Robust unsupervised learning of spiking motifs with optimal transport theory

Antoine Grimaldi[1], Andrea Alamia[1], Matthieu Gilson[2], Laurent Perrinet[2], Boris Sotomayor-Gomez[3], & Martin Vinck[3]
[1]Centre de Recherche Cerveau et Cognition (UMR 5549); Université Paul Sabatier, CNRS; Toulouse, France
[2]Institut de Neurosciences de la Timone (UMR 7289); Aix Marseille Université, CNRS; Marseille, France
[3]Ernst Strüngmann Institute; Frankfurt-am-Main, Germany
antoine.grimaldi@cnrs.fr

Abstract

- Temporal sequences are an important feature of neural information processing in biology. Unsupervised clustering of repeating spike patterns (i.e. spiking motifs) observed in neurobiological has been addressed by several methods [1–4] but remains a challenge.
- We propose a new method to extract spike patterns embedded in raster plots using a 1D convolutional autoencoder.
- We show that using the Earth Mover's Distance as a loss function has interesting properties to extract spiking motifs.

An autoencoder for spiking motifs extraction

Illustration

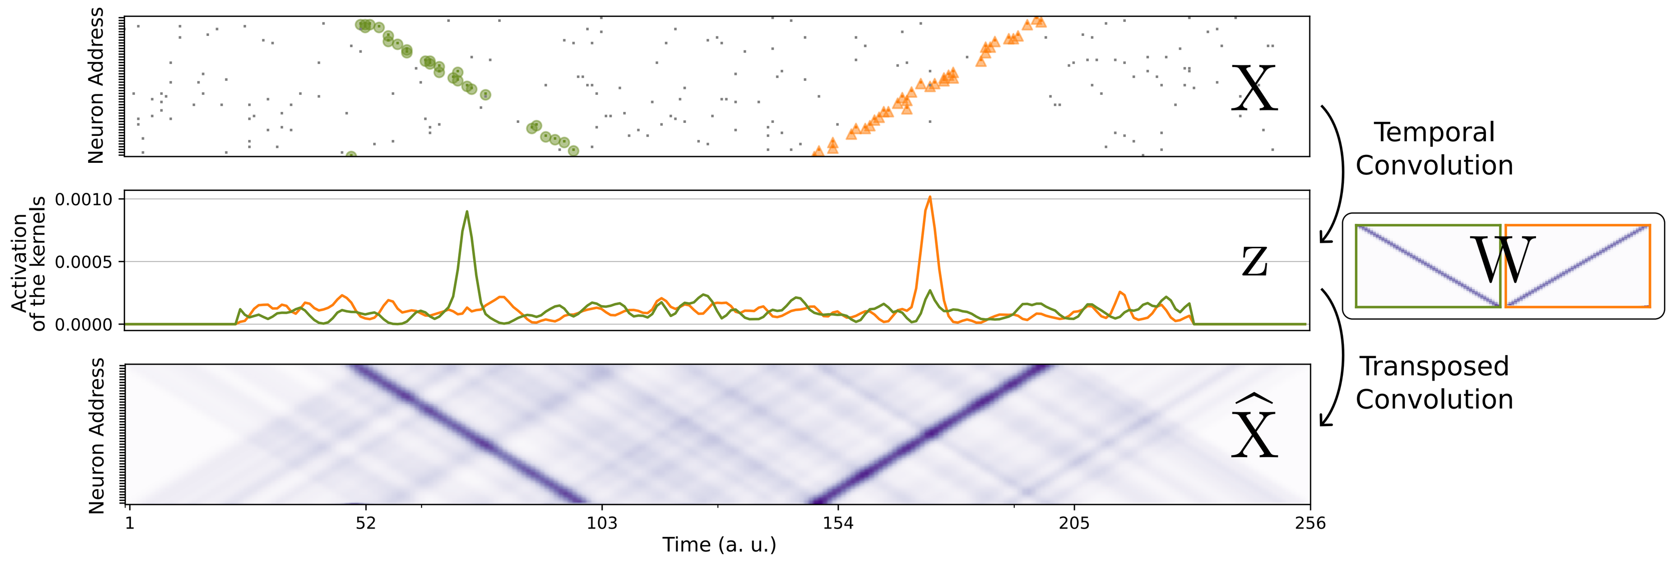

Figure 1: A single-layer autoencoder (AE) learns to represent the spike train as input (X) with a given number of kernels (W). If the decoding weights and the latent variables (z) are forced to be positive, the output of the AE ($\hat{X}$) performs Non-Negative Matrix Factorisation (NMF).

1D Wasserstein distance or Earth Mover's Distance (EMD)

Between two 1D probability distributions μ and ν, the EMD is given by the difference between the two cumulative distributions F_μ and F_ν:

$$\mathrm{EMD}(\mu,\nu,t) = \int_0^t |F_\mu(x) - F_\nu(x)|\,dx$$

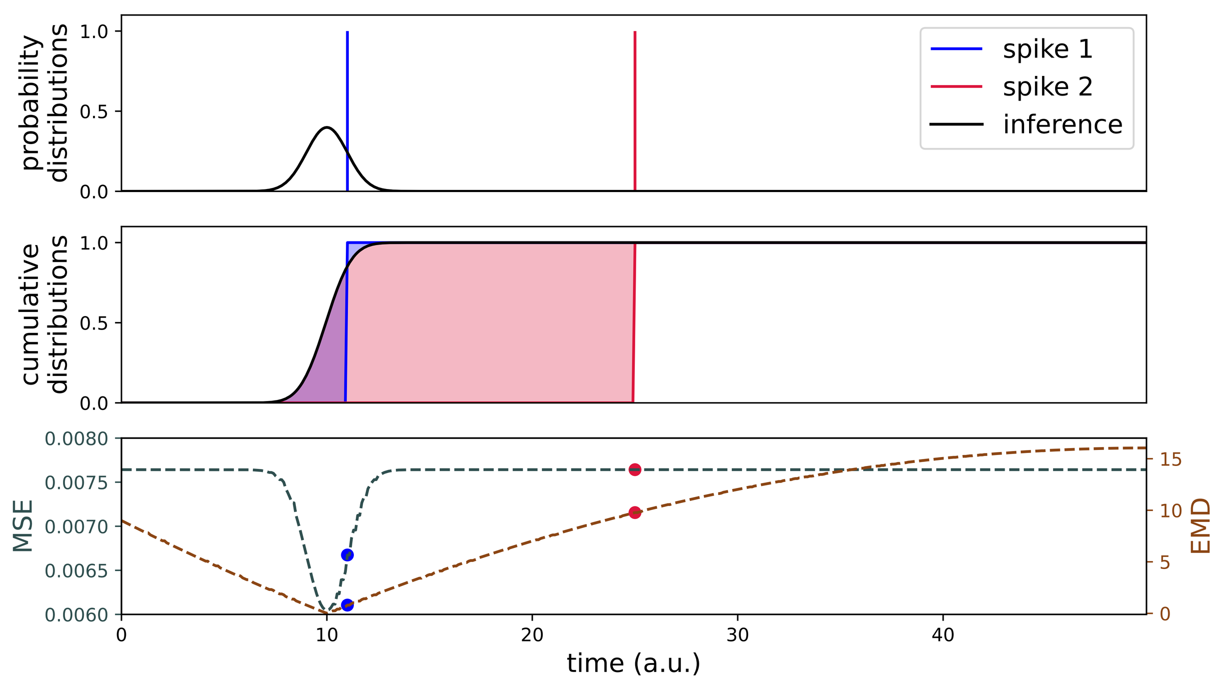

Figure 2: Comparison of the EMD and MSE values (lower panel) between a probability distribution (black) and binary spikes (blue and pink) shown in the upper panel. While the MSE saturates, the EMD continues to increase as the difference between the timing of the spike and the mean of the Gaussian increases.

Results

A qualitative example of spiking motifs clustering in synthetic data

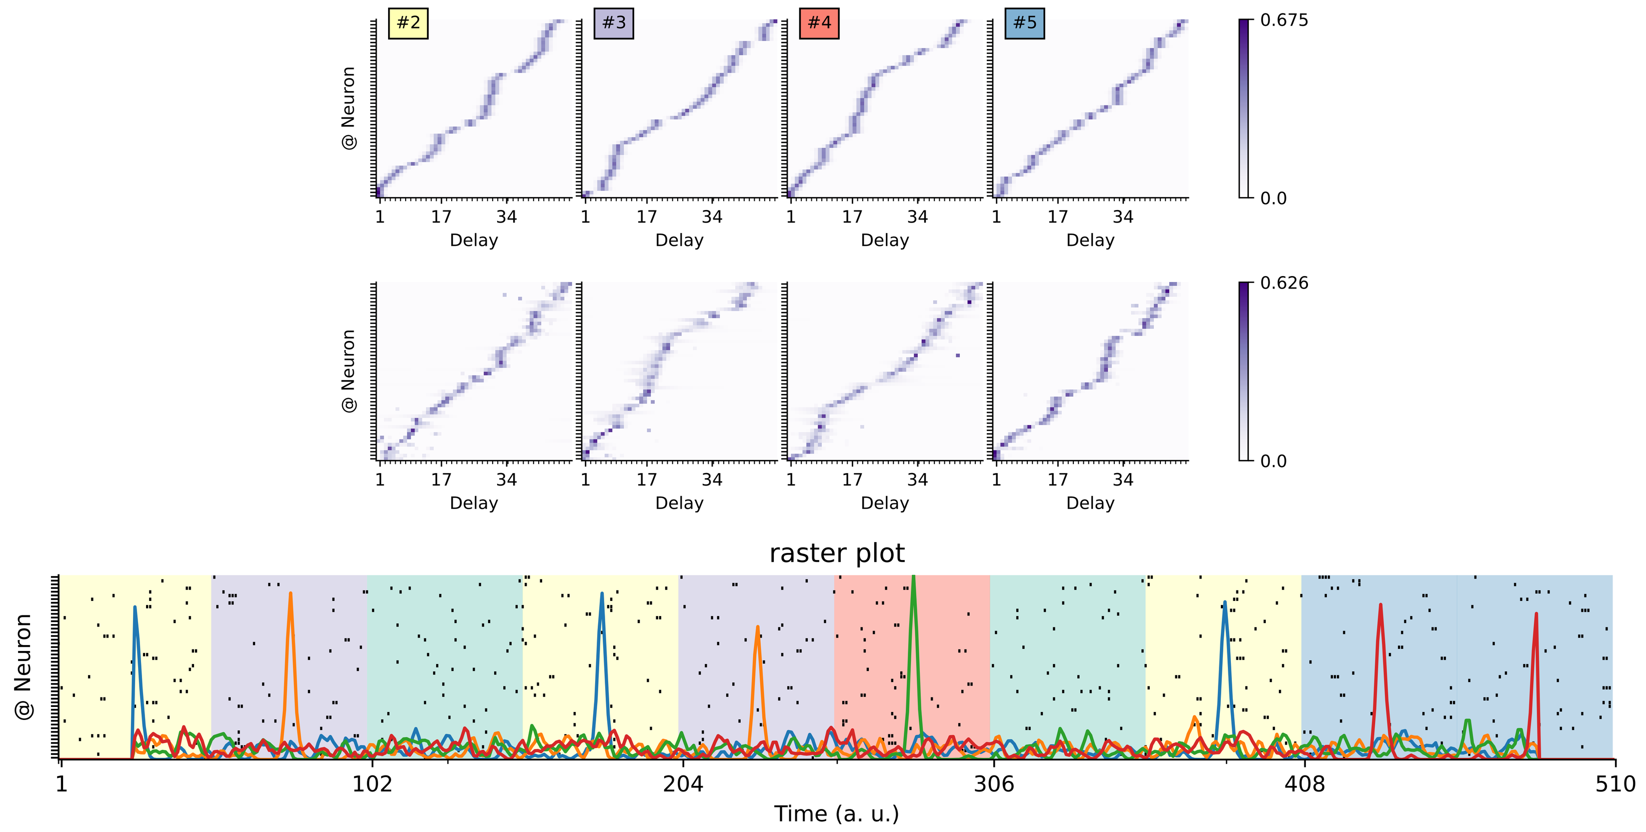

Figure 3: The upper panel represents the ground truth, i.e. 4 structured probability distributions used to generate spike trains (the neurons are sorted to better visualise the patterns). These distributions are randomly concatenated and a Bernoulli trial is applied on the sequence to generate a raster plot. The middle panel represents the different kernels learned by the autoencoder minimizing the EMD on a synthetic dataset of 60 samples. The bottom panel shows a synthetic raster plot with the temporal activation of the different kernels of the AE after training.

Performance as a function of the number of epochs for the static case (no convolution)

- no temporal convolution, i.e. the size of the AE kernel is the same as the size of the samples
- linear AE with tied positive weights
- 4 noise conditions with different jitter (2% or 10% of the temporal length of the sample) and different amount of spontaneous activity or additive noise (10% and 50%)

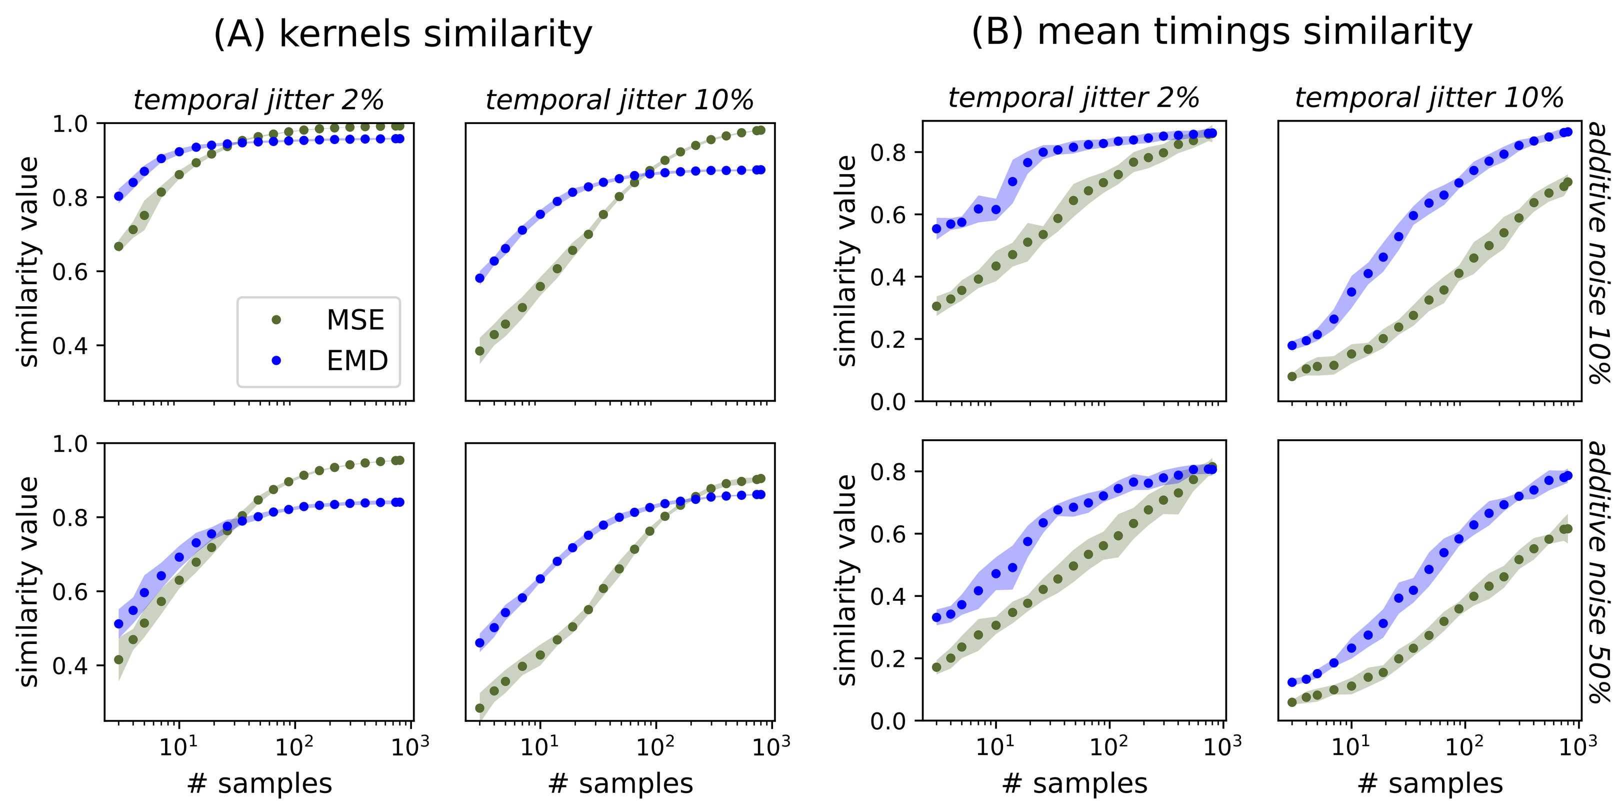

Figure 4: Performance of an AE trained with the MSE (green) and the EMD (blue) as a function of the number of samples used during training. Two metrics are used to assess the similarity of the learned kernel with the ground truth (the *kernels similarity* equivalent to the cosine similarity and the *mean timings similarity* which only takes into account the relative latency of the spikes (mean of the Gaussian distribution used to generate spikes), and ignores their precision (standard deviation of the distribution). We obtain a better similarity values for the method with EMD when a small amount of samples is available.

Robustness to different types of noise

- temporal convolution (fully unsupervised learning)
- nonlinear (sigmoid) AE with positive weights for the decoder
- single spiking motif
- different types of noise are tested: temporal jitter, additive noise, dropout probability

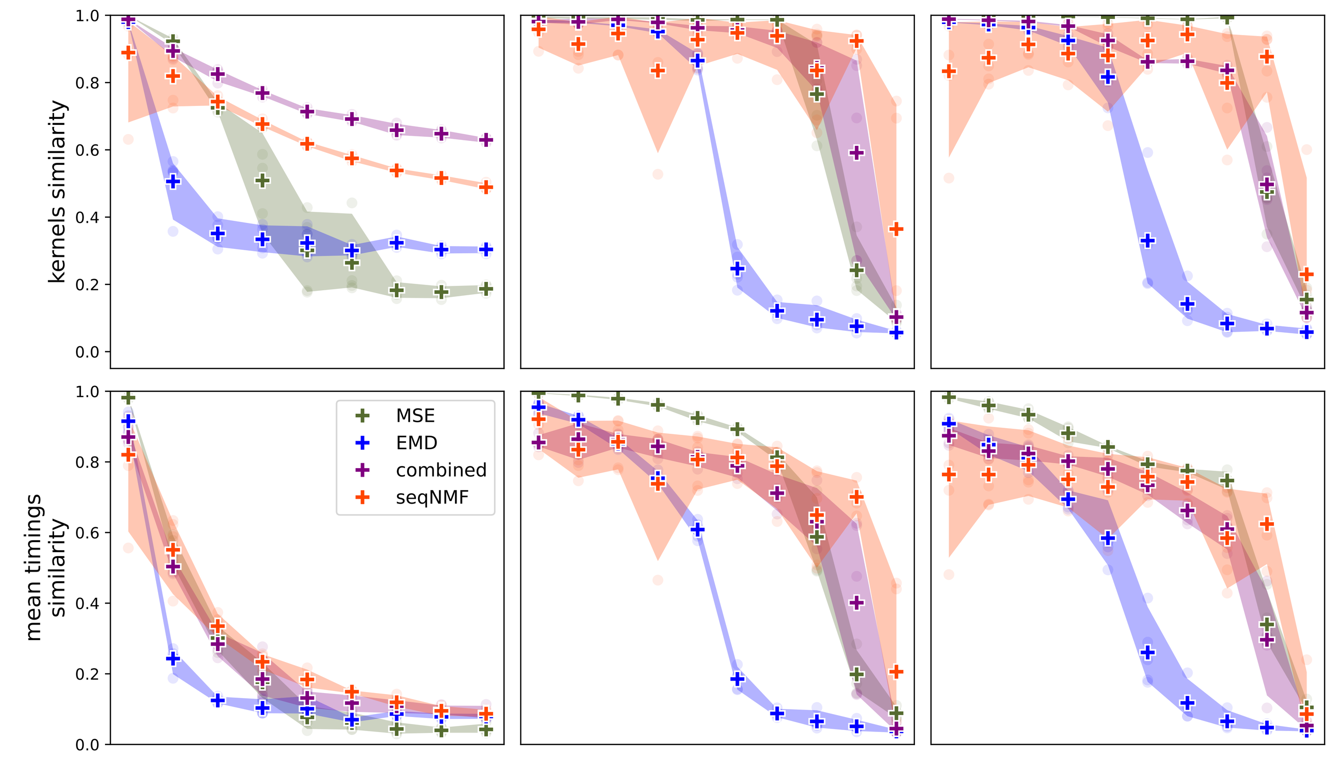

Figure 5: Performance of the AE to extract a single spiking motif embedded in a raster plot with different types of noise. The AE is trained with the EMD (blue), with the MSE (green) and iteratively with both losses (purple). We also show the performance of seqNMF [1] applied on this synthetic dataset. Depending on the type of noise, the loss choice can change to obtain better performance.

Learning multiple spiking motifs

- temporal convolution (fully unsupervised learning)
- nonlinear (sigmoid) AE with positive weights for the decoder
- multiple spiking motif
- relatively low amounts of noise

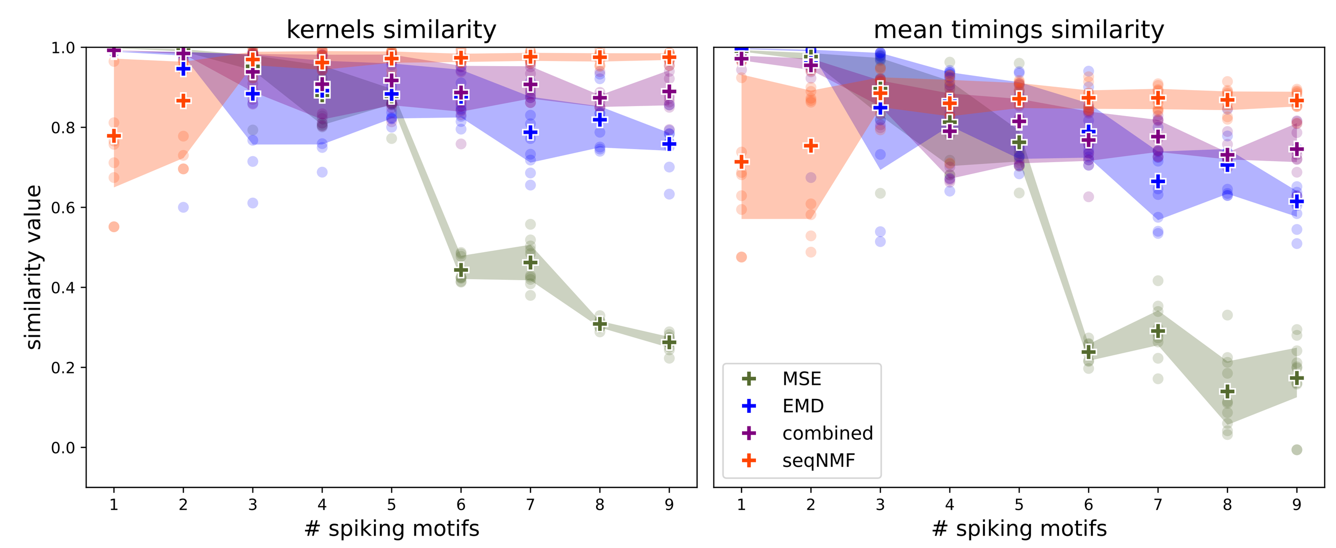

Figure 6: Performance of the extraction of multiple motifs for AEs trained with different losses and seqNMF. For the AE, training with the EMD gets better results when increasing the amount of motifs.

Conclusion

- Novel, simple and scalable method to extract spiking motifs in neural recordings
- Training with the EMD shows promising results at extracting spiking motifs especially when the available number of samples is limited
- The method can be applied to fully unsupervised clustering of multiple spiking motifs with a temporal convolution

References

[1] E. L. Mackevicius et al. "Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience". In: *eLife* 8 (Feb. 2019). Ed. by L. Colgin et al. Publisher: eLife Sciences Publications, Ltd, e38471.

[2] A. Williams et al. "Point process models for sequence detection in high-dimensional neural spike trains". In: *Advances in Neural Information Processing Systems*. Vol. 33. Curran Associates, Inc., 2020, pp. 14350–14361.

[3] C. Stringer et al. *Rastermap: a discovery method for neural population recordings*. en. preprint. Neuroscience, July 2023.

[4] R. Koshkin et al. *convSeq: Fast and Scalable Method for Detecting Patterns in Spike Data*. en. arXiv:2402.01130 [eess]. May 2024.