reflecting costs on sparsity-aware hardware.
What is the 'Model Execution Rate' in the NeuroBench algorithm track?	The rate in Hz at which the model's forward inference pass needs to be executed to process real-time data.
What is the full name of the 'Keyword FSCIL' benchmark task in NeuroBench's algorithm track?	Keyword Few-Shot Class-Incremental Learning.
What key challenge of biological intelligence does the Keyword FSCIL benchmark evaluate?	Learning new tasks from a small number of samples while retaining knowledge of prior tasks.
What dataset is used for the Keyword FSCIL benchmark?	The Multilingual Spoken Word Corpus (MSWC).
What is the correctness metric for the Event Camera Object Detection benchmark in NeuroBench?	COCO mean average precision (mAP).
What is the task of the 'NHP Motor Prediction' benchmark?	To predict two-dimensional fingertip velocity from multi-channel recordings of the sensorimotor cortex.
What correctness metric is used for the NHP Motor Prediction benchmark?	The coefficient of determination ( $R^2$ ) score.
What is the task in the Chaotic Function Prediction benchmark?	To autoregressively forecast the future values of a one-dimensional Mackey-Glass time series.
What correctness metric is used for the Chaotic Function Prediction benchmark?	Symmetric mean absolute percentage error (sMAPE).
What is the purpose of the NeuroBench algorithm track 'benchmark harness'?	To standardize benchmark interfaces for loading data, processing, and calculating metrics to ensure fair comparisons.
The NeuroBench harness standardizes runtime execution using components like ____, ____, and ____.	Dataloaders, Processors, Accumulators
In the Keyword FSCIL baseline results, what incremental learning approach was used that involved feature-based clustering?	Prototypical networks.
In the Event Camera Object Detection task, the RED ANN baseline uses blocks of feed-forward convolutional layers and recurrent ____ layers.	convolution-LSTM (ConvLSTM)
The Hybrid ANN-SNN model for event camera object detection replaces the ConvLSTM layers of the RED ANN with what type of spiking neuron layers?	Leaky integrate-and-fire (LIF) spiking neural layers.
In the NHP motor prediction task, the SNN baseline achieved similar correctness to the ANN but with significantly higher ____, leading to fewer effective operations.	activation sparsity
What type of recurrent ANN, belonging to the class of reservoir computing, was used as a baseline for the Chaotic Function Prediction task?	An Echo State Network (ESN).
How did the high connection sparsity of the ESN model impact its performance on the chaotic prediction task compared to the LSTM?	It led to a significant reduction in synaptic operations, using an order of magnitude fewer effective operations for comparable accuracy.
In the Event Camera Object Detection task, the RED ANN baseline uses blocks of feed-forward convolutional layers and recurrent ____ layers.	NeuroBench system track measurements must include the cost of data pre- and post-processing.
Why must correctness be measured in the NeuroBench system track, unlike in MLPerf's closed category where it's a threshold?	Because the model used to solve a task is unconstrained due to the tight coupling between algorithms and specific neuromorphic hardware.
What is a key hallmark of biological intelligence that makes power and energy consumption first-order metrics in the NeuroBench system track?	Energy efficiency.
What v1.0 system track benchmark challenges systems to classify audio into categories like 'airport' or 'bus' for low-power hearable devices?	Acoustic Scene Classification.
What dataset, from a well-known challenge, is used for the Acoustic Scene Classification system benchmark?	The DCASE challenge dataset (TAU Urban Acoustic Scenes).
What is a key difference in measurement between the NeuroBench system track and conventional system benchmarks like MLPerf?	Quadratic Unconstrained Binary Optimization (QUBO).
In the QUBO system benchmark, what metric measures the relative gap between the system's solution and the best-known solution (BKS)?	BKS-Gap.
For the Acoustic Scene Classification benchmark, the neuromorphic Xylo system showed a dynamic inference energy consumption ____ times less than the Arduino CPU baseline.	33.4
For the QUBO benchmark, the Loihi 2 neuromorphic system solver required ____ times less power compared to the best CPU solver.	37.24
What future direction for NeuroBench involves tasks where a system must sense and act within an environment, demonstrating embodied intelligence?	Closed-loop benchmarks.
In the context of neuromorphic computing, what does SNN stand for?	Spiking Neural Network.
In the context of machine learning models, what does ANN stand for?	Artificial Neural Network.
Which baseline model for the Keyword FSCIL task used Mel-frequency cepstral coefficients (MFCC) for pre-processing?	The M5 ANN baseline.
Which pre-processing algorithm was used to convert audio samples to spike trains for the SNN baseline in the Keyword FSCIL task?	Speech2Spikes (S2S).
The Mackey-Glass differential equation, used in the chaotic prediction benchmark, includes a parameter $\tau$ that represents the system's ____.	delay constant
The expected predictability timescale for a chaotic system, used as a time unit in the Mackey-Glass benchmark, is known as the ____.	Lyapunov time
The LSTM baseline for chaotic prediction utilized an explicit memory buffer, while the ESN model's memory was only implicitly retained within its ____.	hidden layer (reservoir)
In the NHP motor prediction task, what was the effect of flattening input data for the ANN, Flat and SNN, Flat models?	It significantly increased prediction quality ( $R^2$ ) but also came with a significant complexity trade-off (larger footprint or more operations).
Why might a high activation sparsity in an ANN not translate to a proportional reduction in effective synaptic operations?	Processing steps like batch normalization applied before a connection layer can transform sparse activations into dense ones.
The QUBO benchmark workloads search for the maximum independent set of nodes in a graph, which is an ____-hard problem.	NP
The Loihi 2 neuromorphic solver for the QUBO task uses an SNN formulation of which optimization algorithm?	Simulated annealing.
In the Acoustic Scene Classification system benchmark, what is the difference between 'idle power' and 'active power' used to calculate dynamic power?	Idle power is when the system is ready for inference, while active power is measured during pre-processing or inference execution.
For the system track, NeuroBench encourages transparent documentation of methodology. For what reason is this particularly important?	The wide diversity in neuromorphic system implementations, board-level integration, and development maturity makes completely consistent methodology challenging.
What is the proposed model description framework that could be used in future NeuroBench system tracks for a closed-algorithm category?	The Neuromorphic Intermediate Representation (NIR) model description framework.
In the NHP motor prediction SNN baseline, the output neuron's ____ was directly read to produce velocity predictions, functioning as a leaky accumulator.	membrane potential
What is the primary goal of neuromorphic computing?	To unlock hallmarks of biological intelligence by porting brain-inspired principles into engineered computing devices and algorithms.
In the chaotic prediction task, why is the ESN model architecturally unique compared to other ANNs in the study?	It has a randomly generated and fixed reservoir layer, resulting in very high connection sparsity.
What are the three object classes used for evaluation in the Event Camera Object Detection benchmark?	Pedestrian, two-wheeler, and car.
What does the term 'few-shot' signify in the Keyword FSCIL benchmark?	That the model must learn new classes with only a handful of training samples available per class.
What does the term 'class-incremental' signify in the Keyword FSCIL benchmark?	That the model successively incorporates new classes over multiple sessions without forgetting old ones.
The NeuroBench framework is described as iterative and community-driven, analogous to which well-established benchmark framework for machine learning?	MLPerf.
The dataset for the NHP Motor Prediction task consists of recordings from the sensorimotor cortex of what type of animal?	Non-human primates (NHP Indy and NHP Loco).
In the equation for the Mackey-Glass time series, $\frac{dx}{dt} = \beta x(t - \tau) / (1 + x(t - \tau)^\alpha) - \gamma x(t)$ , what does the term $x(t - \tau)$ represent?	The state of the system at a past time, introducing a time delay.
The NeuroBench framework is designed to be ____, encouraging participation of both neuromorphic and non-neuromorphic approaches.	Inclusive
In the Keyword FSCIL benchmark, after each incremental session, the model is tested on which classes?	On all previously learned classes, including the base classes and all incrementally-learned classes.
For the QUBO benchmark, the complexity of workloads is defined by three parameters: number of nodes, density of connections, and a ____.	random seed
What is the key advantage of the SNN's binary activation in the NHP motor prediction task?	It translates high activation sparsity directly into a very low number of effective accumulate (AC) operations.
The RED ANN architecture for object detection uses a ____ head to predict the location and class of bounding boxes.	single-shot detection (SSD)
Why are power and energy considered first-order metrics in the NeuroBench system track but often secondary in conventional benchmarks?	Because energy efficiency is a key hallmark of biological systems and thus a primary focus of neuromorphic research.
In the system track, the QUBO benchmark measures solution optimality after a ____, pre-set runtime.	fixed
What type of search solver was used as a CPU baseline for the QUBO benchmark, which maintains a list of prohibited actions to avoid local minima?	Tabu Search (TABU).