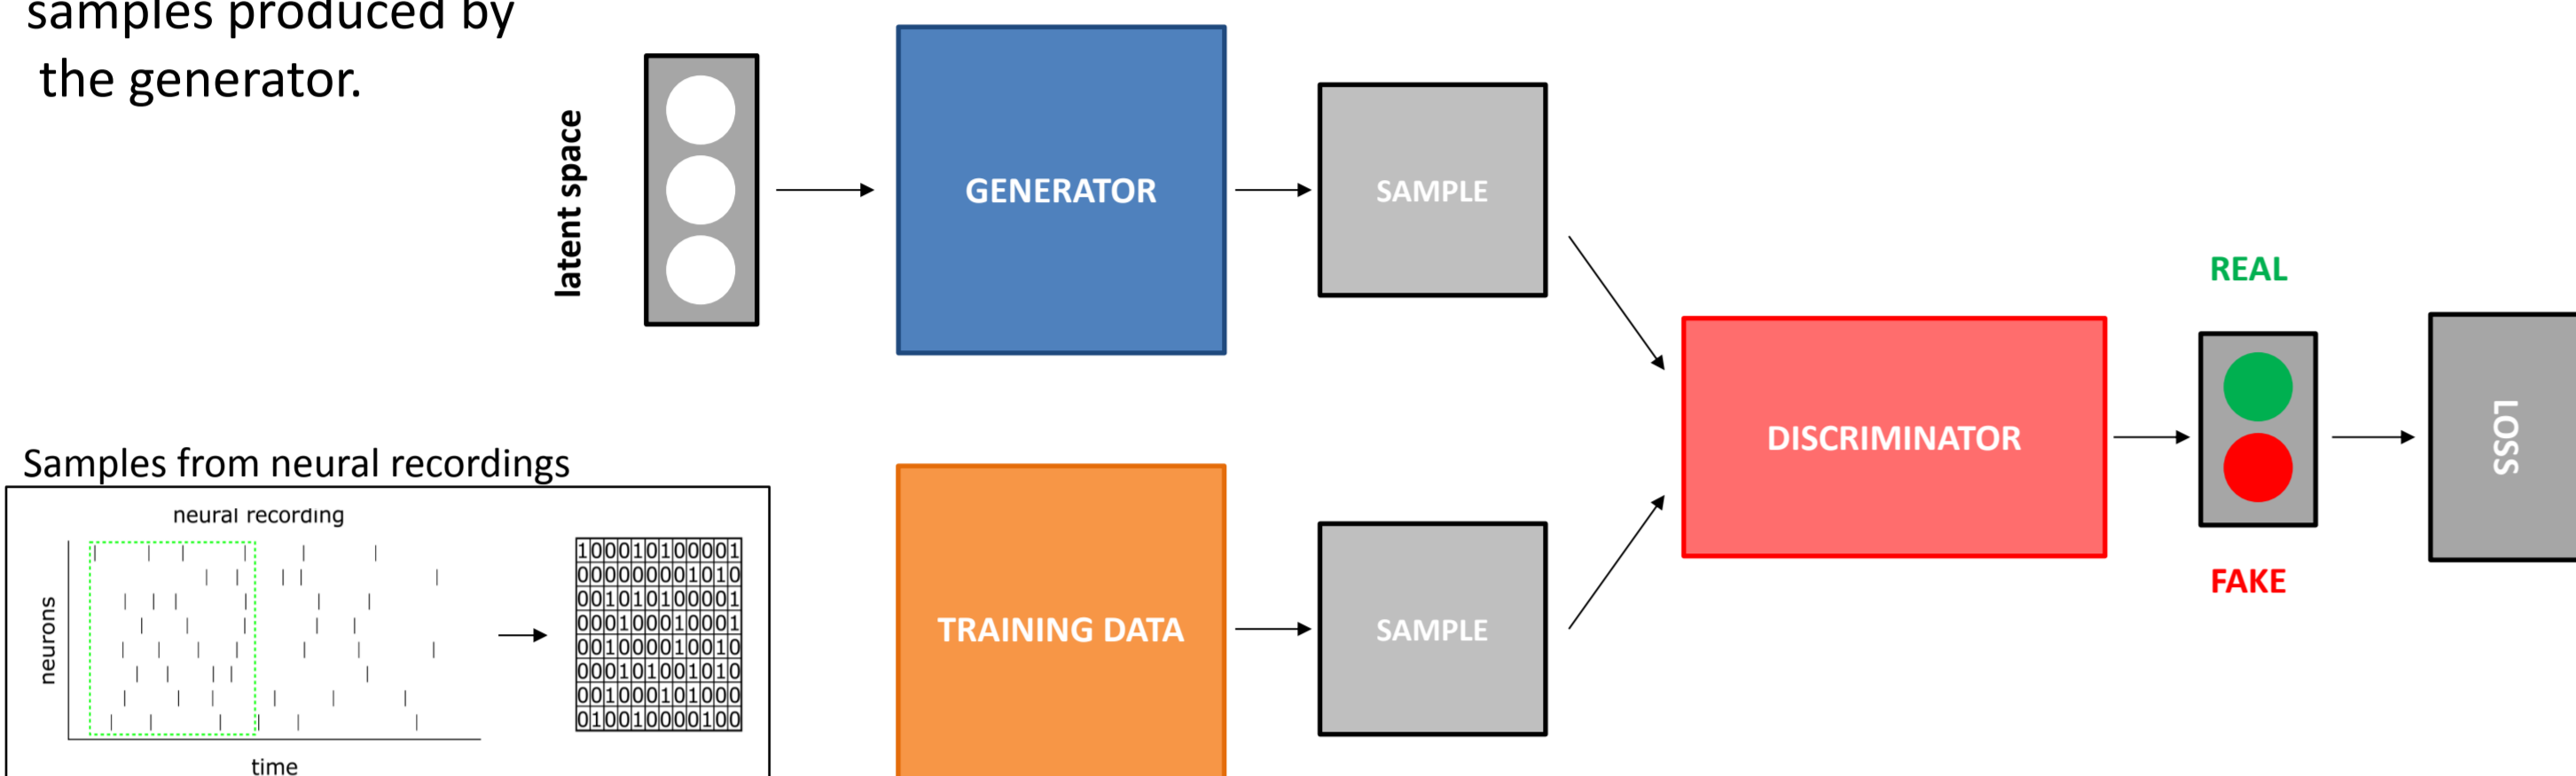


Summary

- We use **GANs** to **simulate the concerted activity of a population of neurons**.
- Spike-GAN** generates spike trains that match the **first- and second-order statistics** of populations of tens of neurons.
- We apply Spike-GAN to a **real dataset recorded from salamander retina** and showed that it performs as well as **state-of-the-art approaches**.
- We exploit a trained Spike-GAN to construct **importance maps** to detect the most relevant statistical structures present in a spike train.

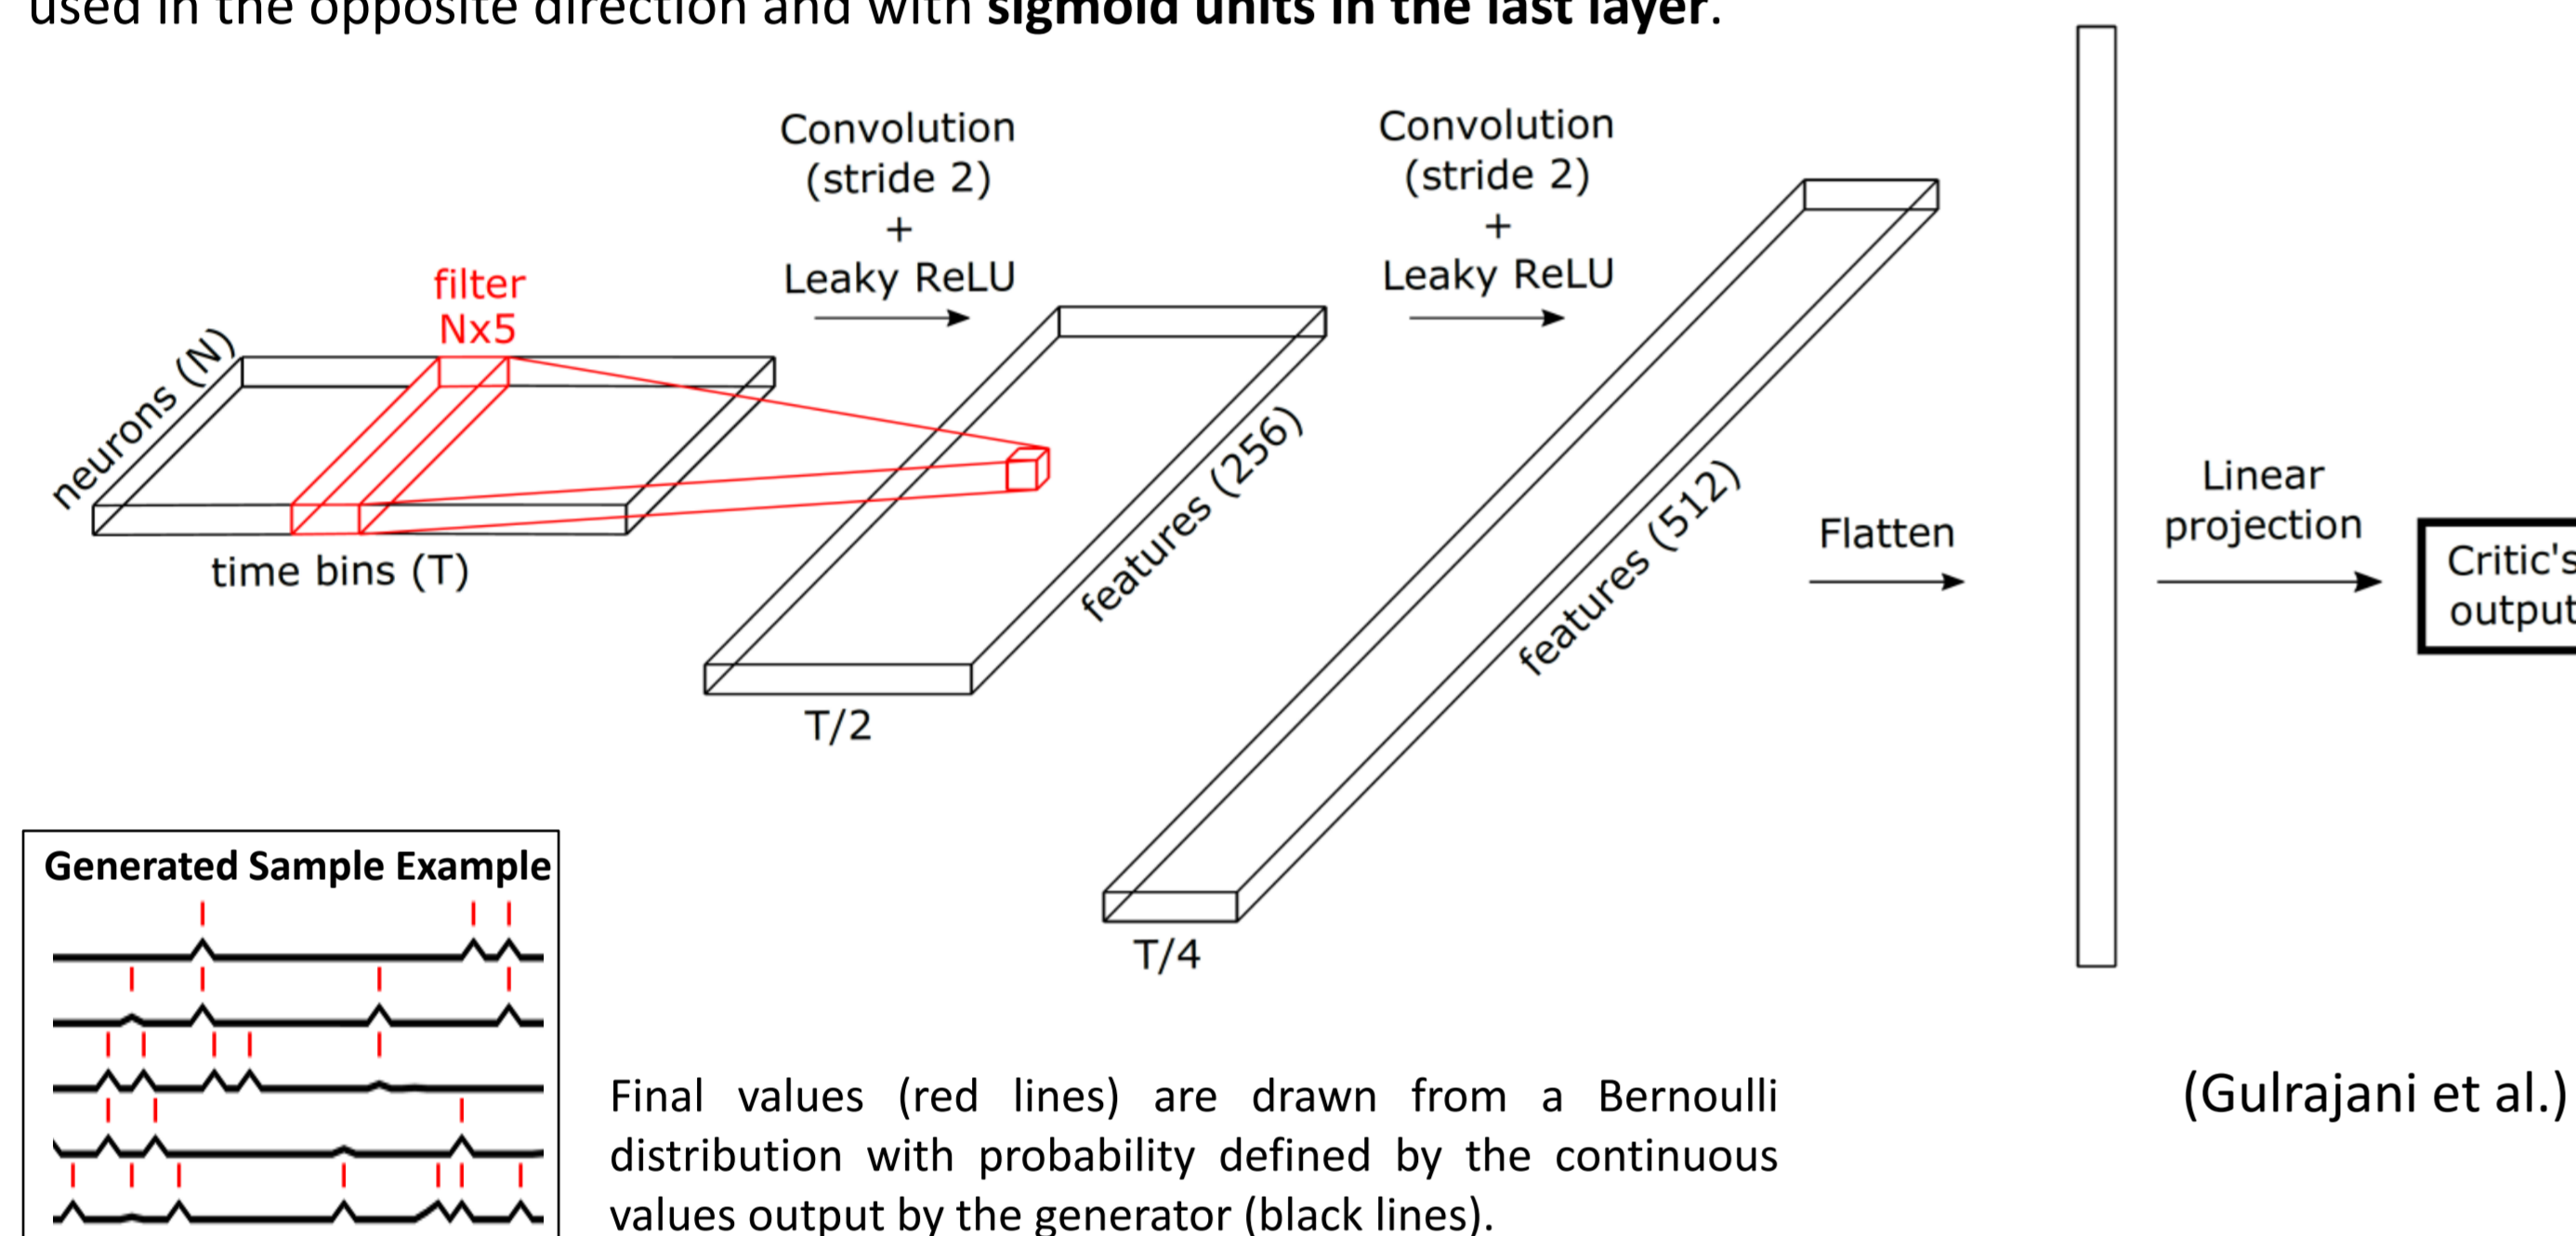
Generative Adversarial Networks

Generative Adversarial Networks (GANs) (Goodfellow et al.) are based on the competition between two deep neural networks: **the generator** tries to produce samples that are indistinguishable from the ones contained in a given training dataset; **the discriminator** aims at distinguishing between samples from the training dataset and samples produced by the generator.



Spike-GAN

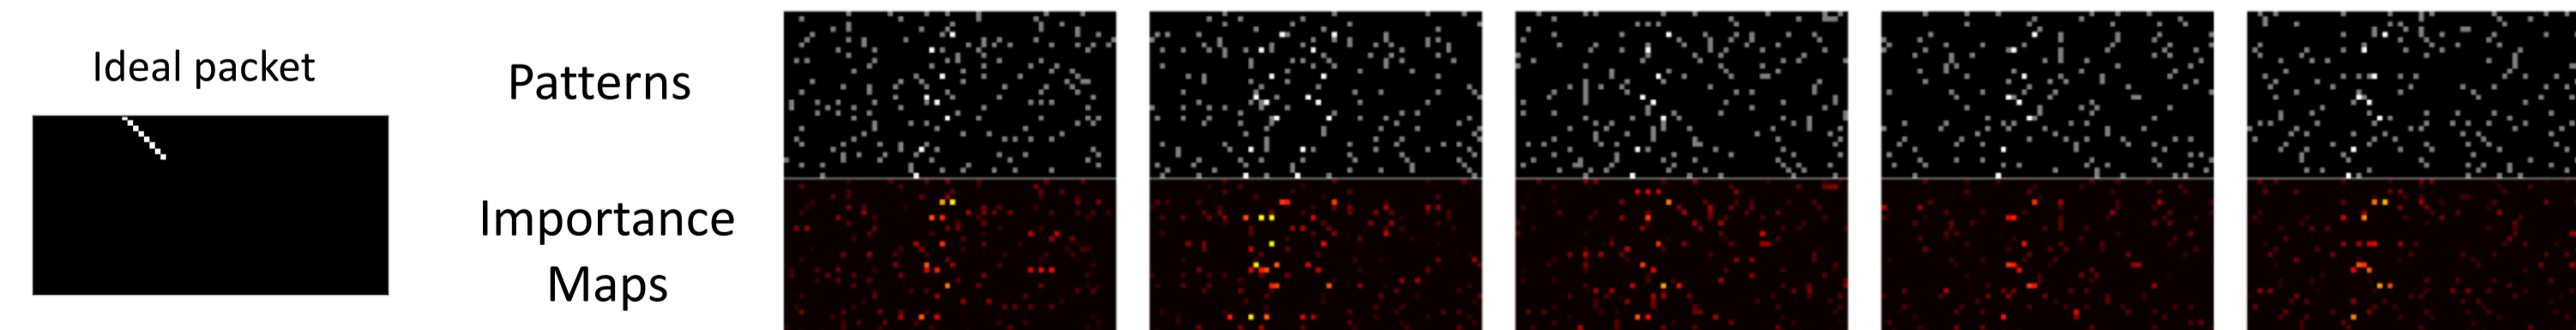
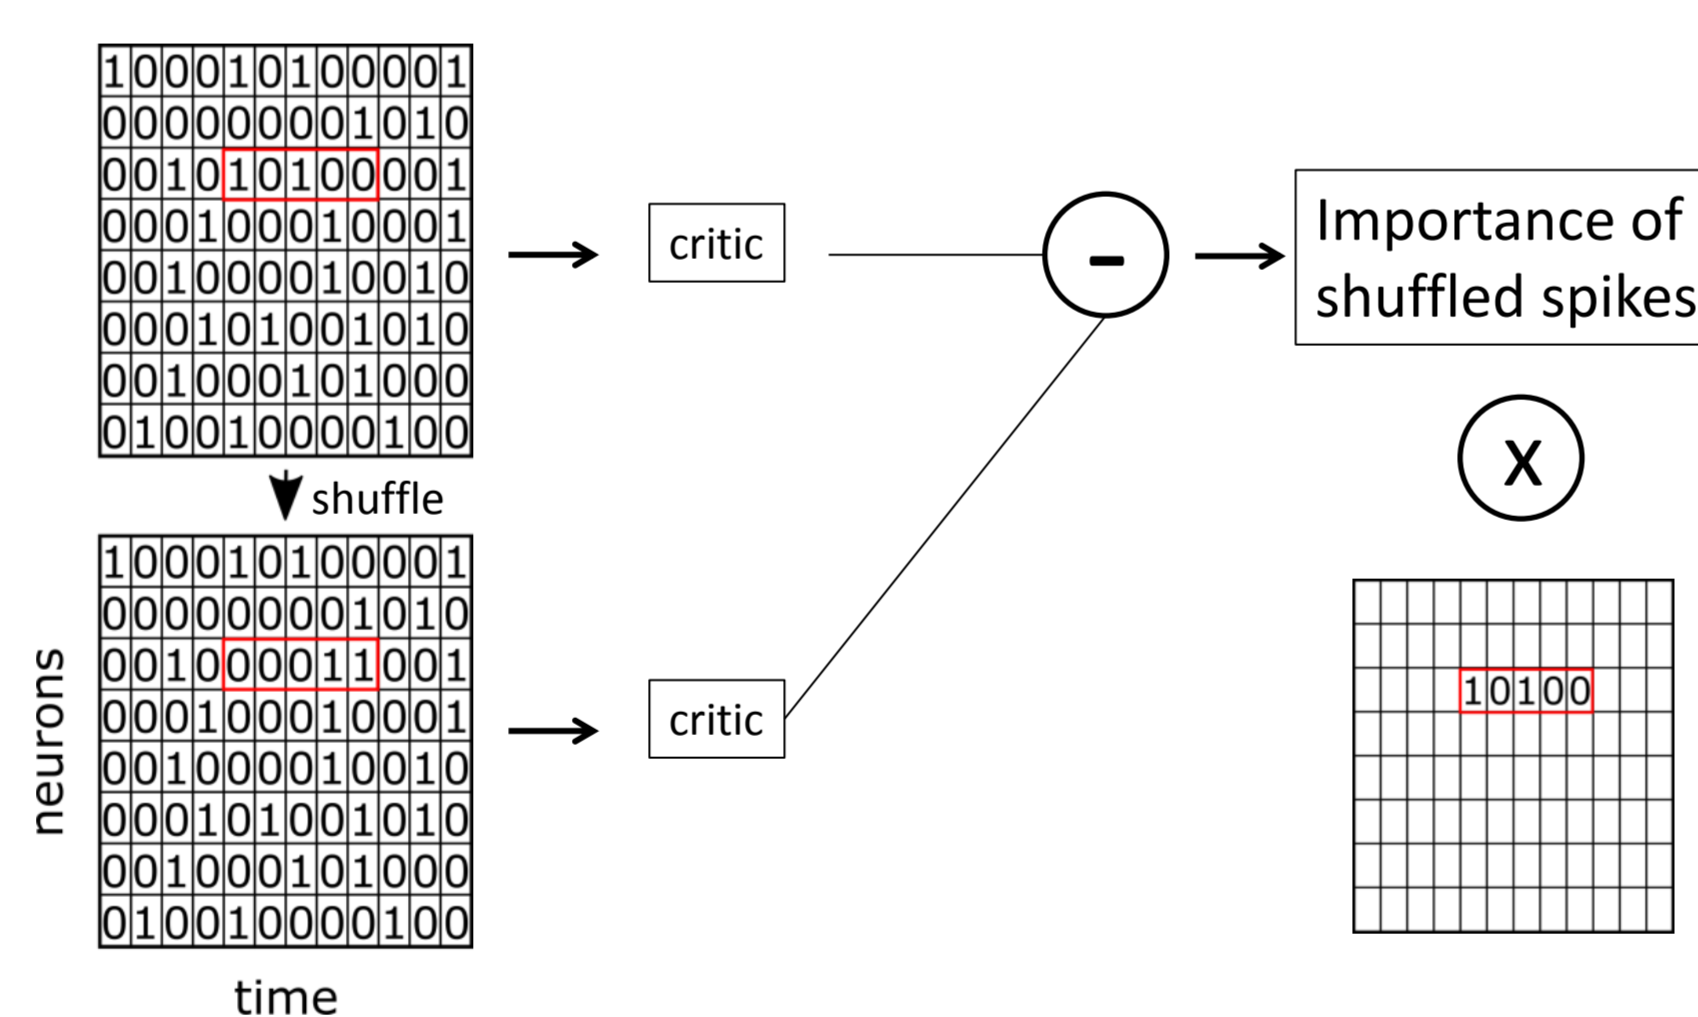
Spike-GAN adapts the Wasserstein-GAN variant described by Arjovsky et al.. **Samples are transposed** so as to input the neurons' activities into different channels. The convolutional filters (red box) **span thus all neurons but share weights across the time dimension**. The architecture of the generator is the same as that of the critic, used in the opposite direction and with **sigmoid units in the last layer**.



Finding relevant patterns of neural activity

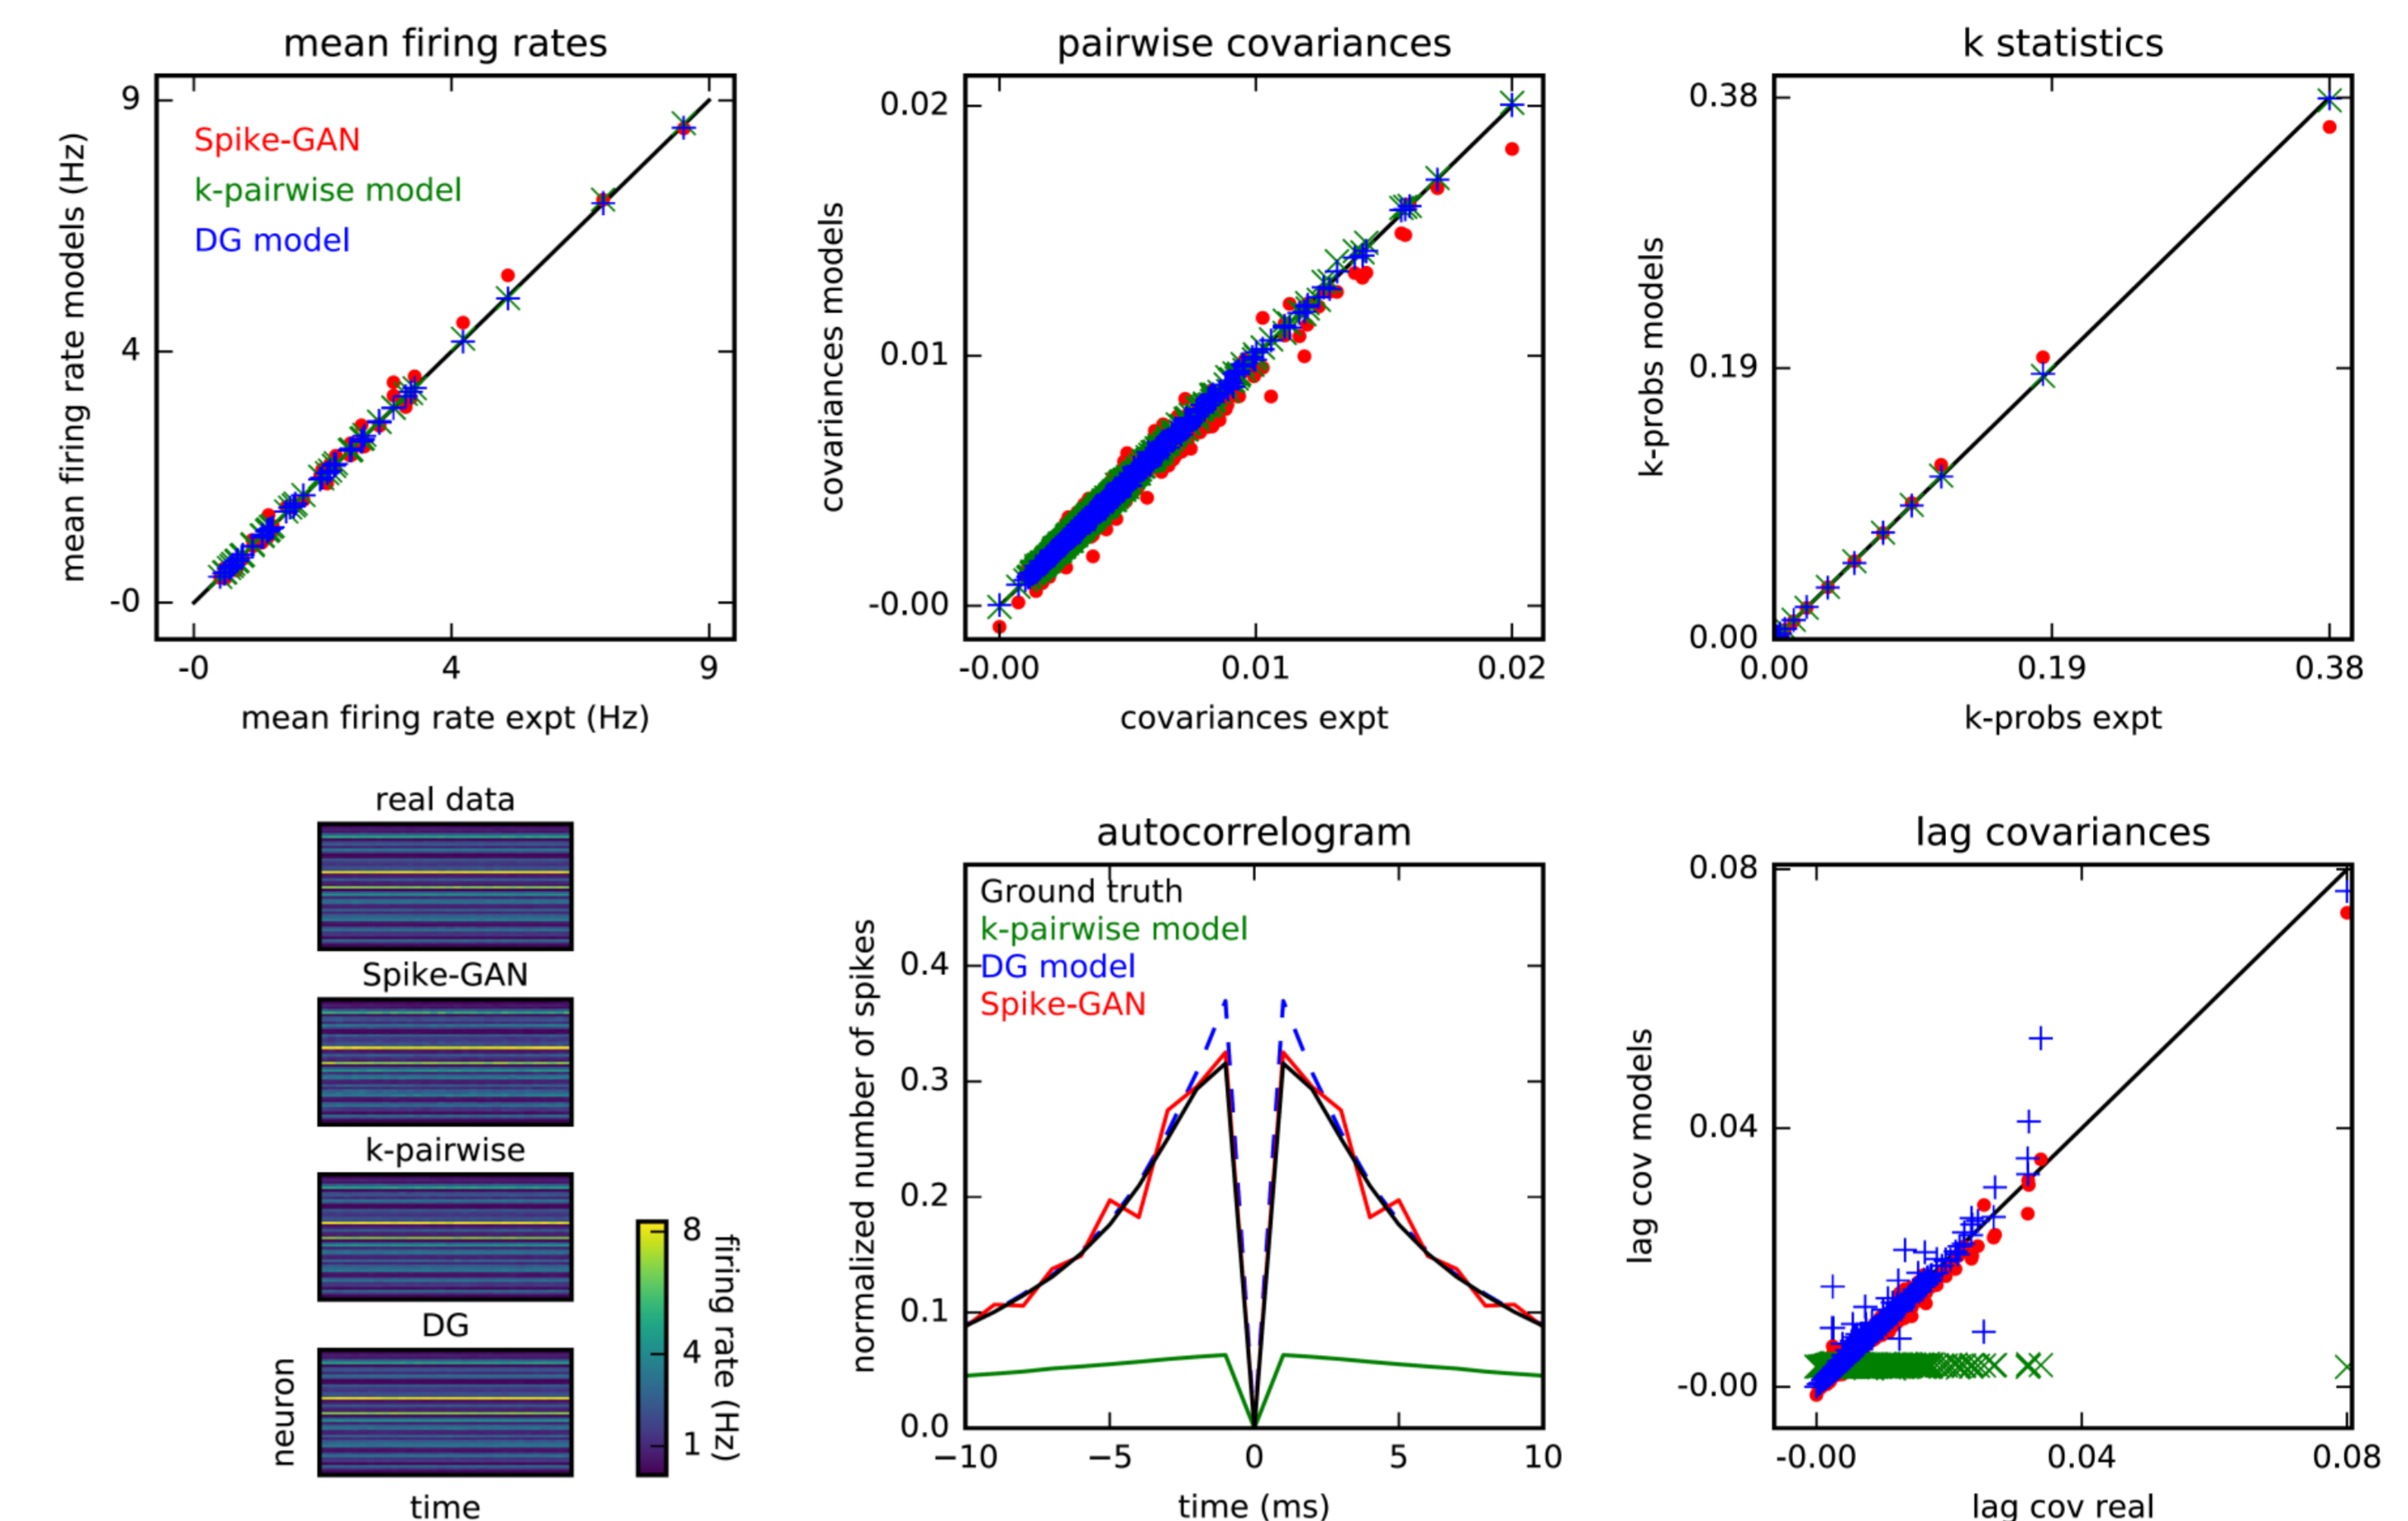
Importance maps: We infer the most relevant features characterizing a given neural activity pattern by enquiring a trained critic:

1. **Compute the output produced by the critic for that particular pattern.**
2. **Shuffle across time the spikes emitted by a neuron during a specific period of time and compute again the output of the critic.**
3. **The absolute difference between the two outputs gives the importance of shuffled spikes.**
4. **Multiply the masked original sample by the importance.**
5. **Sum up all resulting maps for all neurons and time periods to get the importance map.**



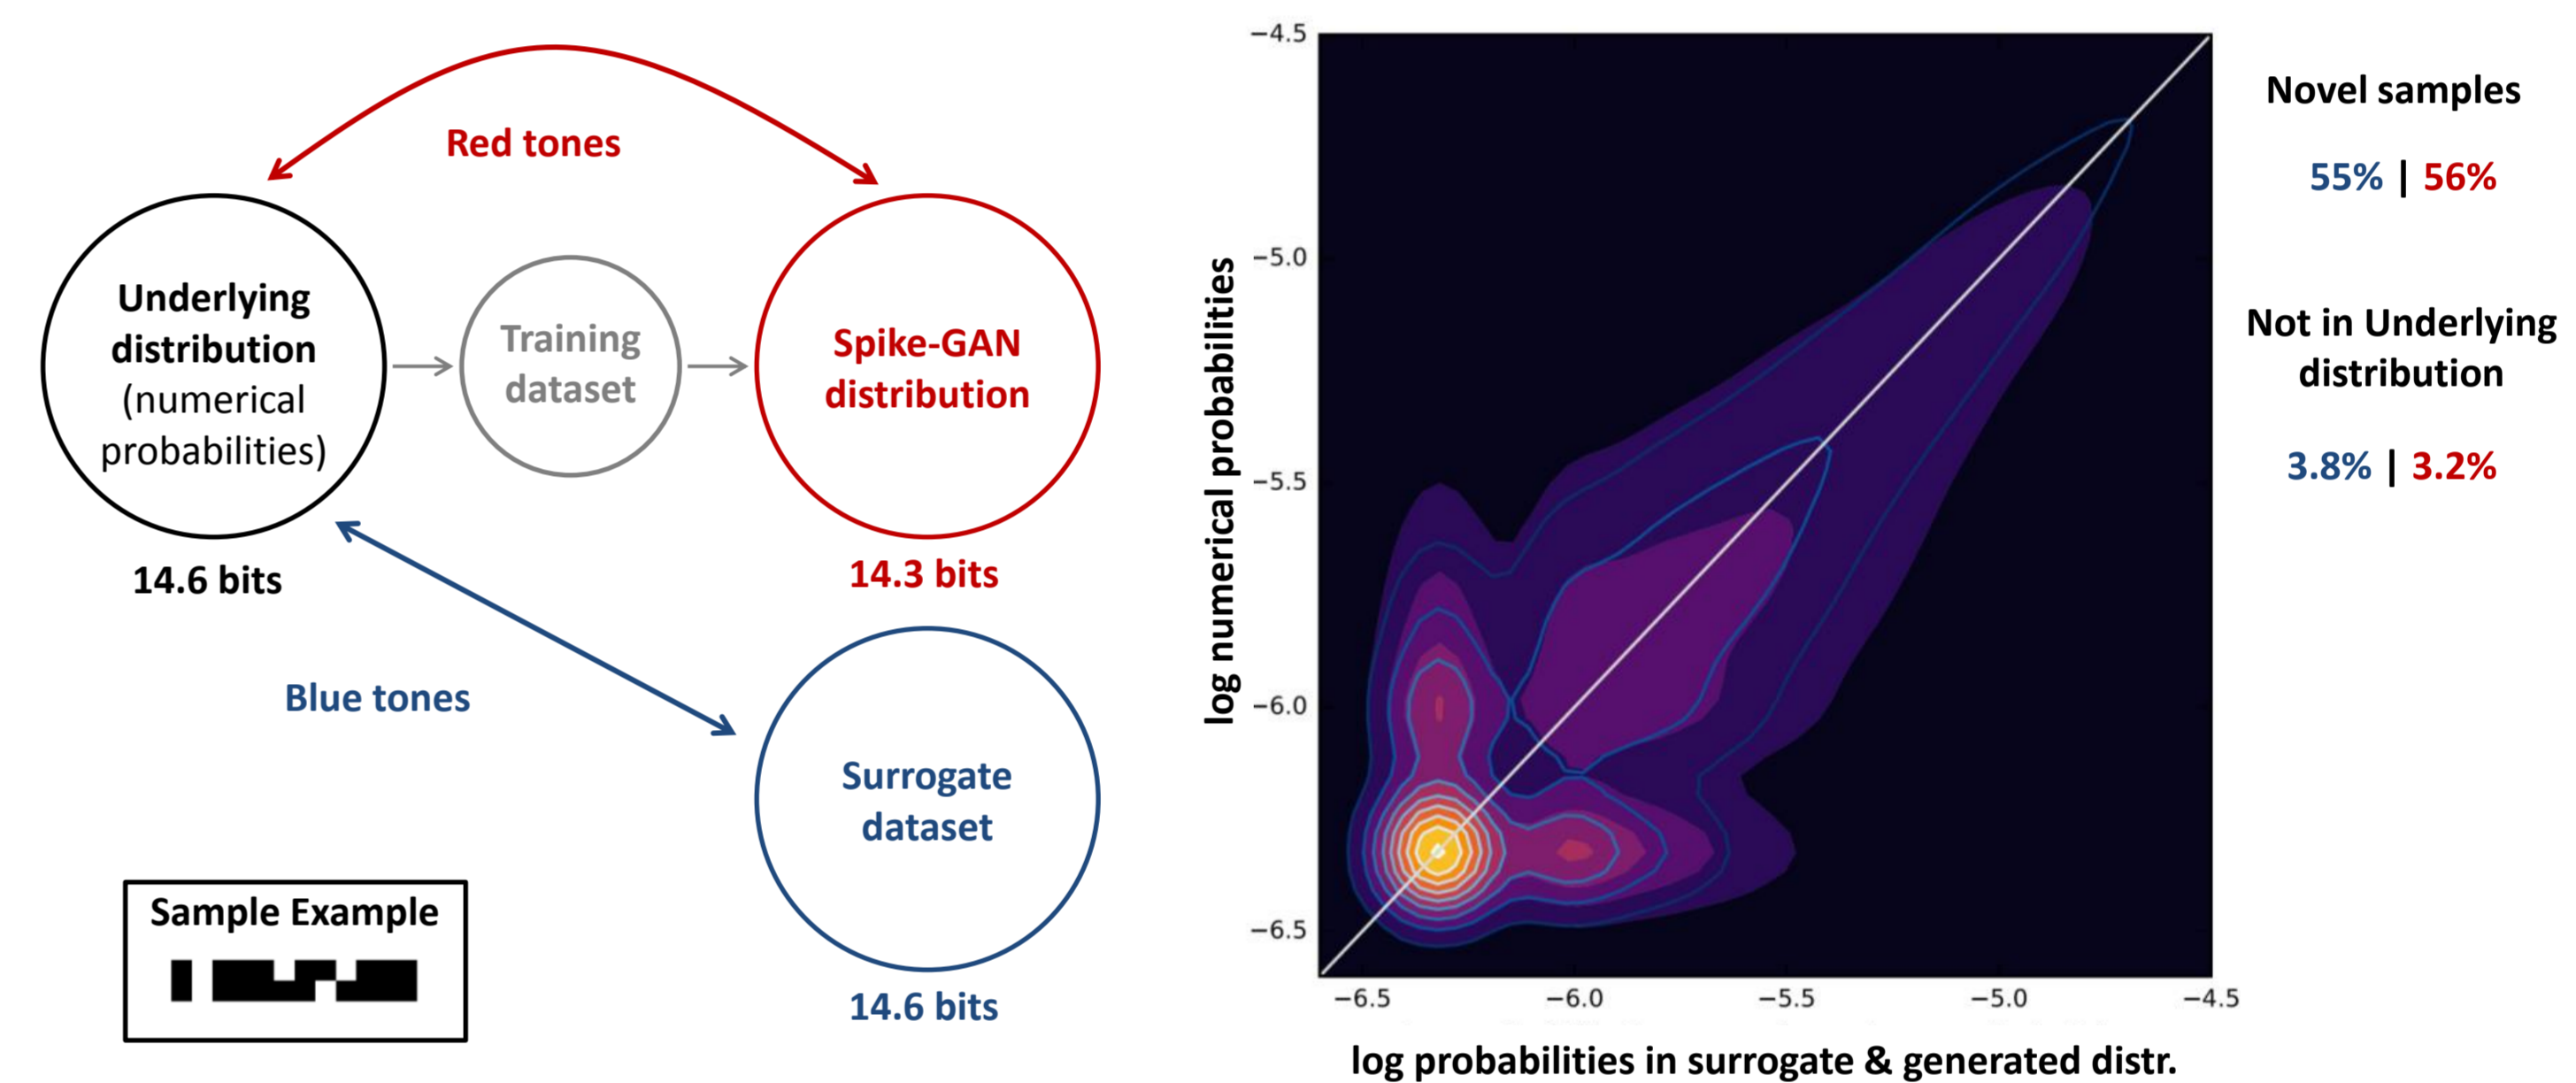
Fitting neural data from salamander retina

We tested Spike-GAN on recordings coming from **the retinal ganglion cells of the salamander retina** (Marre et al.). Spike-GAN performs as well as state-of-the-art approaches based on the **maximum entropy** (Tkacik et al.) and the **dichotomized Gaussian** (Lyamzin et al.) frameworks.

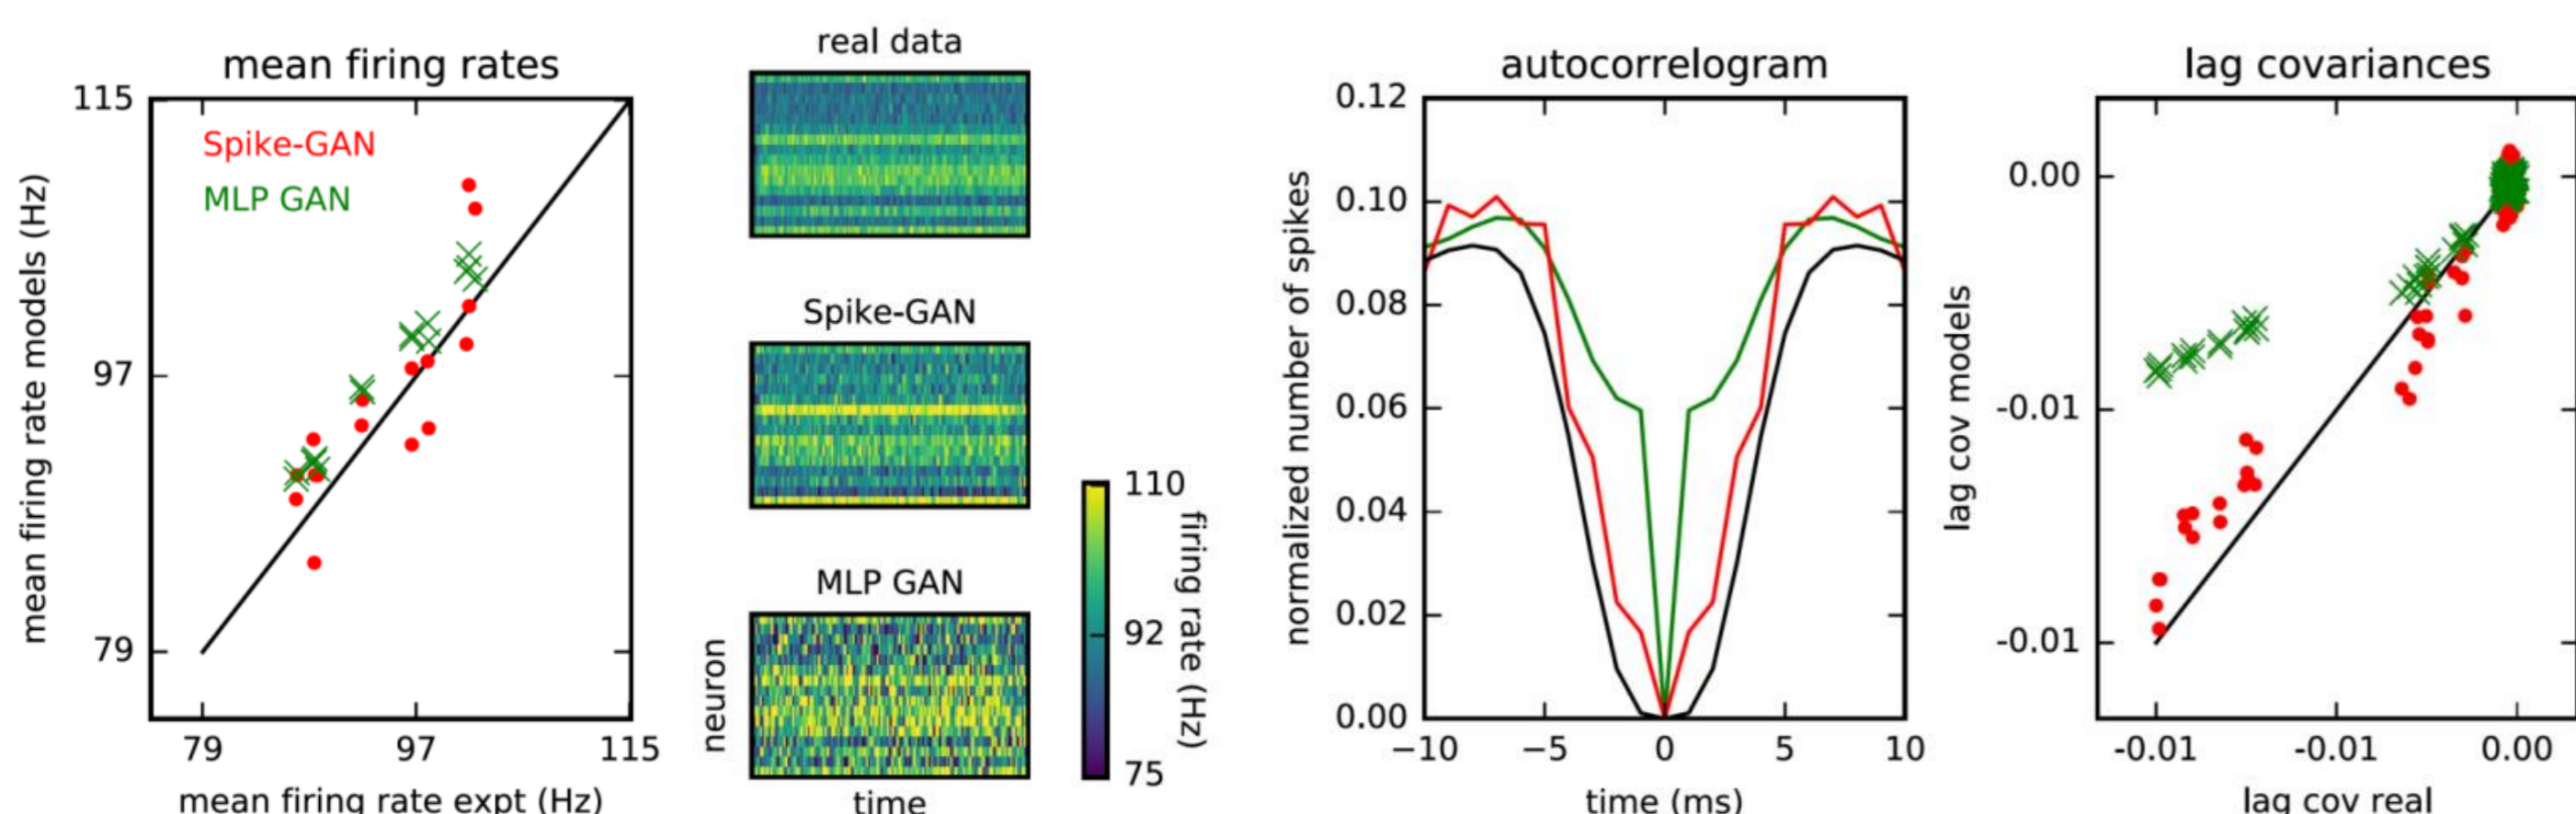


Fitting the whole probability distribution

We simulated activity patterns for a small 'population' of 2 neurons during 12 ms and evaluated how well Spike-GAN fits the whole probability density function from which the patterns are drawn.



Comparison with a multi-layer perceptron



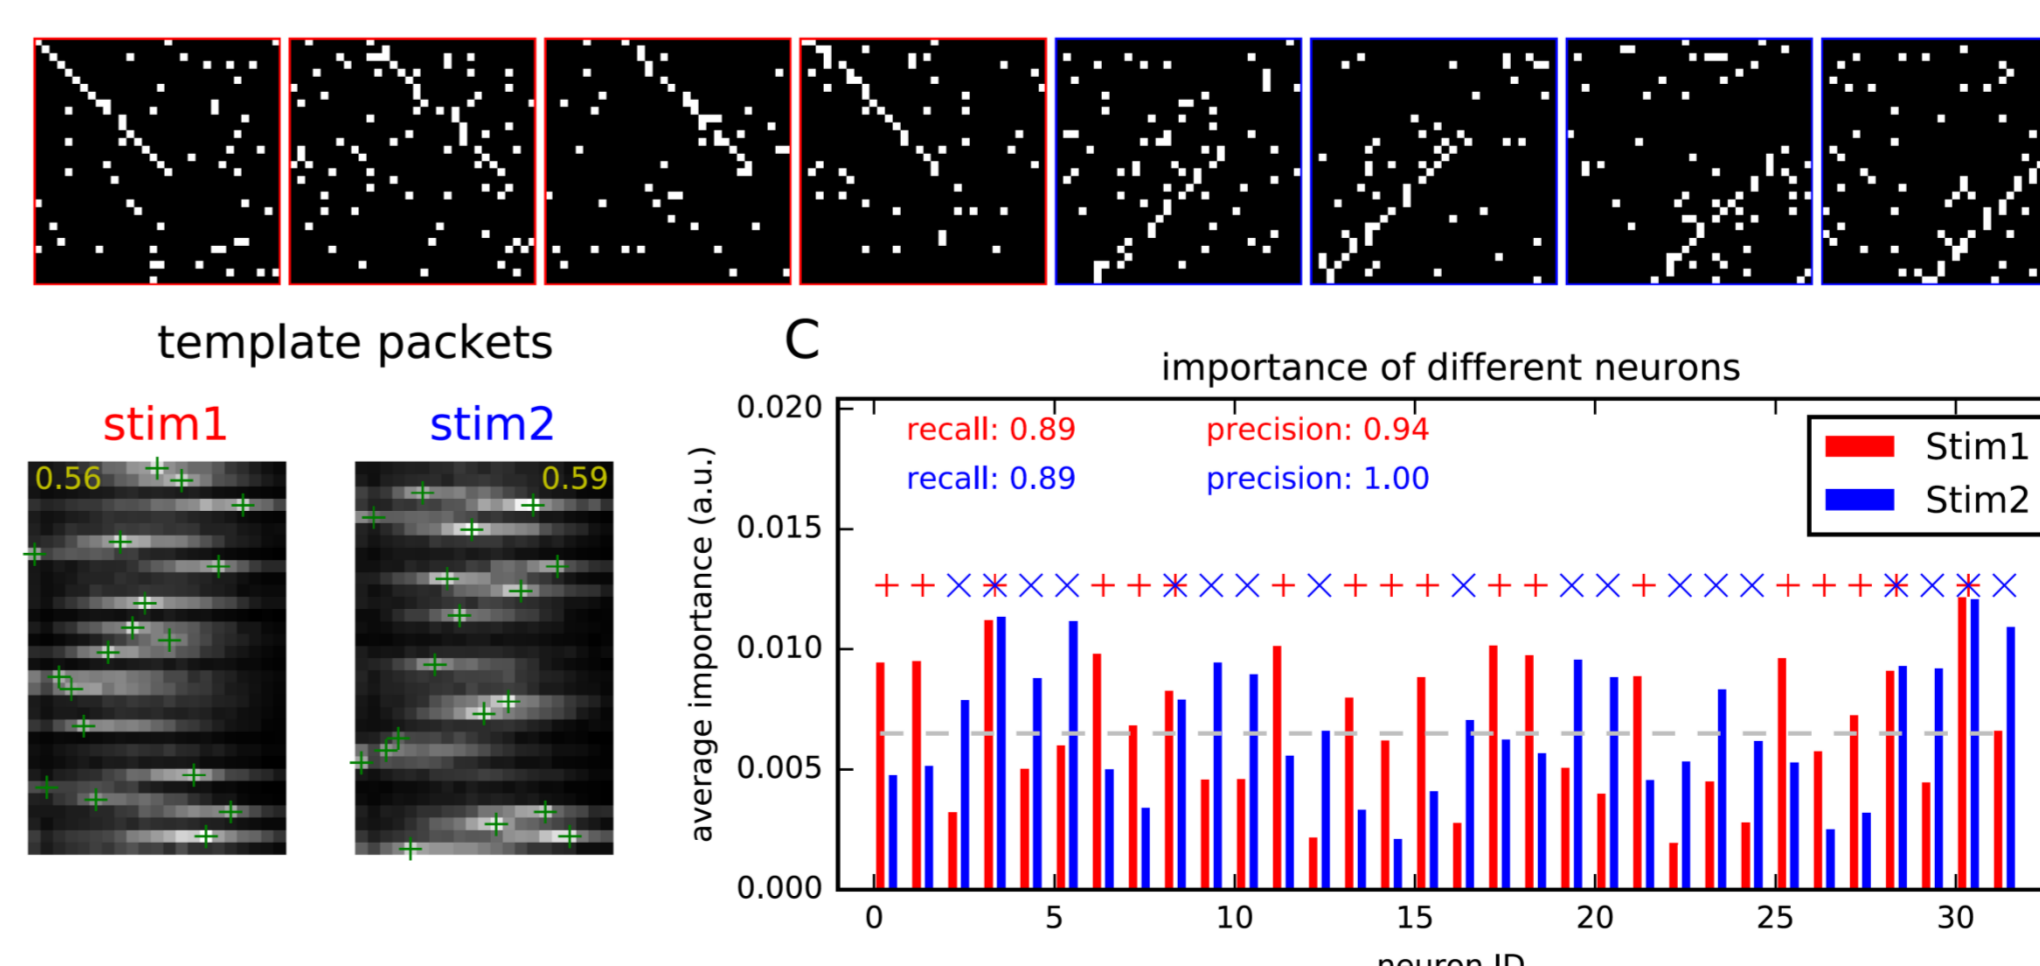
Code

<https://github.com/manuelmolano/Spike-GAN>

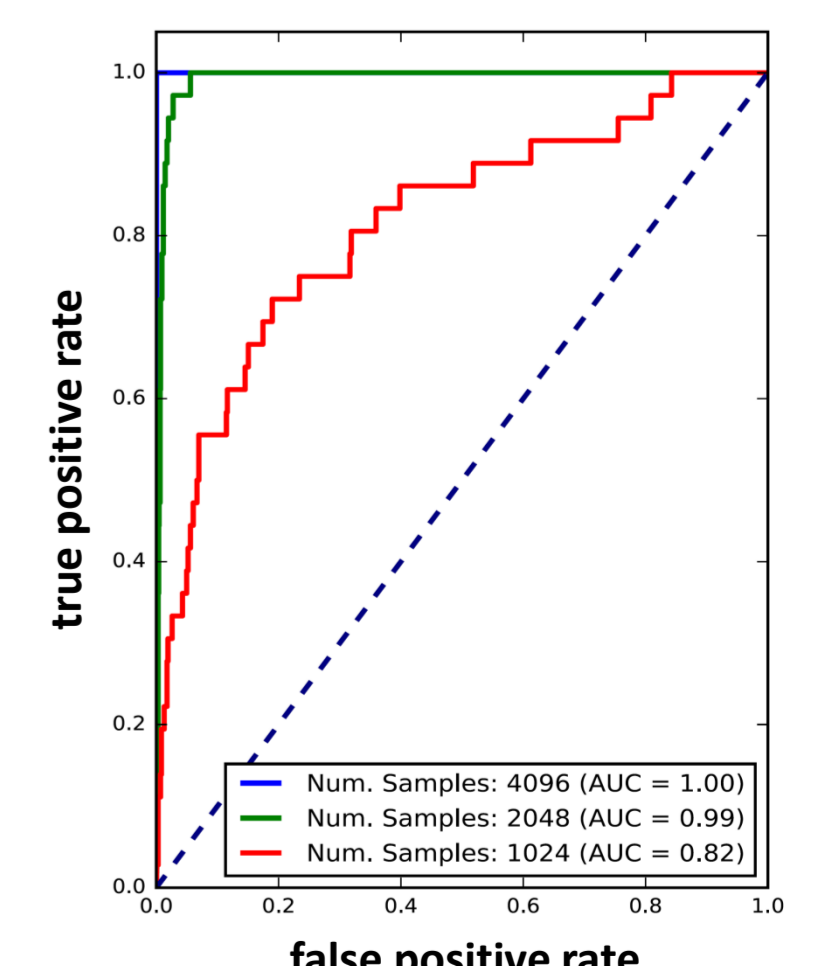
References

1. Arjovsky et al. arXiv 2017.
2. Goodfellow et al. NIPS 2014.
3. Gulrajani et al. NIPS 2017.
4. Odena et al. Distill 2016.
5. Marre et al. IST Austria 2014.
6. Tkacik et al. Plos Comp. Biol. 2014.
7. Lyamzin et al. Front. Comp. Neurosci. 2010.

Less reliable packets



Different number of samples



Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 699829.