SymSeqBench: a unified framework for the generation and analysis of rule-based symbolic sequences and datasets

B. Zajzon(1), Y. Bouhadjar(2), M. Fabre(2, 3), F. Schmidt (2, 4), N. Ostendorf (1, 4) E. Neftci (2, 6), A. Morrison (1, 5), and R. Duarte (7, 8)

(1)Institute for Advanced Simulation (IAS-6), Jülich Research Centre, Germany, (2)Peter-Grunberg Institute, Neuromorphic Software Ecosystems (PGI-15), Jülich Research Centre, Germany, (3)Groningen Cognitive Systems and Materials Center (CogniGron), University of Groningen, (4)RWTH Aachen University, Aachen, Germany, (5)Department of Computer Science 3 - Software Engineering, RWTH Aachen University, Aachen, Germany, (6)Department of Electrical Engineering and Information Technology, RWTH Aachen University, Aachen, Germany, (7)Center for Neuroscience and Cell Biology (CNC-UC), University of Coimbra, Portugal, (8)Center for Innovative Biomedicine and Biotechnology (CIBB-UC), University of Coimbra, Portugal

Sequential structure is a key feature of multiple domains of natural cognition and behavior, as well as an integral property of tasks to which we would like to apply artificial intelligence. Despite this centrality, computational models of sequence learning are typically evaluated on a small number of ad hoc, manually curated sequences, with little regard for sequence complexity, higher-order statistical dependencies, or stimulus variability. This fragmented evaluation landscape makes it difficult to systematically compare models, identify their computational limits, or relate their capabilities to established cognitive tasks and benchmarks.

To address this issue, we introduce two complementary software tools: SymSeq, designed to rigorously generate and analyze structured symbolic sequences, and SeqBench, a comprehensive benchmark suite of rule-based sequence processing tasks to evaluate the performance of artificial learning systems in cognitively relevant domains. In combination, SymSeqBench offers versatility in investigating sequential structure across diverse knowledge domains, including experimental psycholinguistics, cognitive psychology, behavioral analysis, neuromorphic computing and artificial intelligence. Due to its basis in Formal Language Theory (FLT), SymSeqBench provides researchers in multiple domains with a convenient and practical way to apply the concepts of FLT to conceptualize and standardize their experiments, thus advancing our understanding of cognition and behavior through shared computational frameworks and formalisms.

For the neuromorphic and computational neuroscience community, SymSeqBench offers three key capabilities. First, it provides standardized tasks specifically designed to probe sequence-processing architectures on higher-order, non-Markovian dependencies with controllable complexity, conditions under which prior evaluation has been especially sparse. Second, it includes model-to-data comparison tools that permit direct benchmarking of artificial systems against natural behavioral sequences and a wide range of datasets, supporting translation between computational models and empirical neuroscience. Third, by connecting tasks to a principled formal framework, it offers a shared vocabulary for characterizing the computational capacities and limitations of a given architecture, linking questions about biological plausibility to questions about expressive power.

Proof-of-concept results span three domains. In computational neuroscience, we benchmark performance scalability in biologically plausible spiking network models to reveal scaling with sequence complexity and memory horizons achievable with local plasticity rules. In neuromorphic computing, spiking and conventional architectures are systematically compared on spoken digit and speech command datasets structured as sequential classification problems, demonstrating the advantages of adaptive neural dynamics and highlighting performance gaps at higher complexity levels. Finally, multi-scale analyses of empirical animal behavioral sequences provide quantitative evidence for supra-regular generative structure in certain species, connecting neuroethological observations to formal notions of sequential complexity.

SymSeqBench is modular, reproducible, and openly available. It is intended as a community resource for researchers working at the intersection of neural computation, neuromorphic engineering, and the formal theory of learning, providing a shared framework with which biological, computational, and hardware models of sequence processing can be compared on equal footing.