Generating (Medical) Time Series with GANs

Stephanie Hyland

CLS Retreat, Stoos (CH)

4th October 2017

This work was done with Cristóbal Esteban and Gunnar Rätsch

one-slide talk summary

- Use **generative adversarial networks** with **recurrent neural networks** to generate time-series data
    - Look at synthetic data, MNIST, ICU time series
- **Evaluate** generated data:
    - To show it's **realistic**: with maximum mean discrepancy
    - To show it's **useful**: with (seemingly) novel transfer-learning approach
    - To show it's **not** the training data: multiple methods

why generate data?

(our motivation)

1. **Data sharing** - medical data requires protection, but this holds back machine learning research due to:
   a. Lack of reproducibility
   b. Lack of shared tasks on datasets
2. **Data augmentation** - difficulty of inter-hospital sharing means medical datasets can be small/limited
3. **Simulation** - generating realistic data from specific types of patients enables training/education of medical professionals

what is a GAN?

(generative adversarial network)

discriminator network

objective: discriminate if a sample is **real** or **fake**
(binary classification)

generator network

objective: generate fake samples to **trick** discriminator
(maximise probability of discriminator outputting positive label on fake samples)

what is a GAN?

(generative adversarial network)

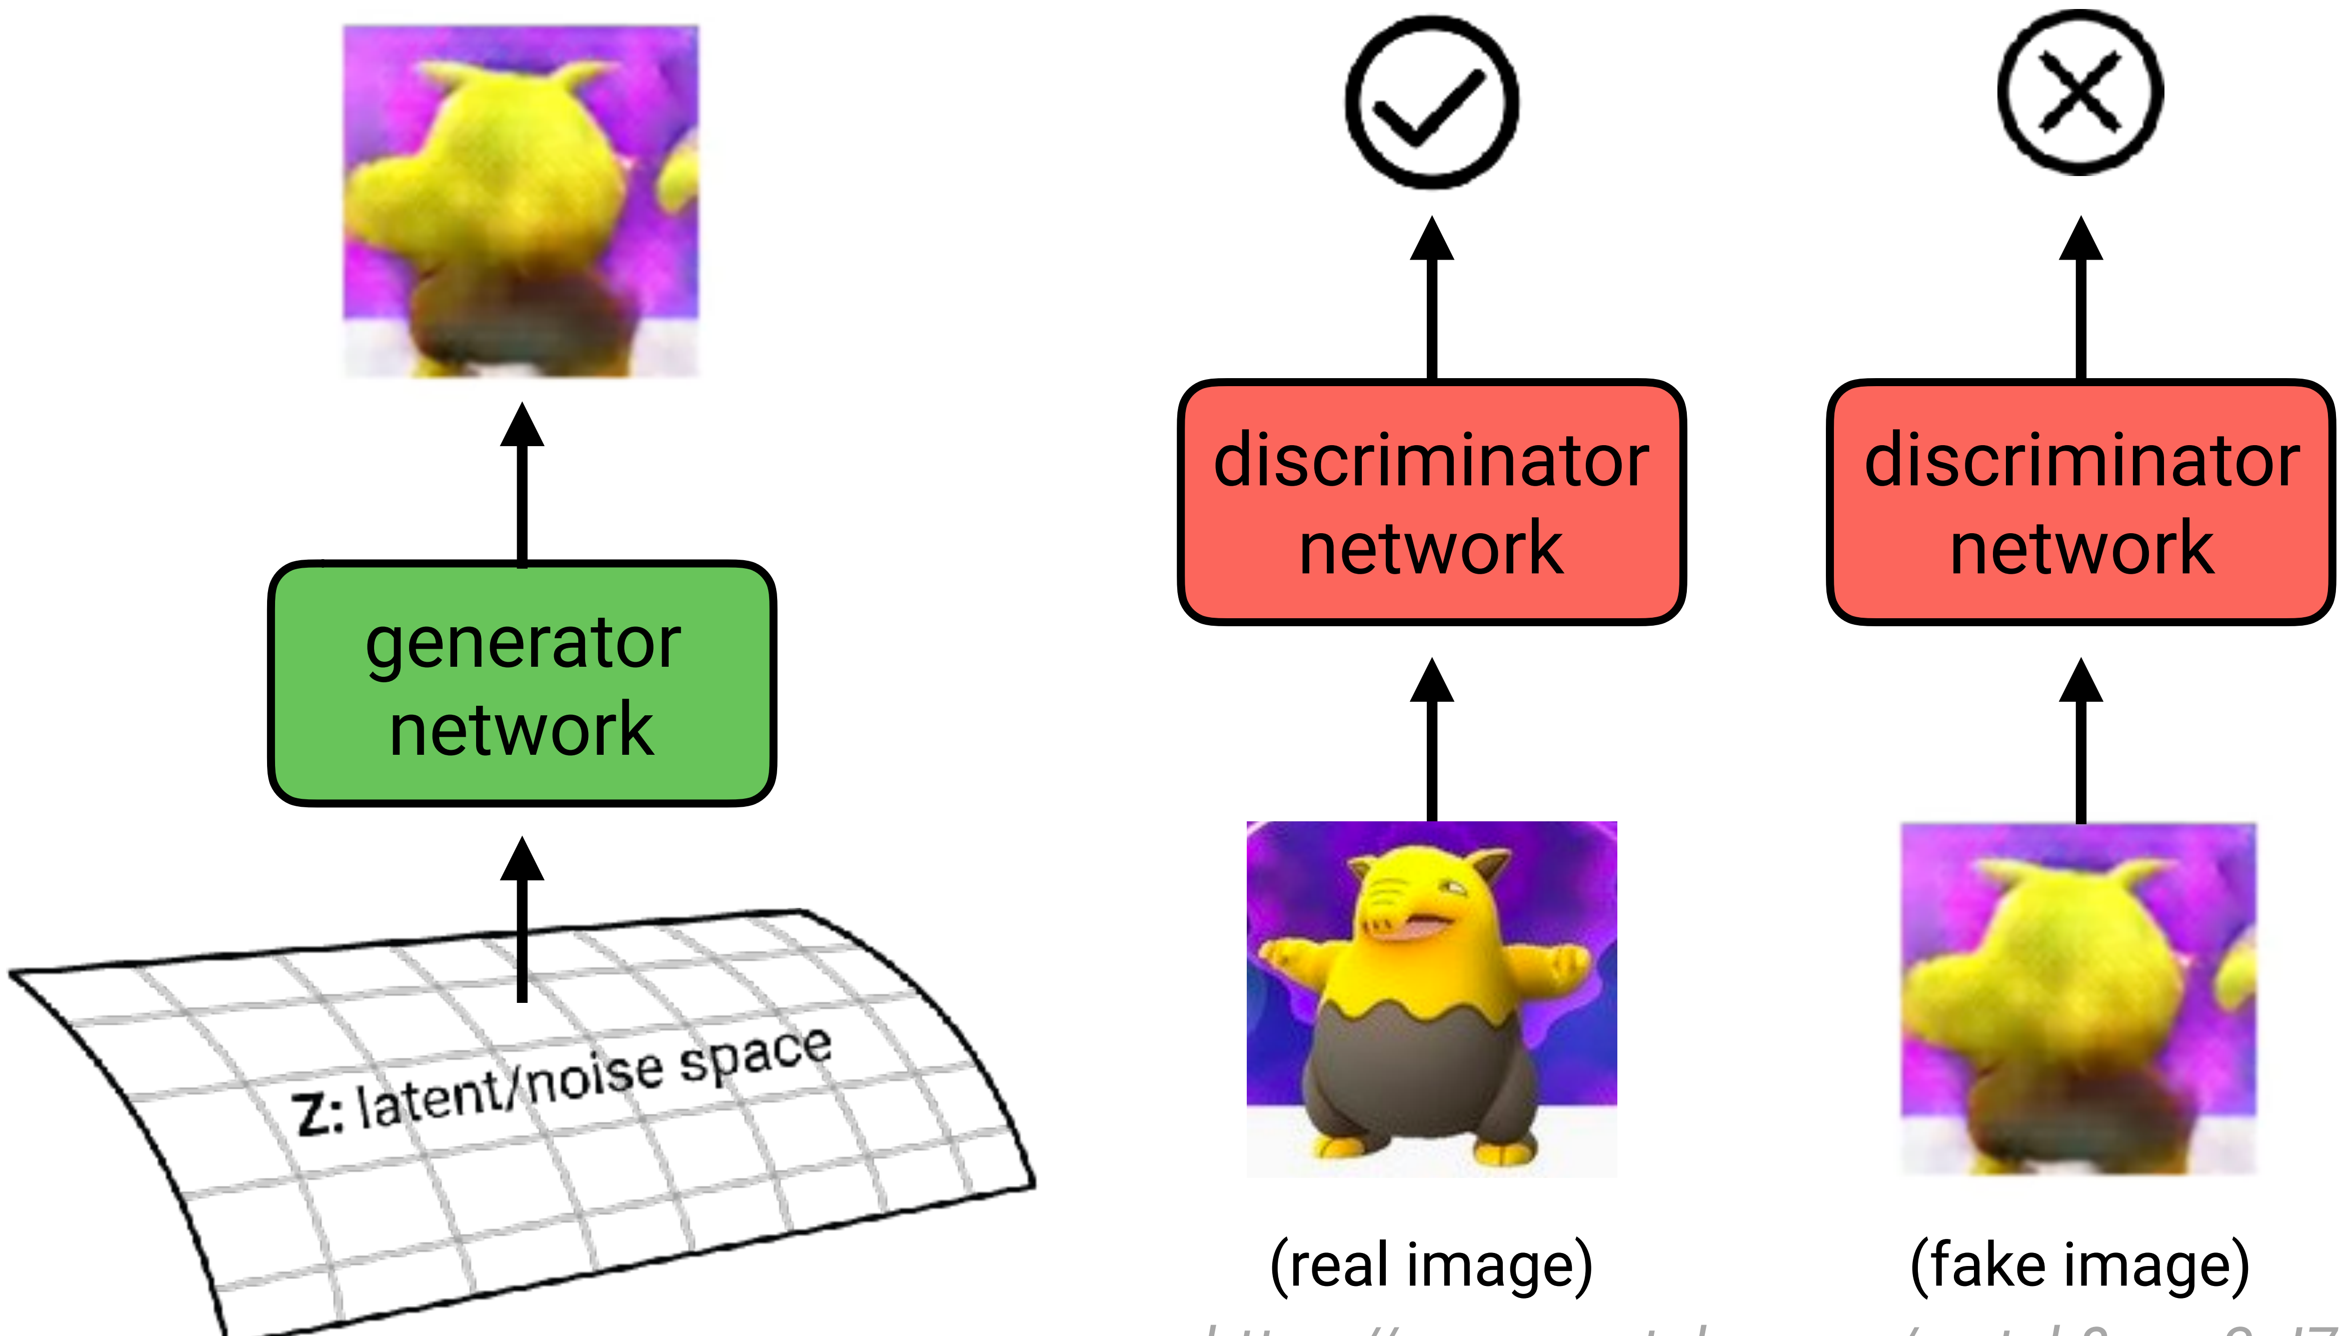

https://www.youtube.com/watch?v=rs3aI7bACGc
DCGAN with Pokemon Go - Yota Ishida

example synthetic images

trained on bedrooms

https://github.com/cameronfabbri/LSGANs-Tensorflow
https://github.com/martinarjovsky/WassersteinGAN

GANs for sequences

- Most GAN work is on *images*, where generator and discriminator are convolutional neural networks
- What about generating *sequences*?
- People have tried it for text (fixed-length "sentence"):

Marks live up in the club comes the handed up moved to a brief d

Language Generation with Recurrent Generative Adversarial Networks without Pre-training, Ofir Press, Amir Bar, Ben Bogin, Jonathan Berant, Lior Wolf, arXiv 17

Busino game camperate spent odea

Improved Training of Wasserstein GANs, Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, Aaron Courville, arXiv 17

Recurrent GAN (RGAN)

- Idea: use **recurrent neural networks** (e.g. LSTM) in both discriminator and generator:

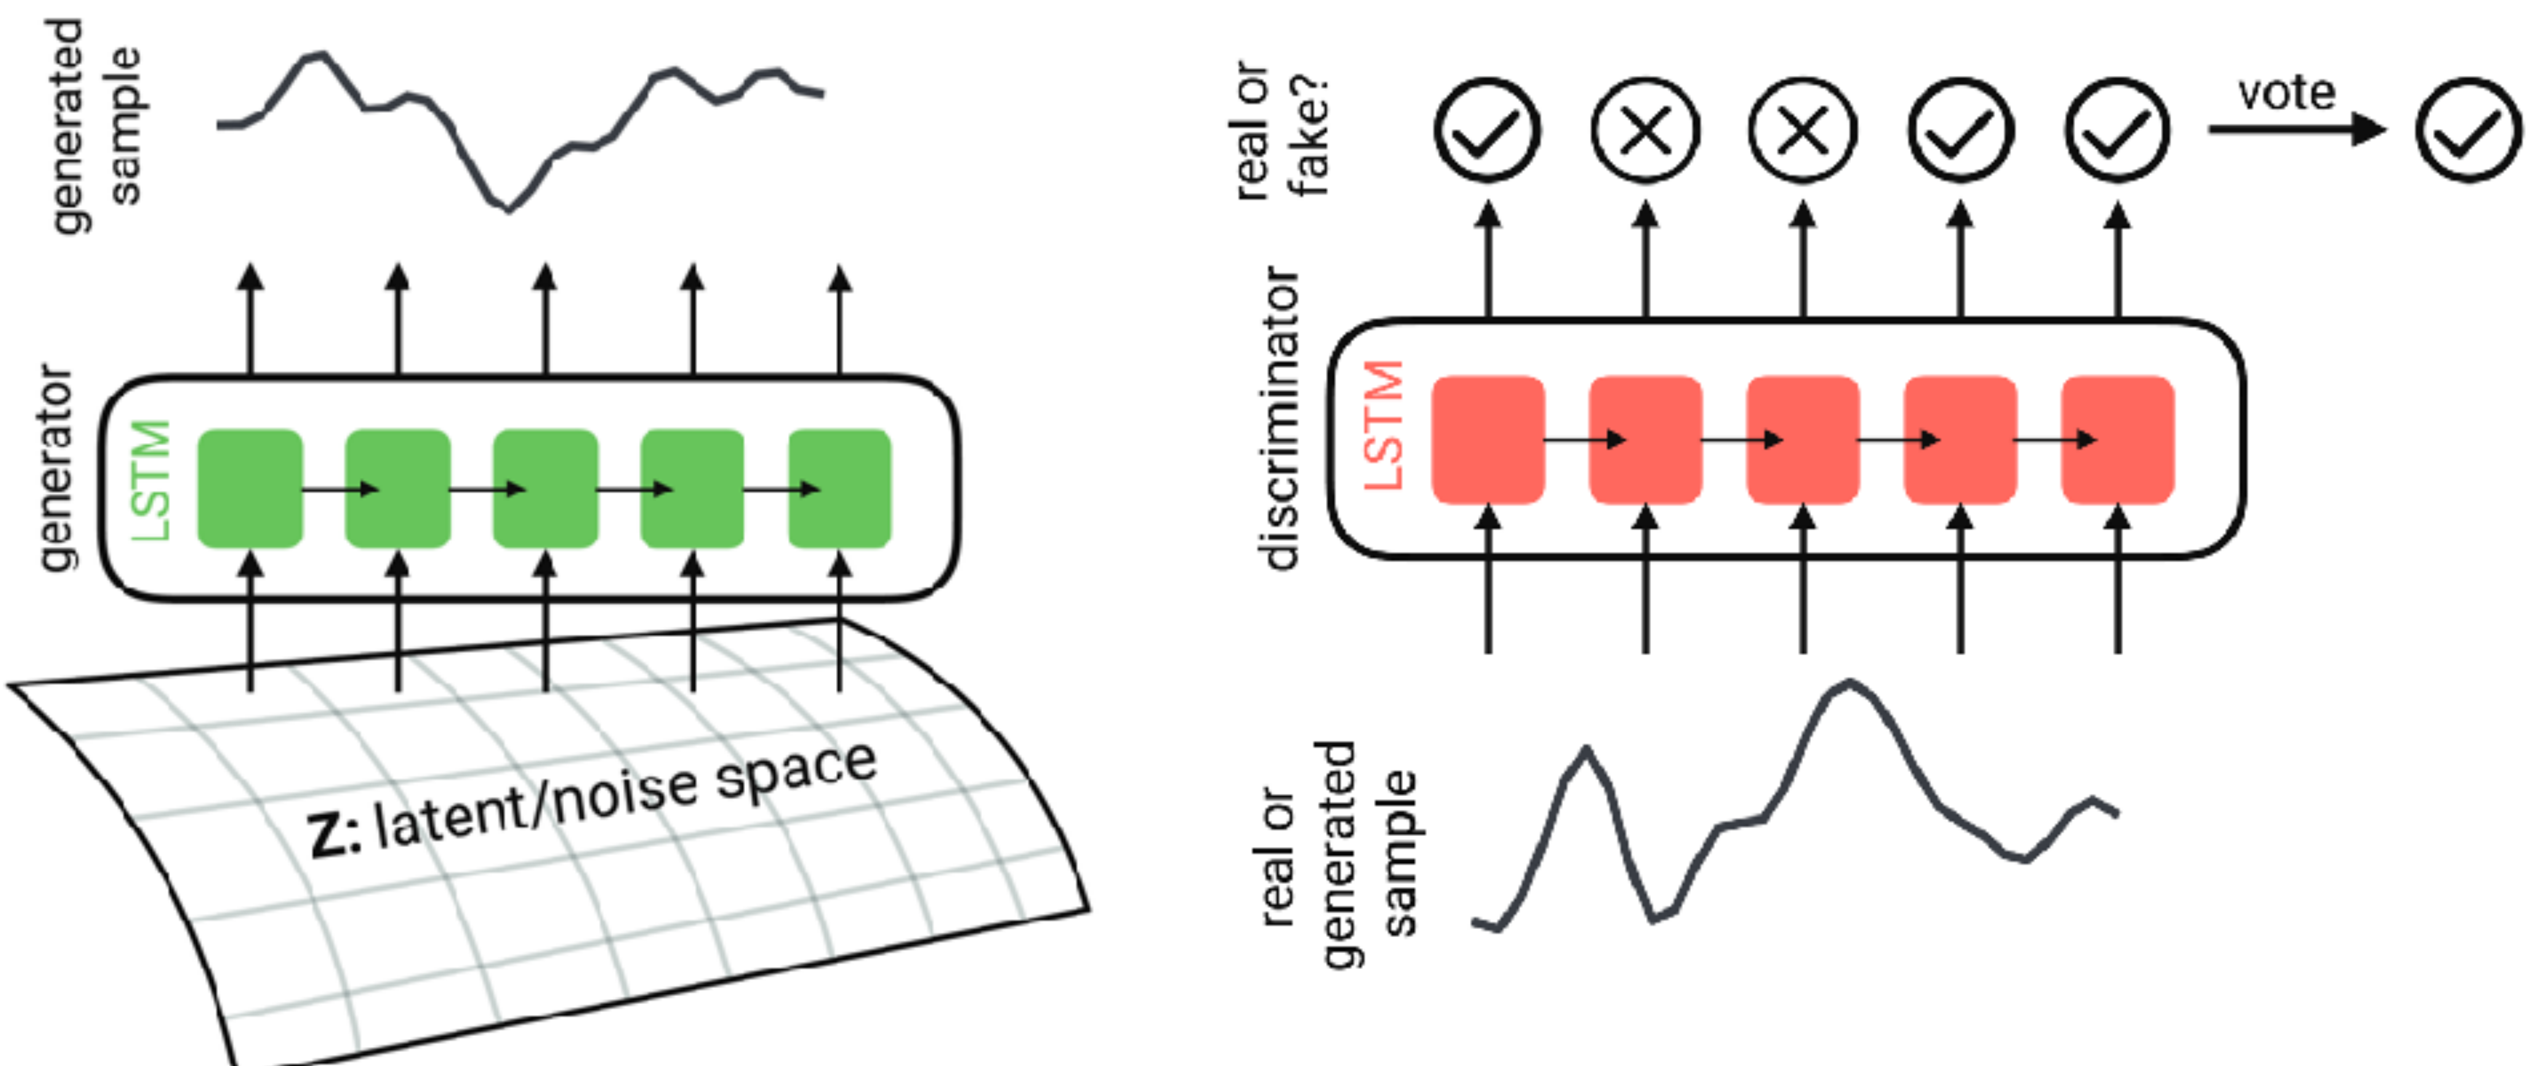

https://github.com/ratschlab/RGAN/blob/master/figures/sine_animation.gif

RGAN - sine waves

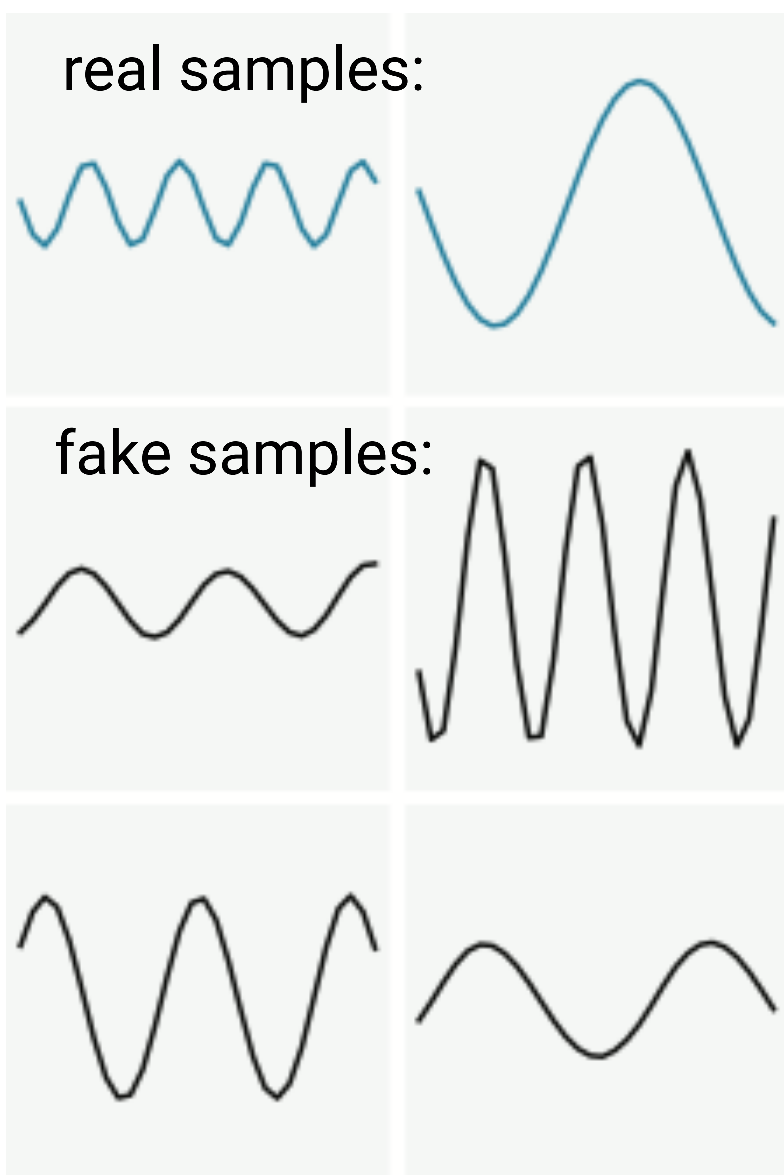

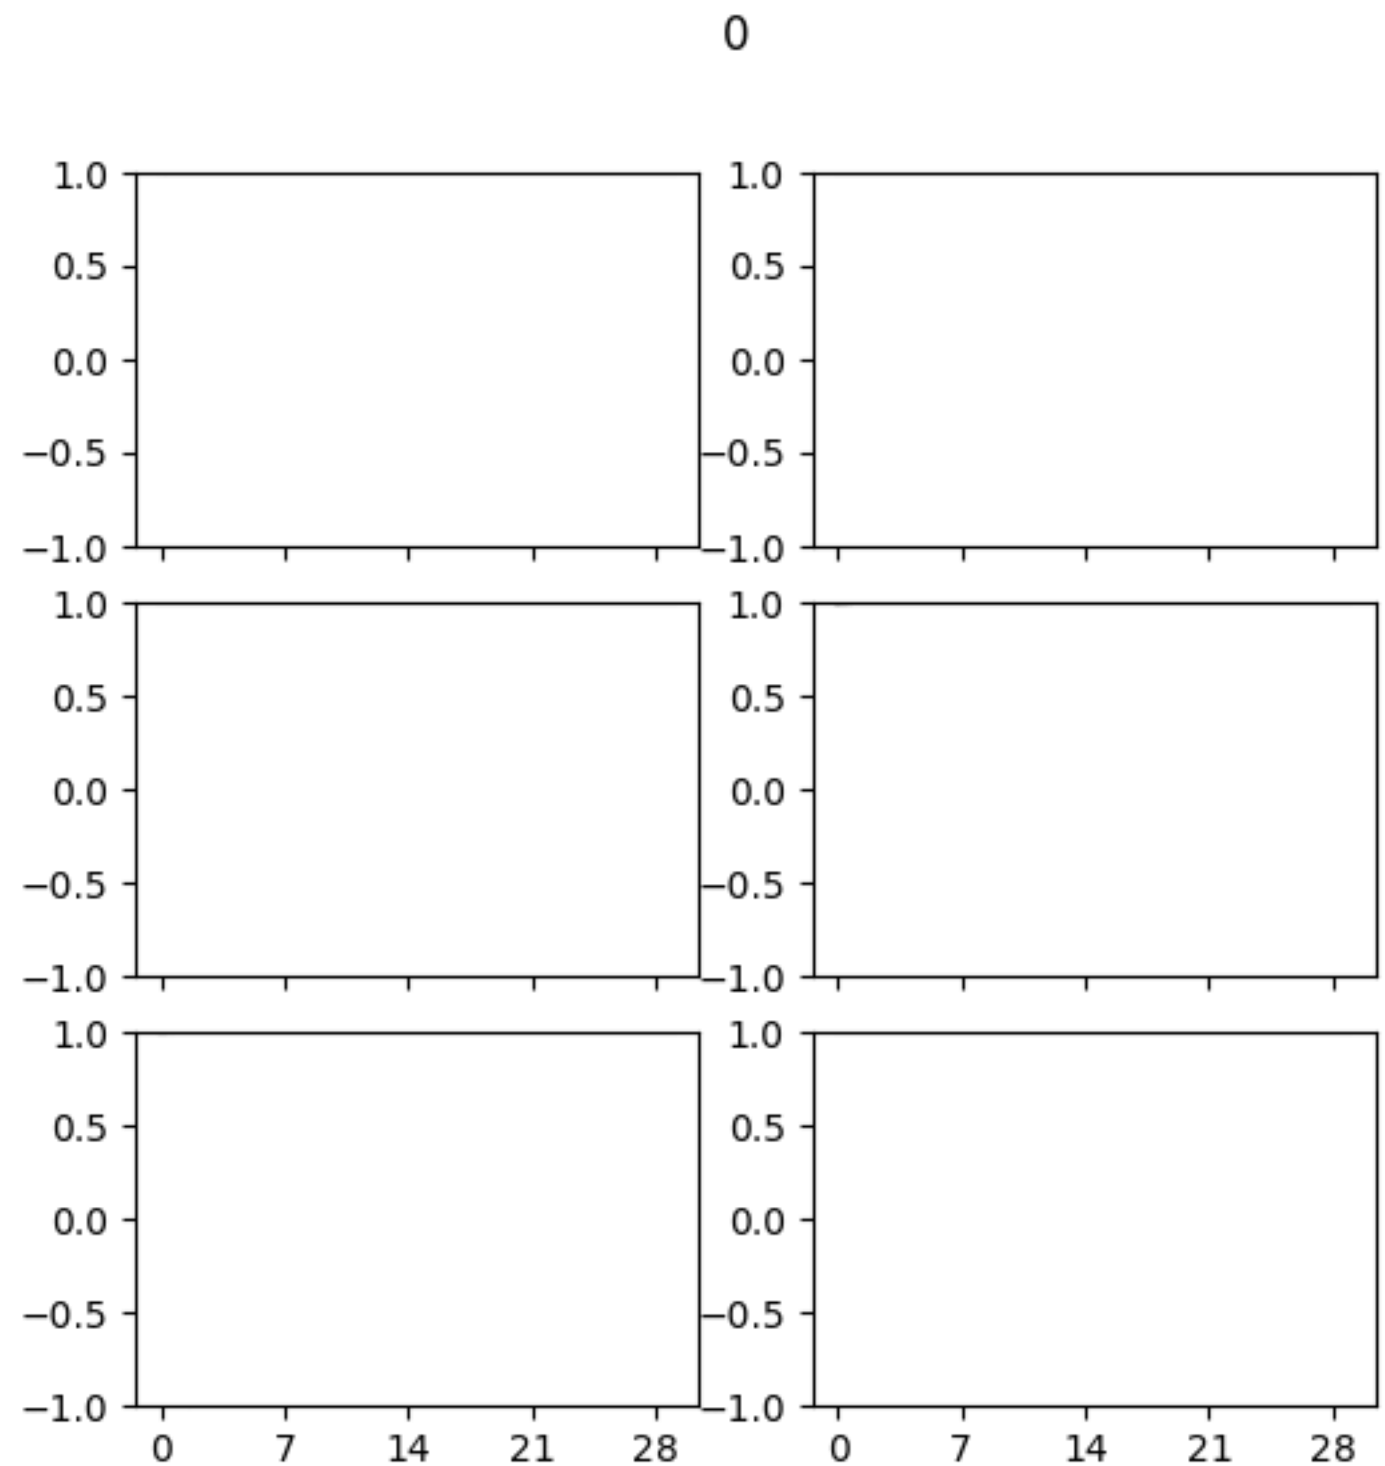

generator output from same 6 latent points throughout training

https://github.com/ratschlab/RGAN/blob/master/figures/rbf_animation.gif

RGAN - smooth functions

(samples from gaussian process)

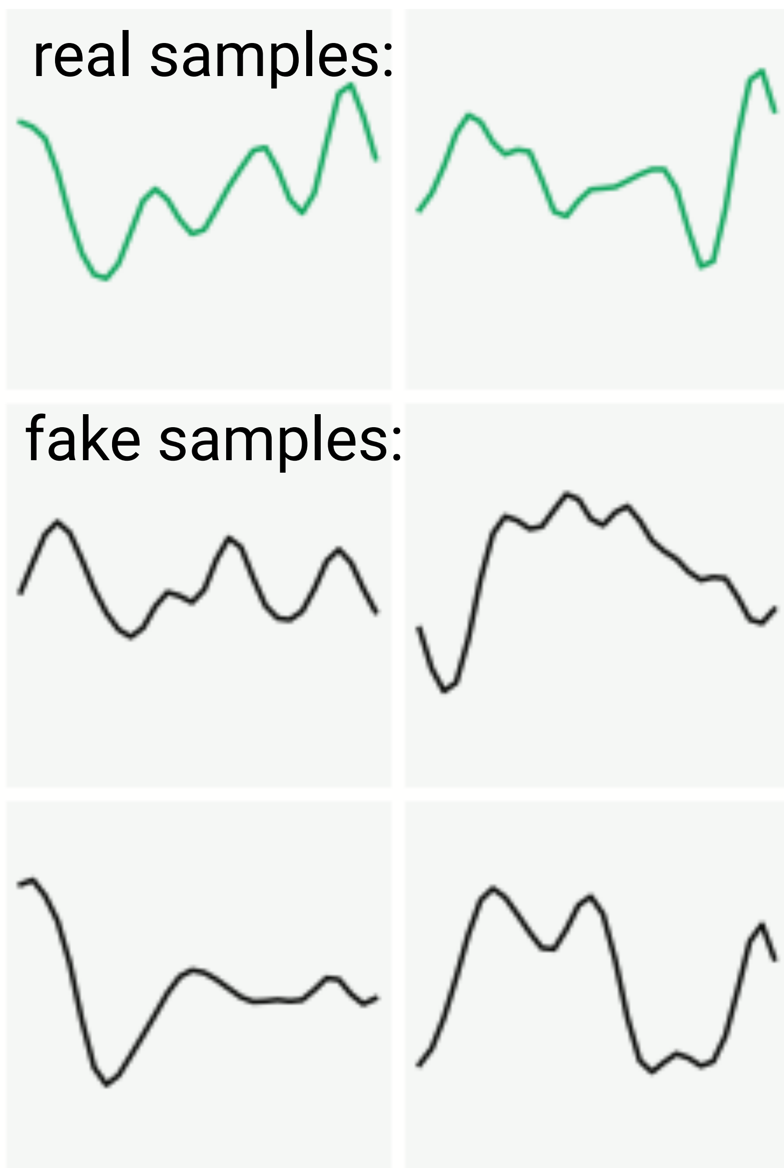

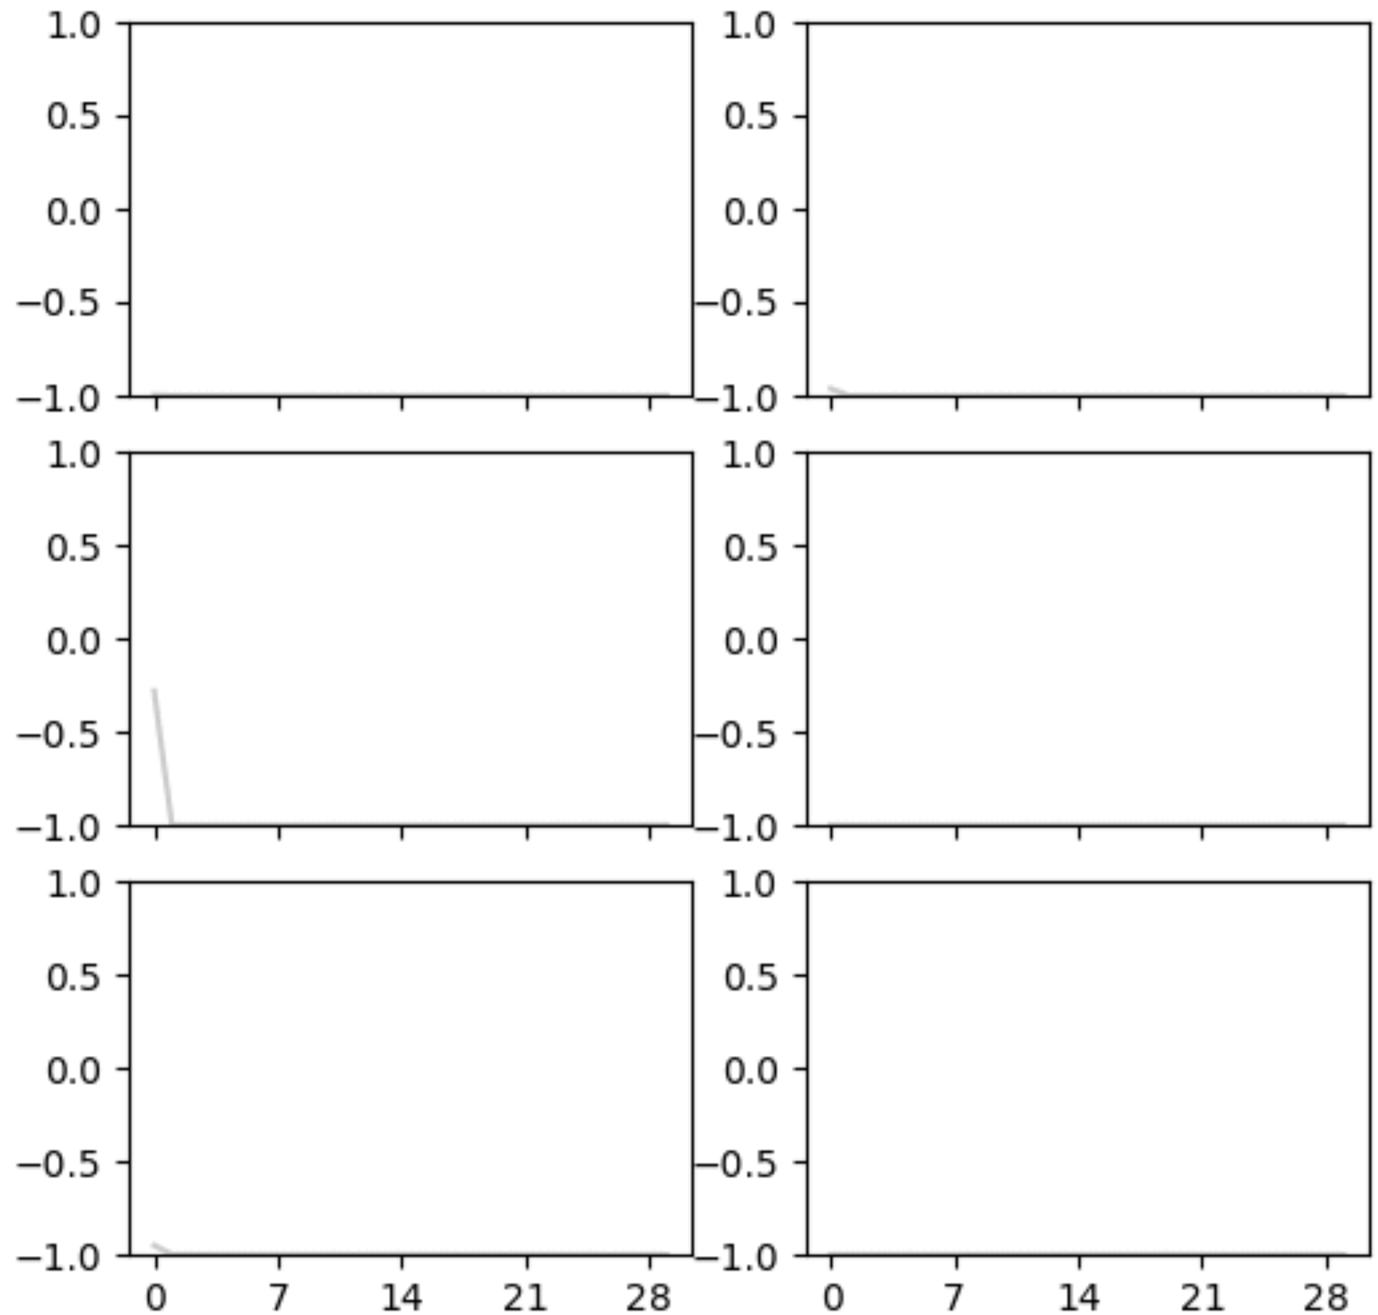

generator output from same 6 latent points throughout training

https://github.com/ratschlab/RGAN/blob/master/figures/mnistfull_a_animation.gif

RGAN - MNIST digits

MNIST as 14x14 sequence

real samples:

fake samples:

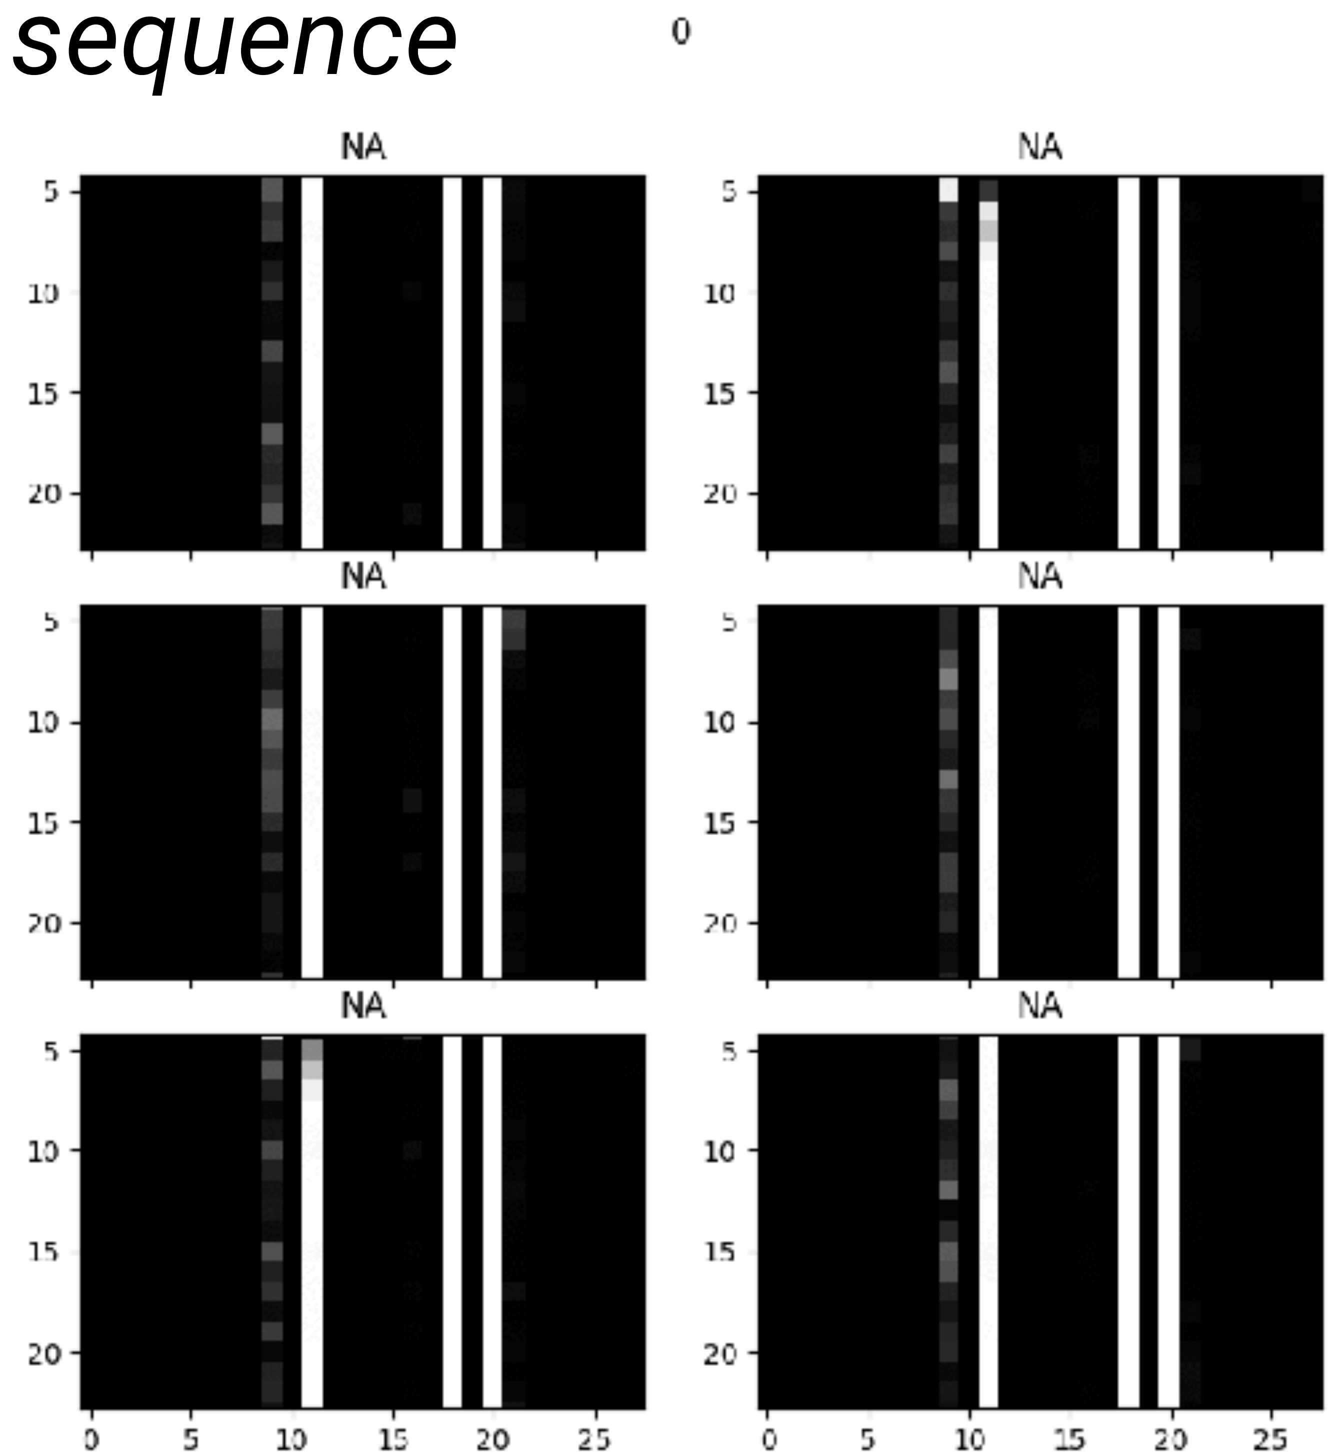

generator output from same 6 latent points throughout training

evaluating GAN samples
(beyond visual inspection…)

- Visual inspection as evaluation: (used in many papers)
  - ⊗ Doesn’t scale
  - ⊗ Uses viewer’s ‘internal discriminator’ (subjective)
- Improvement: use **maximum mean discrepancy** - ask if two sets of samples came from the same *distribution*
  - ⊗ Requires an appropriate kernel (K) on sample space

$$\widehat{\mathrm{MMD}}_u^2 = \frac{1}{n(n-1)}\sum_{i=1}^{n}\sum_{j\neq i}^{n} K(x_i, x_j) - \frac{2}{mn}\sum_{i=1}^{n}\sum_{j=1}^{m} K(x_i, y_j) + \frac{1}{m(m-1)}\sum_{i=1}^{m}\sum_{j\neq i}^{m} K(y_i, y_j)$$

evaluating GAN samples

(with maximum mean discrepancy…)

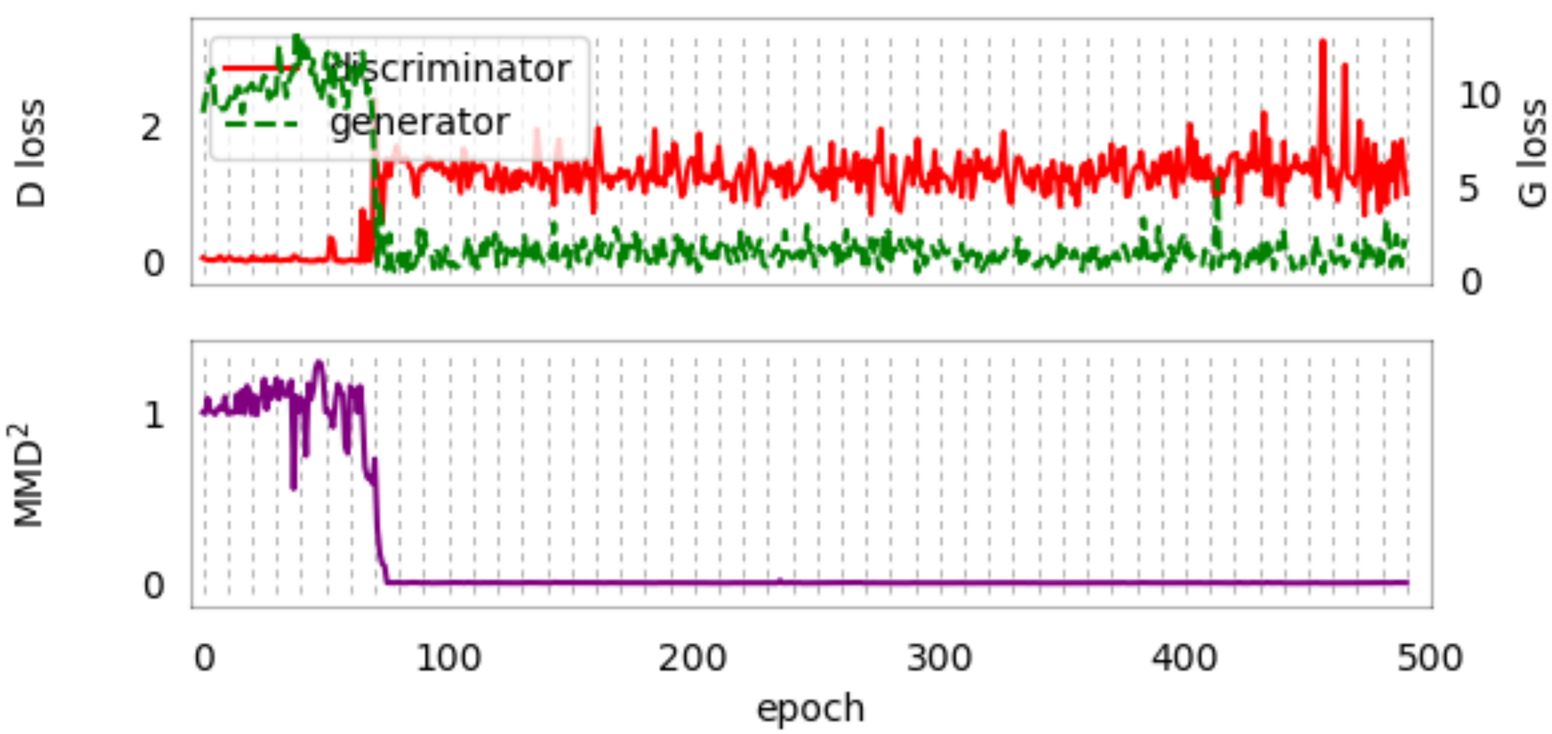

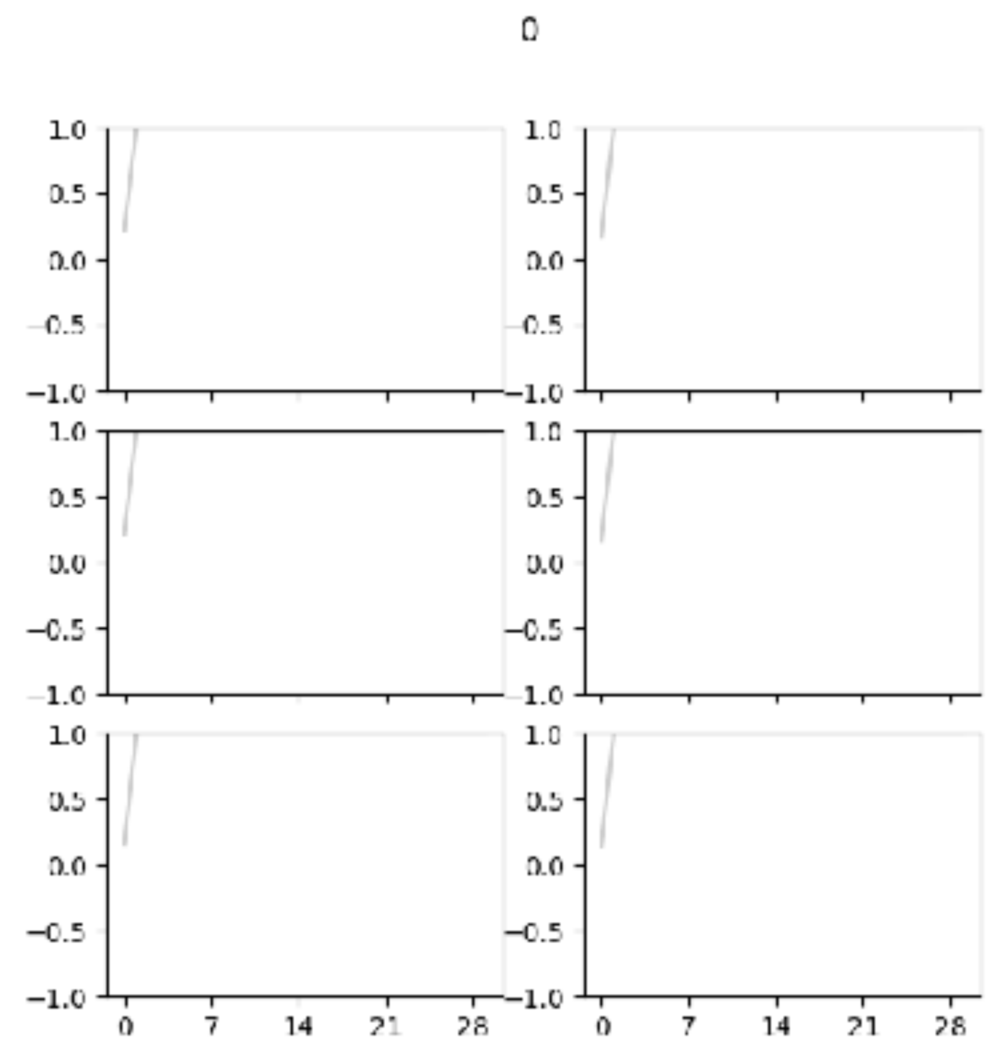

(this example looks worse than normal because it was trained using differentially private stochastic gradient descent)

evaluating GAN samples

*(with a **task**)*

- Idea: (**TSTR**)
    a. train model with **synthetic data**
    b. test that model on **real data**

- Captures the utility of the synthetic data for its intended purpose - training models that can generalise

- Requires **labelled** synthetic dataset:
    - Extend RGAN to be *conditional* (RCGAN)
    - Takes latent point and label to generate sample

a medically relevant task

- **Data:** “eICU collaborative research database”, from Phillips eICU program (open access)
  - ~**140,000** patients from **459** critical care units across the USA (between 2014-2015)
  - Measurements every 5 minutes of various vital signs
- **Task:**
  - Predict if values will become ‘extreme’ (high/low)
  - Focus on first 4 hours of patient’s stay
    - variables: heart rate, respiratory rate, mean arterial pressure, oxygen saturation

TSTR results for eICU task

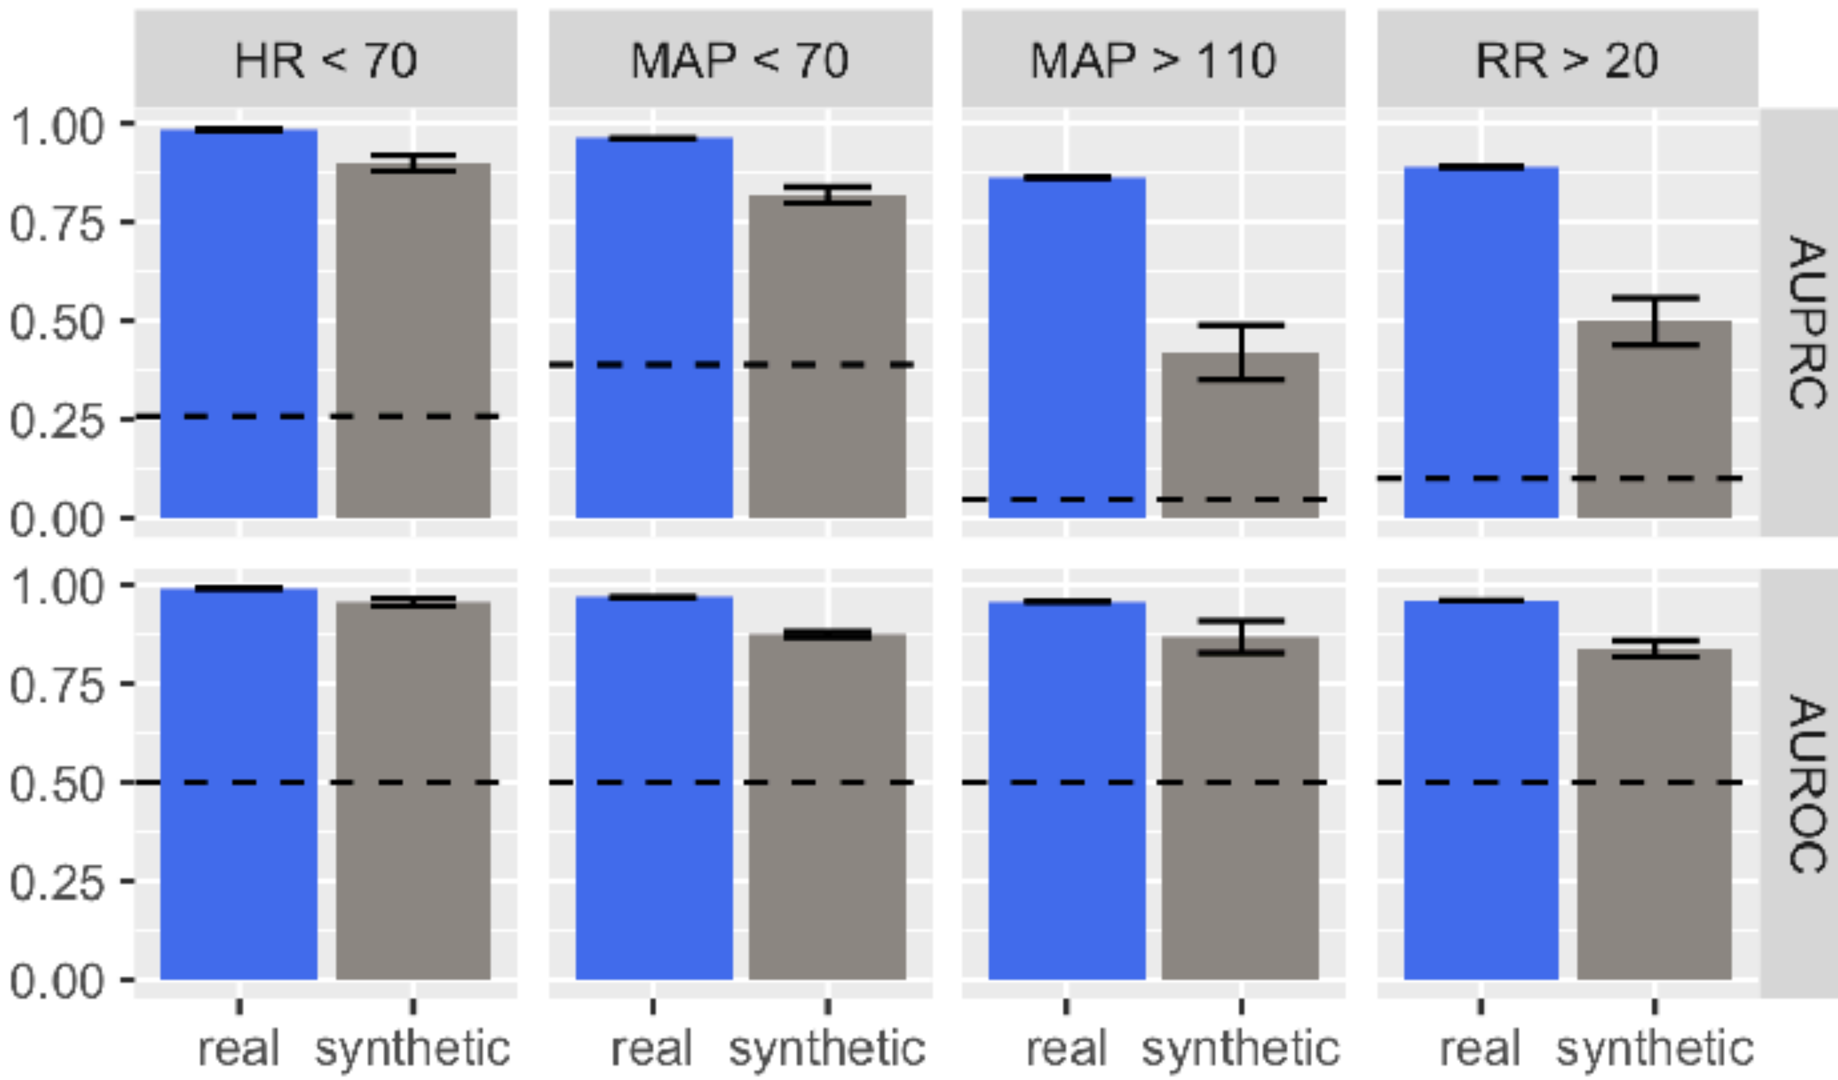

Area under precision-recall curve

Area under ROC curve

Conclusion: model trained on synthetic data can still generalise well to real test set

is it memorising the data?

- Trivially successful generator would simply output training examples
  - Known problem for GANs - ‘mode collapse’
  - Constitutes privacy breach of original training data
- Tested to see if model is memorising:
  1. Do **reconstruction errors** differ between train and test?
  2. **Maximum mean discrepancy:** is MMD(synth, train) < MMD(synth, test)?
  3. **Interpolation**: smoothly vary between latent points corresponding to training examples, look at output

interpolation examples

1. Backwards-project A, B to get point(s) in latent space
2. Produce intermediate latent points
3. Run through generator to create samples

Training example A →

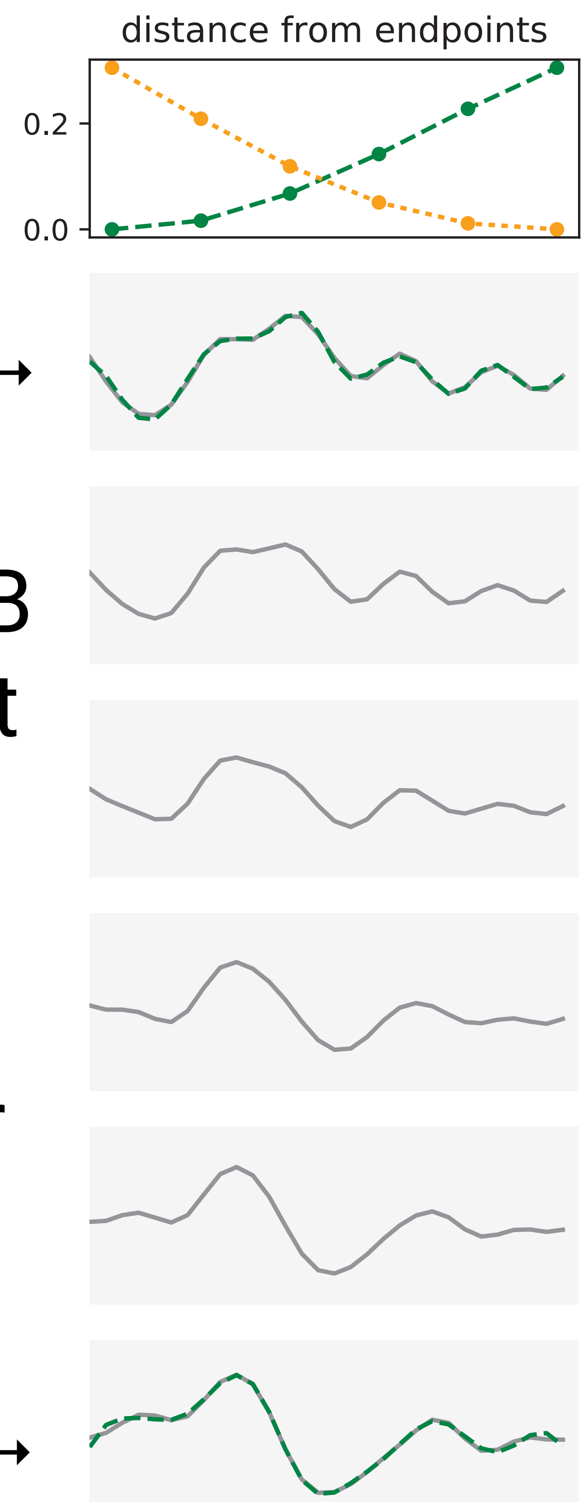

Training example B →

Mont Blanc →

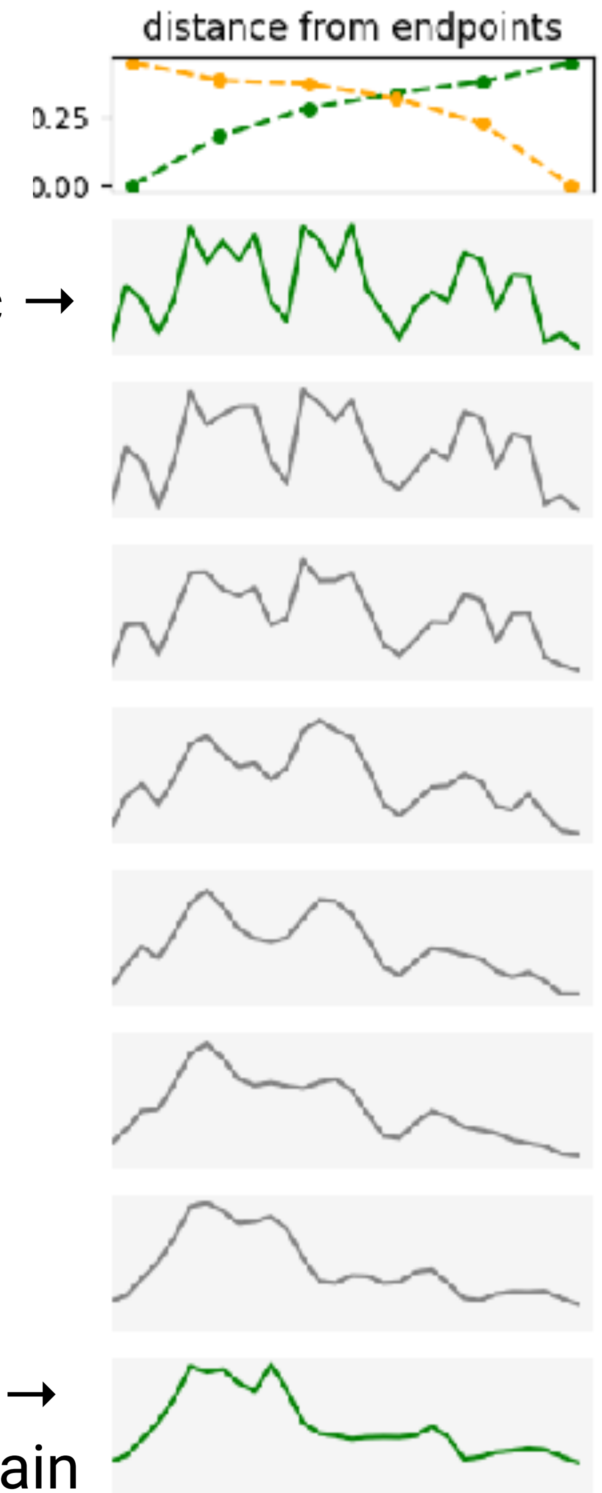

Carrauntoohill →
(highest mountain in Ireland)

in conclusion

- **Recurrent generative adversarial network** to generate synthetic, real-valued time-series data (e.g. ICU data)
- (Seemingly) novel evaluation via **TSTR method**
- Analysed synthetic data for evidence of **memorisation** (found none)
- Future work:
  - Include stronger privacy guarantees (e.g. differential privacy)
  - Comparisons with other methods (e.g. variational recurrent auto-encoders)

Thank you!

With thanks to:
Cristóbal Esteban
Gunnar Rätsch
Francesco Locatello
Xinrui Lyu
The rest of the BMI group

Real-valued (Medical) Time Series Generation with Recurrent Conditional GANs

Stephanie L. Hyland*[1,2], Cristóbal Esteban*[1], Gunnar Rätsch[1]
[1]Department of Computer Science, ETH Zurich, Switzerland
[2]Tri-Institutional Training Program in Computational Biology and Medicine, Weill Cornell Medical
{stephanie.hyland, cristobal.esteban, raetsch}@inf.ethz.ch
*Authors contributed equally.

https://github.com/ratschlab/RGAN

https://arxiv.org/abs/1706.02633

hyland@inf.ethz.ch

@_hylandSL