Coincidence Detection Is All You Need

Celestine Preetham Lawrence[+]
[+]Bernoulli Institute and Groningen Cognitive Systems and Materials Center (CogniGron),
University of Groningen, 9700 AB, Groningen, Netherlands.

Abstract

This paper demonstrates that the performance of coincidence detection - a classic neuromorphic signal processing method found in Rosenblatt's perceptrons with distributed transmission times, can be competitive to a state-of-the-art deep learning method for pattern recognition. Hence, we cannot remain comfortably numb to the prevailing dogma that efficient matrix-vector operations is all we need; but should enquire with greater vigour if more advanced continual learning methods (running on spiking neural network hardware with neuromodulatory mechanisms at multiple timescales) can beat the accuracy of task-specific deep learning methods. With regards to deployability, coincidence detection is an interpretable shallow learning method and its applications provide a commercial use-case for neuromorphic hardware such as Intel Loihi.

1 Introduction

Frank Rosenblatt and his team (1957-1971) built and analyzed several kinds of perceptrons [1][2][3][4] - networks of sensory, association and receptor neurons; which in contemporary deep learning terminology relates to the input, hidden and output layers. The propagating signals were binary (compatible with a spike-based view), the synaptic delays (transmission times) and weights (memory states) could be analog, the network could be recurrent and was often randomly interconnected, and learning often meant tuning the weights of the association-receptor subnetwork by some error-corrective reinforcement. The synaptic delays were not learnt but instead randomly distributed in Rosenblatt's Tobermory perceptrons [5], and this was rich enough to realize concentration-invariant and uniform time-warp invariant spatiotemporal classification by logarithmic encoding and coincidence detection. However, the processing speed of commercial Von Neumann computers advanced exponentially and outperformed neuromorphic hardware on yesterdecade's benchmarks [6]. The Tobermory perceptron was forgotten, nevertheless, the utility of logarithmic encoding and coincidence detection was formalized by John Hopfield [7] as an efficient solution to the *analog match* problem in pattern recognition.

Now, half a century after the accidental demise of Rosenblatt, neuromorphic signal processors are making a comeback. For example, (1) Intel's Loihi with spike-time dependent plasticity mechanisms for learning olfactory pattern recognizers [8]; (2) Physical reservoir computing networks [9] where the interconnectivity of the hidden layer is unchanged, closer to the spirit of Rosenblatt's randomly interconnected sensory-association subnetwork.

Here, to strengthen the case for revisiting classic methods on novel and modern hardware, we evaluate the performance of coincidence detection in comparison to a deep learning method. Nothing more, nothing less, although this work was triggered by a rabid interest in employing artificial intelligence to sniff out infections and prevent future pandemics.

2 Methods

Here, we consider the work [10] of an interdisciplinary team, where a 26 layer convolutional neural network with residual connections (ResNet-26) was successfully trained for classifying pathogenic bacteria by Raman spectroscopy. In their work, there are $N = 30$ classes of bacterial isolates and they begin with a ResNet-26 pre-trained on $N \times 2000$ spectra, then for each class $n = 1 : N$ there are $M = 100$ training spectra, and similarly $N \times M = 3000$ test spectra. Each spectrum $\boldsymbol{x}$ contains 1000 floating-point numbers ranging between 0 and 1. Although compute intensive, their deep learning method proved to be a tool of great convenience for pattern recognition in a challenging dataset, where intra-isolate spectra were often more dissimilar than inter-isolate spectra.

Our method to tackle the above dataset, is inspired by the theory of how coincidence detection [7] in animal brains is fundamental for odour classification in complex and turbulent mixtures. Each class n has a vector representation

Table 1: Test accuracy (%)

ResNet-26	Coincidence detection
82.2± 0.3 (from [10])	82.7 (this work)

$\boldsymbol{w}_n$ that is learnt, and an input vector $\boldsymbol{x}$ results in an output class $y(\boldsymbol{x}) = \arg_n \max(\boldsymbol{x} \bigwedge \boldsymbol{w}_n)$ where we introduce the operator $\bigwedge$ to represent the coincidence between two signals. The analytical nature of coincidence detection depends on the specificities of the ion-channels and the membranes involved [11], and may even incorporate nonlinear leaky-integrate [12] multiple timescale mechanisms. We do not yet have a complete theory of neuromorphic signal processing, so here we introduce an approximation for the translation and scale-invariant property of coincidence detection as

$$\arg_n \max(\boldsymbol{x} \bigwedge \boldsymbol{w}_n) \approx \arg_n \max(\boldsymbol{w}_n \cdot \hat{\boldsymbol{x}}), \tag{1}$$

where $\hat{\boldsymbol{x}}$ is the zero-mean unit-variance normalization of $\boldsymbol{x}$.

Thus, the approximation in Eq. (1) allows $y(\boldsymbol{x})$ to be learnt by a logistic regression on the normalized dataset. We discard the pre-training data, pre-process the training and test spectra by a range-1 mean filter, and use the default method for logistic regression in Wolfram Mathematica (L2-regularization = 0.0001, optimization method = limited-memory BFGS). Code is provided in the supplemental material for reproducibility.

3 Result and outlook

The coincidence detection (via normalized logistic regression) method introduced here achieves a test accuracy greater than ResNet-26 (see Table 1), and it took less than 3 seconds to train the classifier on a modern desktop (without any special-purpose GPUs). Check `https://openreview.net/attachment?id=xT5rDp5VqKO&name=supplementary_material` for Wolfram Mathematica and Python code and plots of the training and test data, and confusion matrices. Note that the training data was fit all at once to a 100% accuracy. With a more neuromorphic coincidence detection method and a learning method that adapts the synaptic delays $\boldsymbol{w}$ continually, to keep track under changing environmental conditions, we may achieve even greater accuracies.

Reviewer contributions

This paper has been previously reviewed at NeurIPS 2022 (`https://openreview.net/forum?id=xT5rDp5VqKO`) but not recommended for immediate publication for reasons including that it has been only tested on a single dataset. I believe it is good to present this work in a reasonable venue and thereby motivate stakeholders to test coincidence detection on more datasets. Here below, I summarize relevant contributions as author responses to a selection of reviews. Note that the review process also revealed a typo in the supplementary material, where it was wrongly commented that "standardization is performed across samples..." - it should instead read as "standardization is performed samplewise - each sample has a zero-mean and unit-variance across its features...".

Reviewer V6Wx: The simple "coincidence" detector gives very good results compared with a deep net. Although this could be demonstrating an advantage of coincidence detection, it may also be that the classification problem is actually not that difficult. Paper [10] seems to only apply a deep net to the problem. The authors only apply a linear function. What do other functions do? k-nearest neighbors, SVMs, ...?

1. Is there no more suitable implementation of coincidence detection, e.g., within a spiking net?
2. Is your model in eq 1 not simply a perceptron? (With normalized inputs and a max on the outputs)

Response: Ref. [10] already explored traditional methods (k-NN, SVM) and justified their choice for a deep learning method.

1. Yes, references [11] and [12] point to this, but are expensive to implement on conventional hardware. Future work should compare how the approximate implementation of coincidence detection compares to more advanced methods on neuromorphic hardware.

2. Yes, is it not beautiful? Did you notice that the normalization is performed across a different axis in comparison to the standard suggestion of Python sklearn for logistic regression? (Conventional wisdom is that it is a bad idea to do a normalization in this way, which is why perceptrons were not employed with this kind of pre-processing until now. This paper instead argues from the theory of coincidence detection that it is actually a

good idea for preprocessing datasets that are compatible with the analog match problem, which turns out to be true upon evaluation in this empirical dataset.)

Reviewer ctyh: There is an interesting empirical observation here, yet the narrative is too shallow...

Response: The result in table-1 speaks for itself (i.e. here is a novel method with better performance in comparison to the impactful deep learning method by a large team of researchers in Stanford university, cited over 300 times). Of course, this novel method will need to be applied to other datasets (which is why it needs to be presented in a conference to gain the attention of fellow researchers). Moreover, references [7], [11], [12] have been thoughtfully chosen as related work.

Reviewer QphW: Authors should consider generating more stats on their accuracy % and provide a more thorough comparison with the baseline (ResNet-26). Further, authors should share additional experiments breaking down the contribution of standardization and smoothing steps. Lastly, explaining why their model fares better than the deep learning model...

Response: The reviewer asks for more stats, but is it not futile? Given that this is anyhow based on performance on a single dataset? The focus of this paper is to demonstrate that the approximation for coincidence detection introduced here is able to solve an analog match problem (discussed insightfully by Hopfield [7], but not as well-known as it should be). That the model fares slightly better is a bonus, actually deep learning methods can surely learn a coincidence detector (albeit in a computationally expensive way). Moreover, in order to ensure reproducibility, the method was tested in two programming languages Mathematica (yielding an accuracy of 82.7% as reported in the main text) and Python (yielding an accuracy of 82.9% as reported in the supplementary material).

References

[1] Frank Rosenblatt. *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical Laboratory, Inc. Report no. 85-460-1, 1957.

[2] Frank Rosenblatt. *The perceptron: A theory of statistical separability in cognitive systems*. Cornell Aeronautical Laboratory, Inc. Report no. VG-1196-G-1, 1958.

[3] Frank Rosenblatt. *Principles of neurodynamics. perceptrons and the theory of brain mechanisms*. Cornell Aeronautical Laboratory, Inc. Report no. 1196-G-8, 1961.

[4] Frank Rosenblatt. Cognitive systems research program. Technical report, Cornell University, Ithaca, New York, 1971.

[5] Frank Rosenblatt. A description of the tobermory perceptron. In *Collected Technical Papers*, volume 2. Cornell University, Ithaca, New York, 1963.

[6] George Nagy. Neural networks-then and now. *IEEE Transactions on Neural Networks*, 2(2):316–318, 1991.

[7] John J Hopfield. Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376(6535):33–36, 1995.

[8] Nabil Imam and Thomas A Cleland. Rapid online learning and robust recall in a neuromorphic olfactory circuit. *Nature Machine Intelligence*, 2(3):181–191, 2020.

[9] G. Tanaka, T. Yamane, J.B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose. Recent advances in physical reservoir computing: A review. *Neural Networks*, 115:100–123, 2019.

[10] Chi-Sing Ho, Neal Jean, Catherine A Hogan, Lena Blackmon, Stefanie S Jeffrey, Mark Holodniy, Niaz Banaei, Amr AE Saleh, Stefano Ermon, and Jennifer Dionne. Rapid identification of pathogenic bacteria using raman spectroscopy and deep learning. *Nature communications*, 10(1):1–8, 2019.

[11] Nelson Spruston. Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9(3):206–221, 2008.

[12] Wondimu Teka, Toma M Marinov, and Fidel Santamaria. Neuronal spike timing adaptation described with a fractional leaky integrate-and-fire model. *PLoS computational biology*, 10(3):e1003526, 2014.