

Neuromorphic Heart Rate Monitors: Neural State Machines for Monotonic Change Detection

Original

Neuromorphic Heart Rate Monitors: Neural State Machines for Monotonic Change Detection / Carpegna, A., De Luca, C., Pozzi, F.E., Savino, A., Di Carlo, S., Indiveri, G., Donati, E.. - ELETTRONICO. - (2024), pp. 1-5. (2024 IEEE Biomedical Circuits and Systems Conference (BioCAS) Xi'an (CHN) 24-26 October 2024) [10.1109/biocas61083.2024.10798178].

Availability:

This version is available at: 11583/2996382 since: 2025-03-27T12:51:03Z

Publisher:

IEEE

Published

DOI:10.1109/biocas61083.2024.10798178

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Neuromorphic Heart Rate Monitors: Neural State Machines for Monotonic Change Detection

Alessio Carpegna*, Chiara De Luca^{†‡}, Federico Emanuele Pozzi[§], Alessandro Savino*, Stefano Di Carlo*, Giacomo Indiveri[†], Elisa Donati[†]

*Department of Control and Computer Engineering, Politecnico di Torino, Torino, Italy

[†]Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland

[‡]Digital Society Initiative, University of Zurich, Zurich, Switzerland

[§]Neurology Department, Fondazione IRCCS San Gerardo Dei Tintori, Monza, Italy

Abstract—Detecting monotonic changes in heart rate (HR) is crucial for early identification of cardiac conditions and health management. This is particularly important for dementia patients, where HR trends can signal stress or agitation. Developing wearable technologies that can perform always-on monitoring of HRs is essential to effectively detect slow changes over extended periods of time. However, designing compact electronic circuits that can monitor and process bio-signals continuously, and that can operate in a low-power regime to ensure long-lasting performance, is still an open challenge. Neuromorphic technology offers an energy-efficient solution for real-time health monitoring. We propose a neuromorphic implementation of a Neural State Machine (NSM) network to encode different health states and switch between them based on the input stimuli. Our focus is on detecting monotonic state switches in electrocardiogram data to identify progressive HR increases. This innovative approach promises significant advancements in continuous health monitoring and management.

I. INTRODUCTION

Detecting and quantifying Heart Rate (HR) has emerged as a critical tool for identifying potential pathologies and providing valuable insights into cardiovascular health [1]. Among variable behaviors, monotonic HR changes indicate unidirectional trends, either increasing or decreasing, in average HR over time. In non-clinical settings, such as general well-being and athletic training, tracking monotonic changes in HR is essential for evaluating physical fitness and recovery rates [2]. In clinical settings, monitoring monotonic changes in HR is crucial for medical diagnosis and patient monitoring. A consistent monotonic increase or decrease in heart rate can be an early indicator of cardiac conditions such as arrhythmia, bradycardia, or tachycardia. Early detection of these trends enables timely medical intervention, preventing more severe complications and improving patient outcomes [3].

Understanding monotonic increases in HR is also particularly important in patients with dementia, as it can help detect physiological stress or discomfort that could precede or accompany agitation states [4], [5]. By monitoring trends in dementia care and detecting early signs of agitation, healthcare providers

and caregivers can intervene more promptly and effectively. This may potentially reduce the severity and frequency of agitation episodes, thereby improving the quality of care and the quality of life for dementia patients and reducing the burden on caregivers [6].

Due to the importance and, at the same time, the difficulty of detecting these changes for human interpreters, an always-on system capable of continuously monitoring and detecting changes in average HRs can be extremely impactful. Always-on devices for health monitoring must meet stringent requirements, including low power consumption and real-time processing. For patients with dementia, these devices must operate for extended periods without frequent recharging, as subjects might not remember to charge the device regularly. Therefore, a low-power device that can be used as a “wear and forget” system is essential to ensure effective health management.

To this end, neuromorphic technology offers a compelling solution, providing energy-efficient and reliable processing capabilities [7]. Sensory-processing systems built using mixed-signal neuromorphic circuits are well-suited to the demands of continuous health monitoring [8]. Example solutions have already been successfully applied in a wide range of wearable applications, such as electrocardiogram (ECG) anomaly detection [9], [10], High Frequency Oscillation (HFO) detection [11], and Electromyography (EMG) decoding [12], [13].

In this paper, we present for the first time a neuromorphic implementation of an on-line signal processing system to specifically detect monotonic changes in average HRs over a long period of time. We deploy computational primitives of analog neuron circuits, such as recurrent neural network models of finite state machines (i.e., NSMs) [14], [15] that can switch between states, each encoding different average HR conditions, independent of the time elapsed between state changes. Here, we focus on a monotonic state switch, tested on ECG, to detect a progressive HR increase. We validate our model’s ability to track HR changes during activities (i.e., walking, cycling) by testing it on a dataset with varying patterns. We demonstrate how the model accurately follows monotonic changes (steady increase or decrease) while remaining inactive for non-monotonic signals. We further show how the robust computational properties of NSMs allow the mixed-signal neuromorphic processor to produce accurate and reliable results.

This paper has received funding from: The NEURO PULS project in the European Union’s Horizon Europe research and innovation programme under grant agreement No. 101070238; Bridge Fellowship funded by the Digital Society Initiative at University of Zurich (grant no.G-95017-01-12). We also thank the University of Zurich for supporting this project.

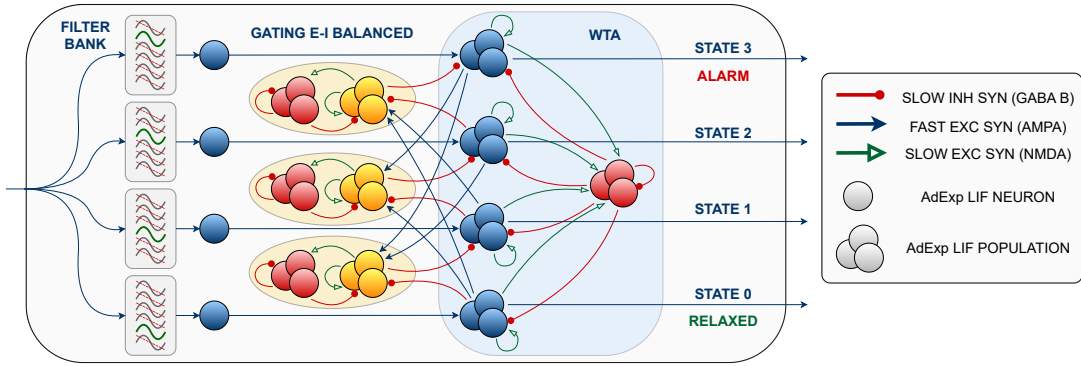


Fig. 1: Monotonic Neural State Machine (NSM). The input signal is filtered through four 4th order Butterworth bandpass filters. Each filtered component is converted into spikes through a adaptive exponential integrate-and-fire model (AdEx I&F) neuron (blue). The input spikes are fed into populations of AdEx I&F neurons (blue), encoding different states of the network, interconnected in a winner-take-all (WTA) configuration via a common inhibitory population (red). States are connected to gating populations (yellow) to implement the monotonic computation. These are organized in a Excitatory-Inhibitory (EI)-balanced configuration with an inhibitory population (red) limiting the overall activity.

The low power consumption, 90 μ W, and real-time processing features of these neuromorphic circuits make them ideal candidates for building continuous, long-term health monitoring devices in both clinical and non-clinical settings.

II. MATERIALS AND METHODS

A. Neuromorphic Hardware

The neuromorphic processor used in this study is the DYNAP-SE chip [16]. It is a custom-designed asynchronous mixed-signal processor that features analog spiking neurons and synapses that mimic the biophysical properties of their biological counterparts in real-time. The chip comprises four cores, each containing 256 AdEx I&F neurons. Each synapse can be configured as one of four types: slow/fast and inhibitory/excitatory. Each neuron includes a Content Addressable Memory (CAM) block with 64 addresses, representing the pre-synaptic neurons to which it is connected. Digital peripheral asynchronous input/output logic circuits receive and transmit spikes via the Address-Event Representation (AER) communication protocol [17]. In this system, each neuron is assigned a unique address encoded as a digital word, which is transmitted using asynchronous digital circuits as soon as an event is generated. The chip features a fully asynchronous inter-core and inter-chip routing architecture, allowing flexible connectivity with microsecond precision even under heavy system loads.

B. Dataset and signal processing

To test our model, we selected an available dataset that includes left wrist photoplethysmogramss (PPGs) and chest ECG recordings taken while participants used an indoor treadmill and exercise bike, accompanied by simultaneous motion data from accelerometers and gyroscopes [18]. Participants performed various exercises, such as walking, light jogging/running on a treadmill, and pedaling at low and high resistance, each for up to 10 minutes. The dataset includes records from 8 participants,

with most participants spending 4 to 6 minutes per activity. The signals were sampled at 256 Hz, and the ECG records were processed with a 50 Hz notch filter to remove mains interference.

An energy-based approach was devised for signal-to-spike conversion [19]. The proposed method comprises two stages: bandpass filtering and Leaky Integrate-and-Fire (LIF) neurons [20]–[22]. Each input channel was processed through a series of bandpass filters. We used four bands: 60–82 (#0), 82–105 (#1), 105–128 (#2), and 128–150 (#3) beats per minute (bpm), employing fourth-order Butterworth filters. This covers a reasonable range of HR variation, spanning from a relaxed state to a tachycardiac one. Afterward, the signal was full-wave rectified and injected as a time-varying current into a simple LIF neuron.

C. Network on chip

Figure 1 shows the designed NSM. The architecture is based on EI-balanced populations of AdEx I&F neurons. These populations can maintain sustained activity for extended periods, much larger than the time constants of the neurons themselves. This allows the network to implement a working memory capable of processing signals that change slowly compared to the time scales used within the chip.

The core of the network is a set of four EI-balanced populations, organized in a winner-take-all (WTA) architecture. The shared inhibition population (red) encodes the state of the network by sustaining the activity of a specific population while silencing all the others. This guarantees that when receiving the input signal, such as an ECG recording, only the population corresponding to the most active frequency range is activated.

A second set of EI-balanced populations, named Gating EI-balanced populations, is used to control the transition between different states. The goal is to target a progressive increase in HR, making the system robust to short fluctuations or temporary bpm decreases, following only the average trend




Fig. 2: Network behavior when stimulated with 50Hz Poissonian sequences testing all the possible transitions: (a) non-monotonic WTA dynamics; (b) monotonic NSM response

of the input. The mechanism of dis-inhibition [15] is exploited to ensure the network switches only to increasing states. The gating populations continuously inhibit the inactive elements of the WTA network, making them insensitive to any input stimulus. When a state is activated, it turns off (inhibits) the gating population of the subsequent state. This leads to the dis-inhibition of the next state, which makes it sensitive to the input. At the same time, it activates all the gating populations of previous states and the states following the next. This guarantees that the network follows only monotonic transitions between states and does not require precise tuning of the populations and connection synapses to make the system work. What matters is (i) that populations are sufficiently stable to maintain a sustained activity and (ii) that the inhibition between the gating populations and the WTA states is strong enough to silence any spiking activity completely. Finally, state 0 has no gating populations connected to it. This allows it to be used as a reset state: if the monotonic increase is only partial, and the HR returns to the relaxation range before reaching the alarming threshold, the network restarts, waiting for a new ramp-up in the input.

III. EXPERIMENTAL RESULTS

To obtain reliable computation in the NSM while minimizing the overall network size, we used 16 neurons per population, with a total requirement of 176 neurons. Table I shows the average connection probabilities for different populations. The resulting network is compact, fitting well within a single core of the target chip [16]. Each population exhibits an average firing rate of approximately 50Hz. The total power consumption can be estimated by integrating the various contributions required for spike generation and communication [23], [24]. The obtained average power consumption is around $90\mu W$, indicating that our network successfully balances stable, sustained activity with low power consumption.

The behavior of the hardware system was first tested using control stimuli produced as 50Hz Poisson input spike trains. Figure 3 shows the sustained activity of one EI-balanced

TABLE I: Connection types and average probabilities

	Connection	Synapse type	Probability
WTA	state -> inh	NMDA	60%
	inh -> state	GABA B	60%
	state -> state	NMDA	83%
	inh -> inh	GABA B	20%
GATING	gate -> inh	NMDA	30%
	inh -> gate	GABA B	30%
	gate -> gate	NMDA	50%
	inh -> inh	GABA B	50%
MONOTONIC	gate -> state	GABA B	100%
	state -> gate	GABA B	100%
	state -> gate	AMPA	100%
INPUT	lif -> state	AMPA	100%




Fig. 3: Stable sustained activity of one of the EI-balanced primitives included in the networks

population stimulated for one second. As shown, the network can keep a stable activity even when the input is removed.

A. WTA network: non-monotonic state transitions

As a second test, we evaluated the dynamics of the WTA network. Figure 2(a) shows the response of the network when stimulated with Poisson spike trains through a protocol covering all the possible transitions. In this case, each state is activated when receiving the corresponding input, regardless of the order of arrival. Once stimulated, the population becomes active, it

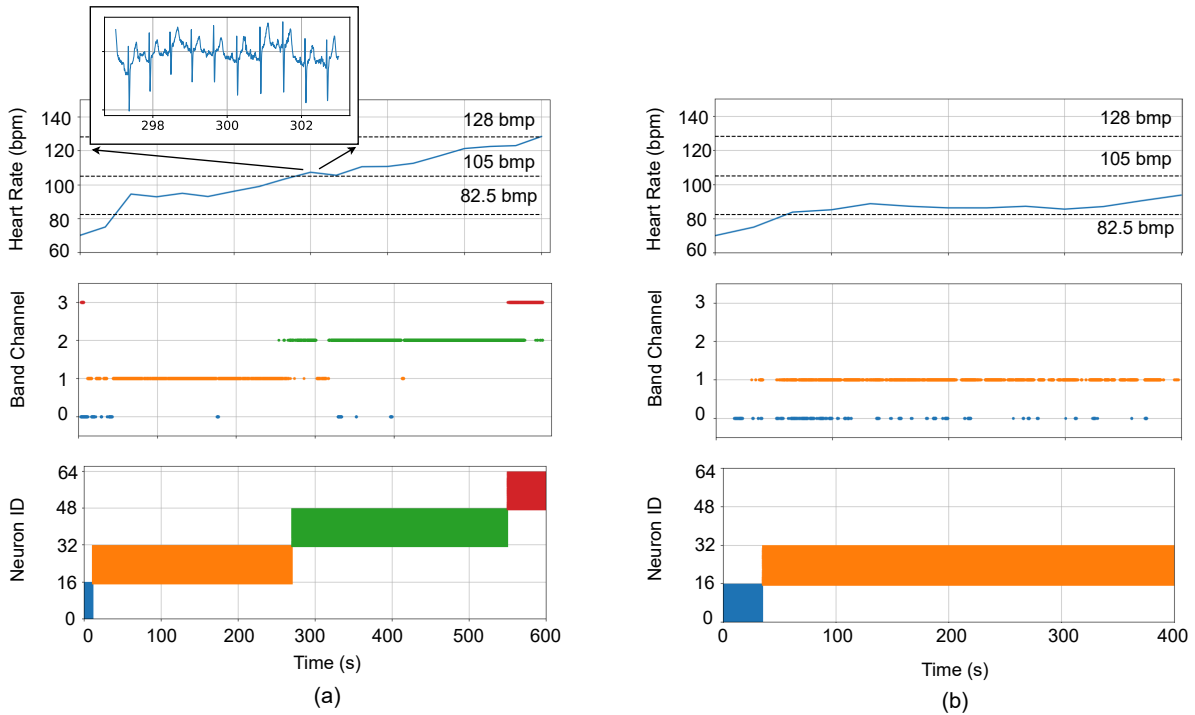


Fig. 4: Response of the network when stimulated with a real ECG signal: (a) intense 10 minutes bike session; (b) 6 minutes and 40 seconds of walk

suppresses the activity of all other populations, and stays active until a new input is provided to a different population.

B. NSM network: monotonic state transitions

When adding the gating connections, the response of the network becomes the one shown in Fig. 2(b). The stimulation protocol is the same as in Section III-A. In this case, however, the network switches state only (i) towards a higher state, following the expected monotonic behavior, or (ii) towards state 0, which behaves as a reset state for the network, as mentioned in Section II-C

C. Network dynamics on real ECG signal

Finally, we evaluated the network performance using real ECG signals. Here, we demonstrate its robust and coherent ability to monitor and detect monotonic increases in HRs. To achieve this, we fed the network pre-processed ECG recordings from the dataset described in Section II-B. Specifically, we tested the network under two different conditions: during an intense physical activity, namely cycling, lasting approximately 10 minutes (Fig. 4(a)), and during a walking session of around 6 minutes (Fig. 4(b)). The first scenario correctly detects a complete ramp-up of the HR, leading to the activation of the alarming state. These results show how the network is robust to spurious transitions: at time zero, the stimulation on input 3, caused by noise in the recording upper band, is ignored since state 3 is completely inhibited by its gating population. The same effect can be observed between 300 and 400 seconds with the network ignoring noisy stimuli from input 1. Note that

also the weak stimuli coming from input 0 are ignored both in state 1 and 2, despite the absence of a gating population in this case. This shows that the WTA dynamics by themselves are robust to spurious transitions and a complete relaxation is required to reset the network, thus restarting the monotonic increase. This is even more evident in the walking session: in this case, the ramp-up is only partial, given the lower effort required. The network remains stable in state 1, waiting for a further increase in the HR and ignoring the noise on input 0.

IV. CONCLUSION

Our results show that a small and simple network, implemented with mixed-signal analog/digital neuromorphic circuits, can reliably monitor a monotonic HR trends over extended periods, paving the way for always-on health monitoring systems that are both efficient and long-lasting. The low power consumption of the neuromorphic circuits enables continuous operation for extended amounts of time, making it ideal for wearable devices for health monitoring. Future research will focus on replacing ECG with PPG signals measured at the wrist and addressing the challenges of noise and movement artifacts. Additionally, we will transition from healthy subjects to patients affected by neuropathologies, such as dementia, which impact HR over long periods, despite the absence of cardiac pathology.

This technology promises significant improvements in patient care and health management, especially in scenarios requiring constant monitoring and rapid response, thus enhancing overall quality of life.

REFERENCES

- [1] U. Rajendra Acharya, K. Paul Joseph, N. Kannathal, C. M. Lim, and J. S. Suri, "Heart rate variability: A review," *Medical and biological engineering and computing*, vol. 44, pp. 1031–1051, 2006.
- [2] Z. Zhu, H. Li, J. Xiao, W. Xu, and M.-C. Huang, "A fitness training optimization system based on heart rate prediction under different activities," *Methods*, vol. 205, pp. 89–96, 2022.
- [3] P. A. Heidenreich, J. G. Trogdon, O. A. Khavjou, *et al.*, "Forecasting the future of cardiovascular disease in the united states: A policy statement from the american heart association," *Circulation*, vol. 123, no. 8, pp. 933–944, 2011.
- [4] K. Y. Liu, E. A. Whitsel, G. Heiss, *et al.*, "Heart rate variability and risk of agitation in alzheimer's disease: The atherosclerosis risk in communities study," *Brain Communications*, vol. 5, no. 6, fcad269, 2023.
- [5] H. Davidoff, N. Van Helleputte, M. Vandenbulcke, *et al.*, "Unraveling the relationship between heart rate and agitation in people with dementia," *Alzheimer's & Dementia*, vol. 19, e078916, 2023.
- [6] S. Caspar, E. D. Davis, A. Douziech, and D. R. Scott, "Nonpharmacological management of behavioral and psychological symptoms of dementia: What works, in what circumstances, and why?" *Innovation in Aging*, vol. 1, no. 3, igy001, 2017.
- [7] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri, "Neuromorphic electronic circuits for building autonomous cognitive systems," *Proceedings of the IEEE*, vol. 102, no. 9, pp. 1367–1388, 2014, issn: 0018-9219.
- [8] E. Donati and G. Indiveri, "Neuromorphic bioelectronic medicine for nervous system interfaces: From neural computational primitives to medical applications," *Progress in Biomedical Engineering*, vol. 5, no. 1, p. 013002, 2023.
- [9] F. Bauer, D. Muir, and G. Indiveri, "Real-time ultra-low power ECG anomaly detection using an event-driven neuromorphic processor," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 13, no. 6, pp. 1575–1582, 2019.
- [10] A. Das, P. Pradhapan, W. Groenendaal, *et al.*, "Unsuper-vised heart-rate estimation in wearables with liquid states and a probabilistic readout," *Neural networks*, vol. 99, pp. 134–147, 2018.
- [11] M. Sharifshazileh, K. Burelo, J. Sarnthein, and G. Indiveri, "An electronic neuromorphic system for real-time detection of high frequency oscillations (HFOs) in intracranial EEG," *Nature Communications*, vol. 12, no. 1, pp. 1–14, 2021.
- [12] Y. Ma, E. Donati, B. Chen, *et al.*, "Neuromorphic implementation of a recurrent neural network for EMG classification," in *International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, 2020, IEEE, 2020, pp. 69–73.
- [13] A. Vitale, E. Donati, R. Germann, and M. Magno, "Neuromorphic edge computing for biomedical applications: Gesture classification using emg signals," *IEEE Sensors Journal*, vol. 22, no. 20, pp. 19490–19499, 2022.
- [14] E. Neftci, J. Binas, U. Rutishauser, *et al.*, "Synthesizing cognition in neuromorphic electronic systems," *Proceedings of the National Academy of Sciences*, vol. 110, no. 37, E3468–E3476, 2013.
- [15] D. Liang and G. Indiveri, "A neuromorphic computational primitive for robust context-dependent decision making and context-dependent stochastic computation," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 66, no. 5, pp. 843–847, 2019.
- [16] S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri, "A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs)," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 12, no. 1, pp. 106–122, 2018.
- [17] S. Deiss, R. Douglas, and A. Whatley, "A pulse-coded Communications infrastructure for neuromorphic systems," in *Pulsed Neural Networks*, W. Maass and C. Bishop, Eds., MIT Press, 1998, ch. 6, pp. 157–78.
- [18] D. Jarchi and A. J. Casson, "Description of a database containing wrist ppg signals recorded during physical exercise with both accelerometer and gyroscope measures of motion," *Data*, vol. 2, no. 1, p. 1, 2016.
- [19] S. Narayanan, M. Cartiglia, A. Rubino, *et al.*, "Spaic: A sub- μ w/channel, 16-channel general-purpose event-based analog front-end with dual-mode encoders," in *2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, IEEE, 2023, pp. 1–5.
- [20] M. Zanghieri, S. Benatti, L. Benini, and E. Donati, "Event-based low-power and low-latency regression method for hand kinematics from surface emg," in *2023 9th International Workshop on Advances in Sensors and Interfaces (IWASI)*, IEEE, 2023, pp. 293–298.
- [21] R. Sava, E. Donati, and G. Indiveri, "Feed-forward and recurrent inhibition for compressing and classifying high dynamic range biosignals in spiking neural network architectures," in *2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, IEEE, 2023, pp. 1–5.
- [22] M. Zanghieri, P. M. Rapa, M. Orlandi, *et al.*, "Event-based estimation of hand forces from high-density surface emg on a parallel ultra-low-power microcontroller," *IEEE Sensors Journal*, 2024.
- [23] N. Risi, A. Aimar, E. Donati, S. Solinas, and G. Indiveri, "A spike-based neuromorphic architecture of stereo vision," *Frontiers in Neurobotics*, vol. 14, p. 93, 2020, issn: 1662-5218.
- [24] J. Zhao, M. Monforte, G. Indiveri, C. Bartolozzi, and E. Donati, "Learning inverse kinematics using neural computational primitives on neuromorphic hardware," *npj Robotics*, vol. 1, no. 1, p. 1, 2023, issn: 2731-4278.