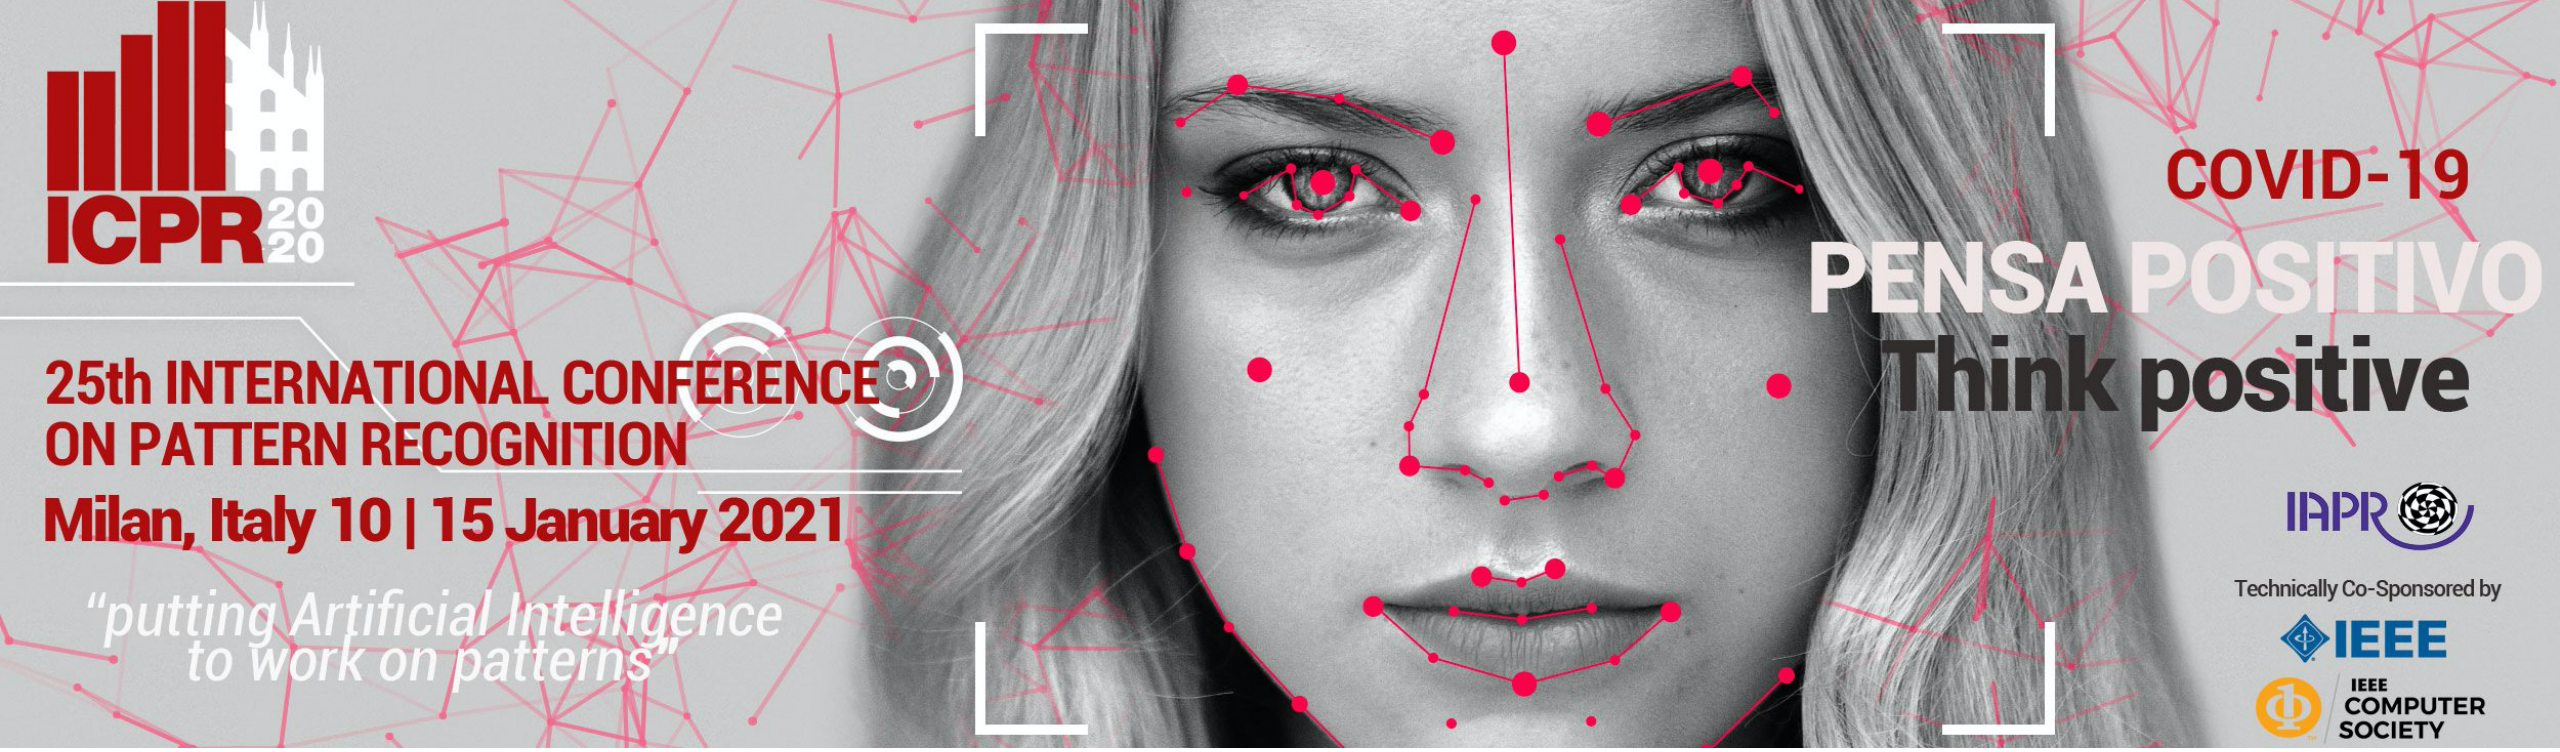

On the use of Benford's law to detect GAN-generated images

Nicolò Bonettini, Paolo Bestagini, Stefano Tubaro – Politecnico di Milano

Simone Milani – Università degli Studi di Padova

POLITECNICO MILANO 1863

UNIVERSITÀ DEGLI STUDI DI PADOVA

Detection of GAN-generated images

Problem

- Images generated by GANs can be very realistic

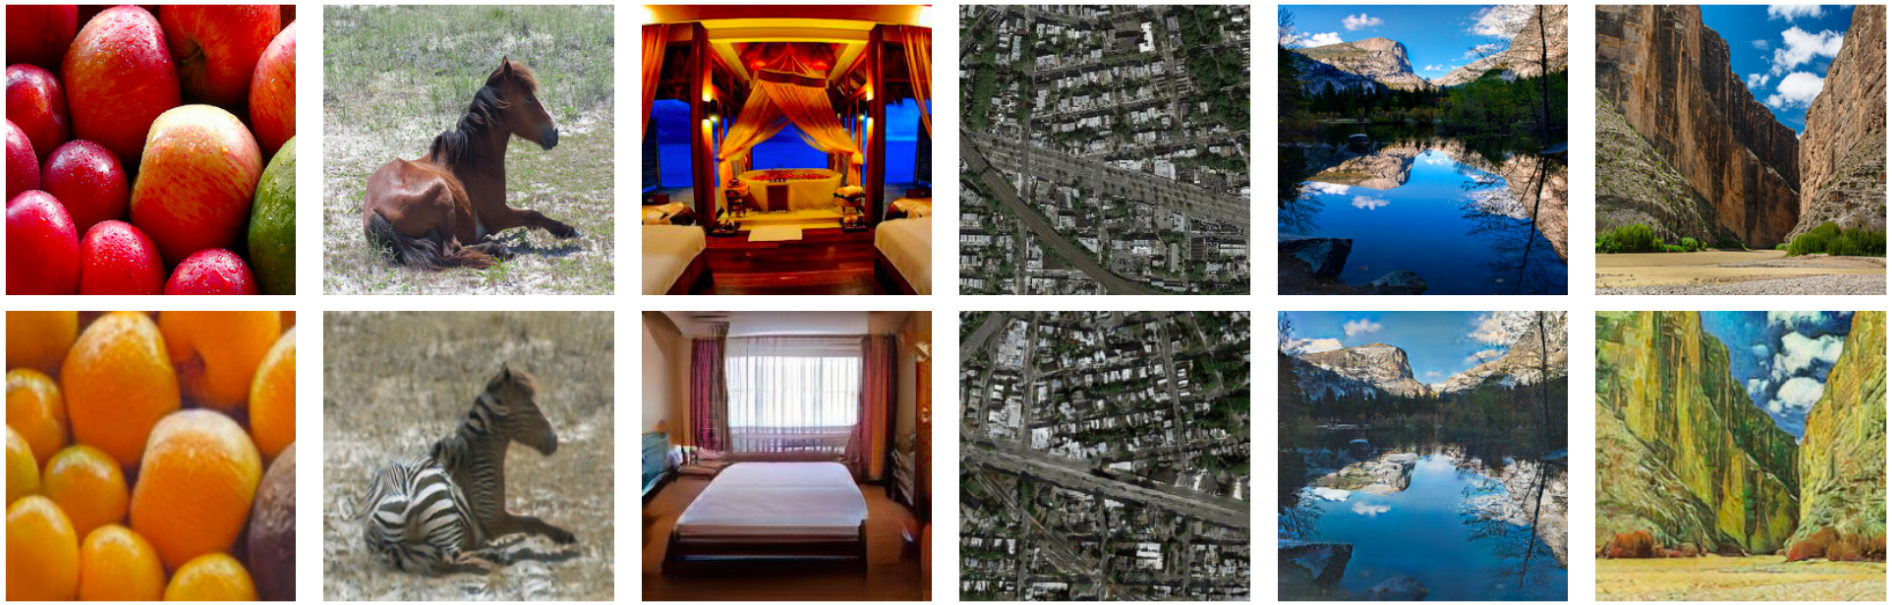

Goal

- To detect whether a picture is a natural one or it has been generated by a neural network

Main idea

- GAN images may have different statistics from natural images
- Benford's Law can capture these traces
- Given a natural image (JPEG compressed), compute the first digit (FD) of quantized DCT coefficients. It is known that distribution of the FDs follows Benford's law.

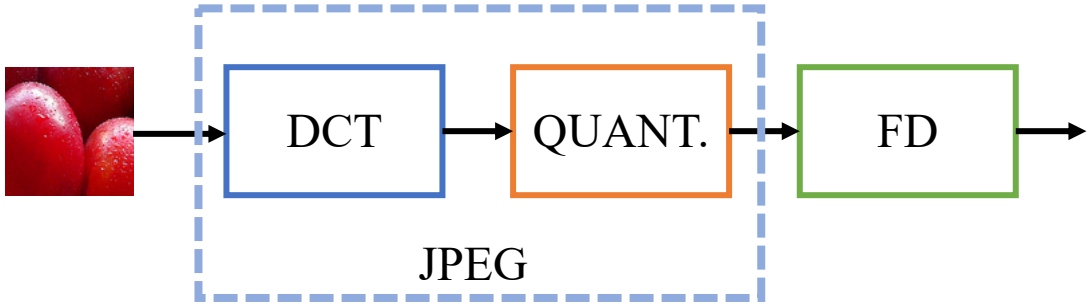

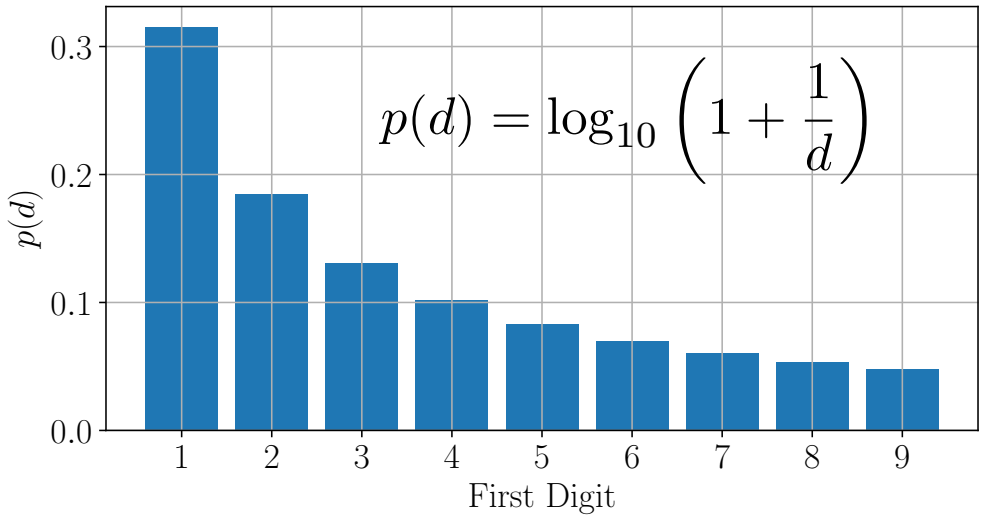

Proposed solution

1. Given a query image, compute FD of quantized DCT coefficients.
2. Compute probability distribution through histogram computation $\hat{p}$
3. Fit theoretical Benford's curve p
4. Use deviation between p and $\hat{p}$ as an element of the feature vector

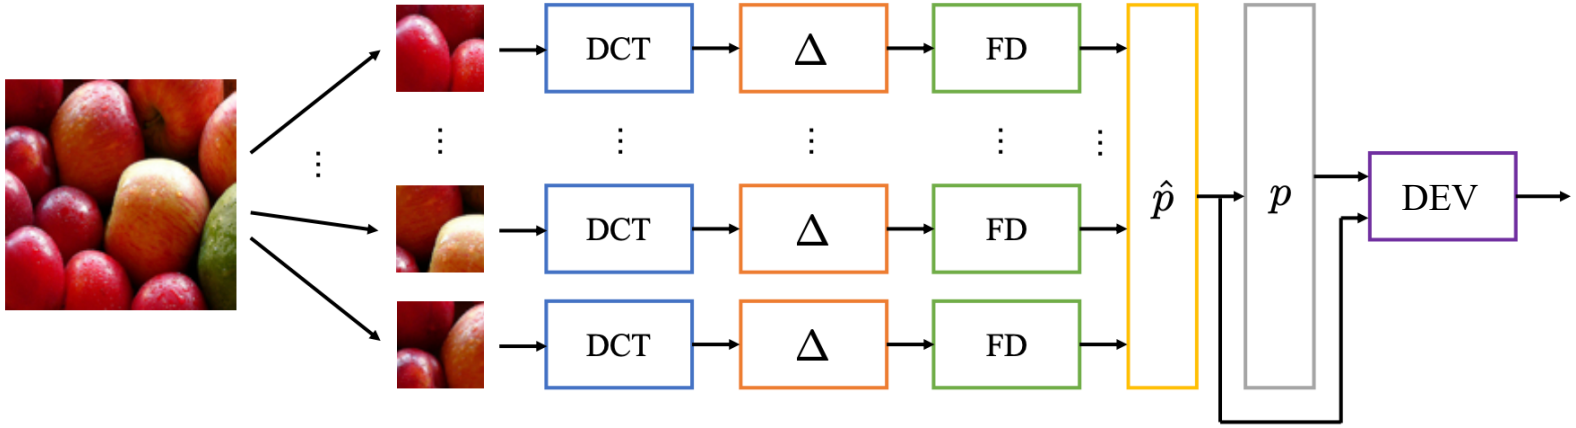

Train a simple classifier (Random Forest) to discriminate between real and generated images

Building the feature vector

- Different feature vectors generated by combining:
  - 9 possible DCT frequencies
  - 5 possible JPEG quality factors (QF)
  - 4 possible bases for computing the first digits

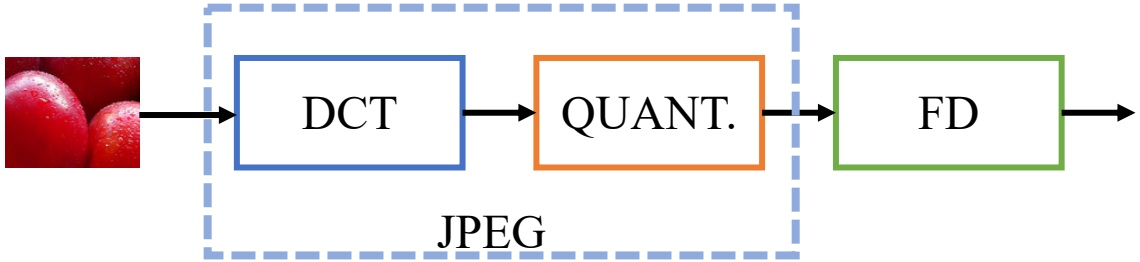

- We end up with 675 possible feature vectors (setups)

- Depending on the number of DCT coefficients, JPEG QF and bases, feature vector length can vary from 3 to 540 elements

Dataset

- Publicly available images from [1]

Architecture	Dataset	Number of images
Cycle-Gan	orange2apple	1280
	photo2ukiyoe	4072
	winter2summer	1484
	zebra2horse	1670
	photo2cezanne	3978
	photo2vangogh	4099
	photo2monet	4765
	facades	259
	cityscapes	1996
	sats	684
ProGAN	lsun_bedroom	30770
	lsun_bridge	28768
	lsun_churchoutdoor	29120
	lsun_kitchen	42706
	lsun_tower	29020

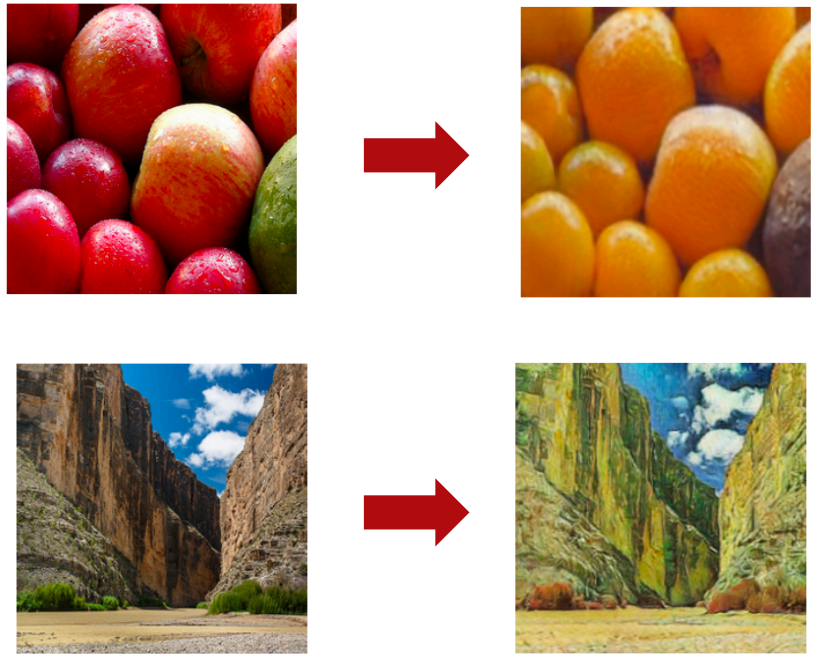

[1] F.Marra, D.Gragnaniello, L.Verdoliva, G.Poggi, "Do GANs Leave Artificial Fingerprints?" *IEEE International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2019

Results on uncompressed images

- Train a Random Forest classifier for each different setup, with Leave One Group Out policy
- Comparison with [2] and with a baseline Xception network trained for the purpose

Dataset	Proposed	Xception	Steganalysis SVM	Steganalysis RF
orange2apple	**98.13**	97.64	88.80	76.49
photo2ukiyoe	**100.00**	97.41	86.78	87.90
winter2summer	**100.00**	68.33	77.96	68.89
zebra2horse	**99.69**	89.58	91.01	77.00
photo2cezanne	**99.97**	95.91	95.88	93.17
photo2vangogh	**100.00**	93.75	94.68	92.93
photo2monet	**99.84**	94.08	94.80	89.87
facades	**100.00**	99.84	73.93	76.06
cityscapes	**100.00**	**100.00**	**100.00**	**100.00**
sats	**99.69**	73.00	90.92	96.93
lsun_bedroom	**100.00**	76.22	98.92	99.25
lsun_bridge	**99.89**	82.49	95.90	95.16
lsun_churchoutdoor	**99.99**	99.79	98.81	99.12
lsun_kitchen	**99.99**	87.26	99.49	99.59
lsun_tower	**99.98**	95.45	98.87	99.19
avg	**99.83**	89.64	91.03	90.11

[2] F. Marra, D. Gragnaniello, D. Cozzolino, and L. Verdoliva, "Detection of GAN-Generated Fake Images over Social Networks," *IEEE International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2018.

Results on uncompressed images (2)

- The length of the feature vector can be tailored to specific needs of accuracy or time constraints

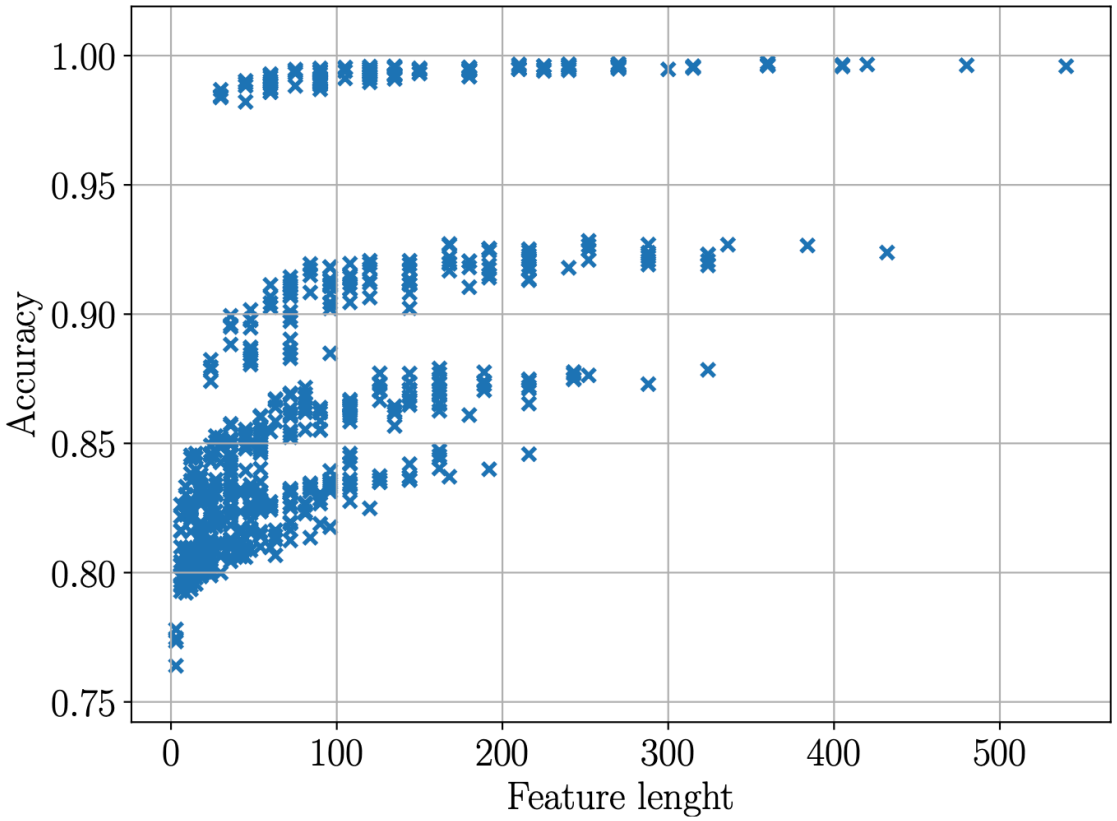

Results on compressed images

- To investigate a more realistic scenario, we compressed the images with different JPEG QF and retrained
- Comparison with a baseline Xception network trained for the purpose

QF	Dataset	Proposed	Xception
100	orange2apple	**94.50**	92.56
	photo2ukiyoe	**100.00**	98.50
	cityscapes	**100.00**	100.00
	lsun_tower	**100.00**	94.64
95	orange2apple	82.01	**90.66**
	photo2ukiyoe	97.00	**98.42**
	cityscapes	**99.99**	99.32
	lsun_tower	**99.80**	99.48
90	orange2apple	65.93	**85.61**
	photo2ukiyoe	92.01	**98.17**
	cityscapes	**100.00**	99.66
	lsun_tower	**99.60**	98.86

Preliminary results on faces

- Dataset composed only by human faces:
  - All faces from [1] (ProGAN, StarGAN, GlowGAN)
  - Additional faces generated by recent StyleGan2



Dataset	Proposed
progan_celeba	79.75
stargan_black_hair	**97.26**
stargan_blond_hair	**96.56**
stargan_brown_hair	**96.76**
stargan_male	**96.24**
stargan_smiling	**96.06**
glow_black_hair	**86.56**
glow_blond_hair	**88.26**
glow_brown_hair	**86.18**
glow_male	**87.11**
glow_smiling	83.04
stylegan2-0.5	77.18
stylegan2-1	72.63
avg	**87.96**

[1] F.Marra, D.Gragnaniello, L.Verdoliva, G.Poggi, "Do GANs Leave Artificial Fingerprints?" *IEEE International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2019

Conclusions

- We propose a handcrafted feature extraction pipeline to perform GAN-generated image detection
- Our method requires low computational power, and the feature vector length can be tailored to the desired accuracy and amount of time
- Our method can be trained when a limited set of data is available, e.g., only a few images from a brand-new GAN
- We achieve the best accuracies on uncompressed images, and we are still competitive when dealing with further JPEG compression
- Preliminary results on faces are promising, we need to improve on newer GAN architectures

GAN-generated images are realistic, but they still have statistical inconsistencies

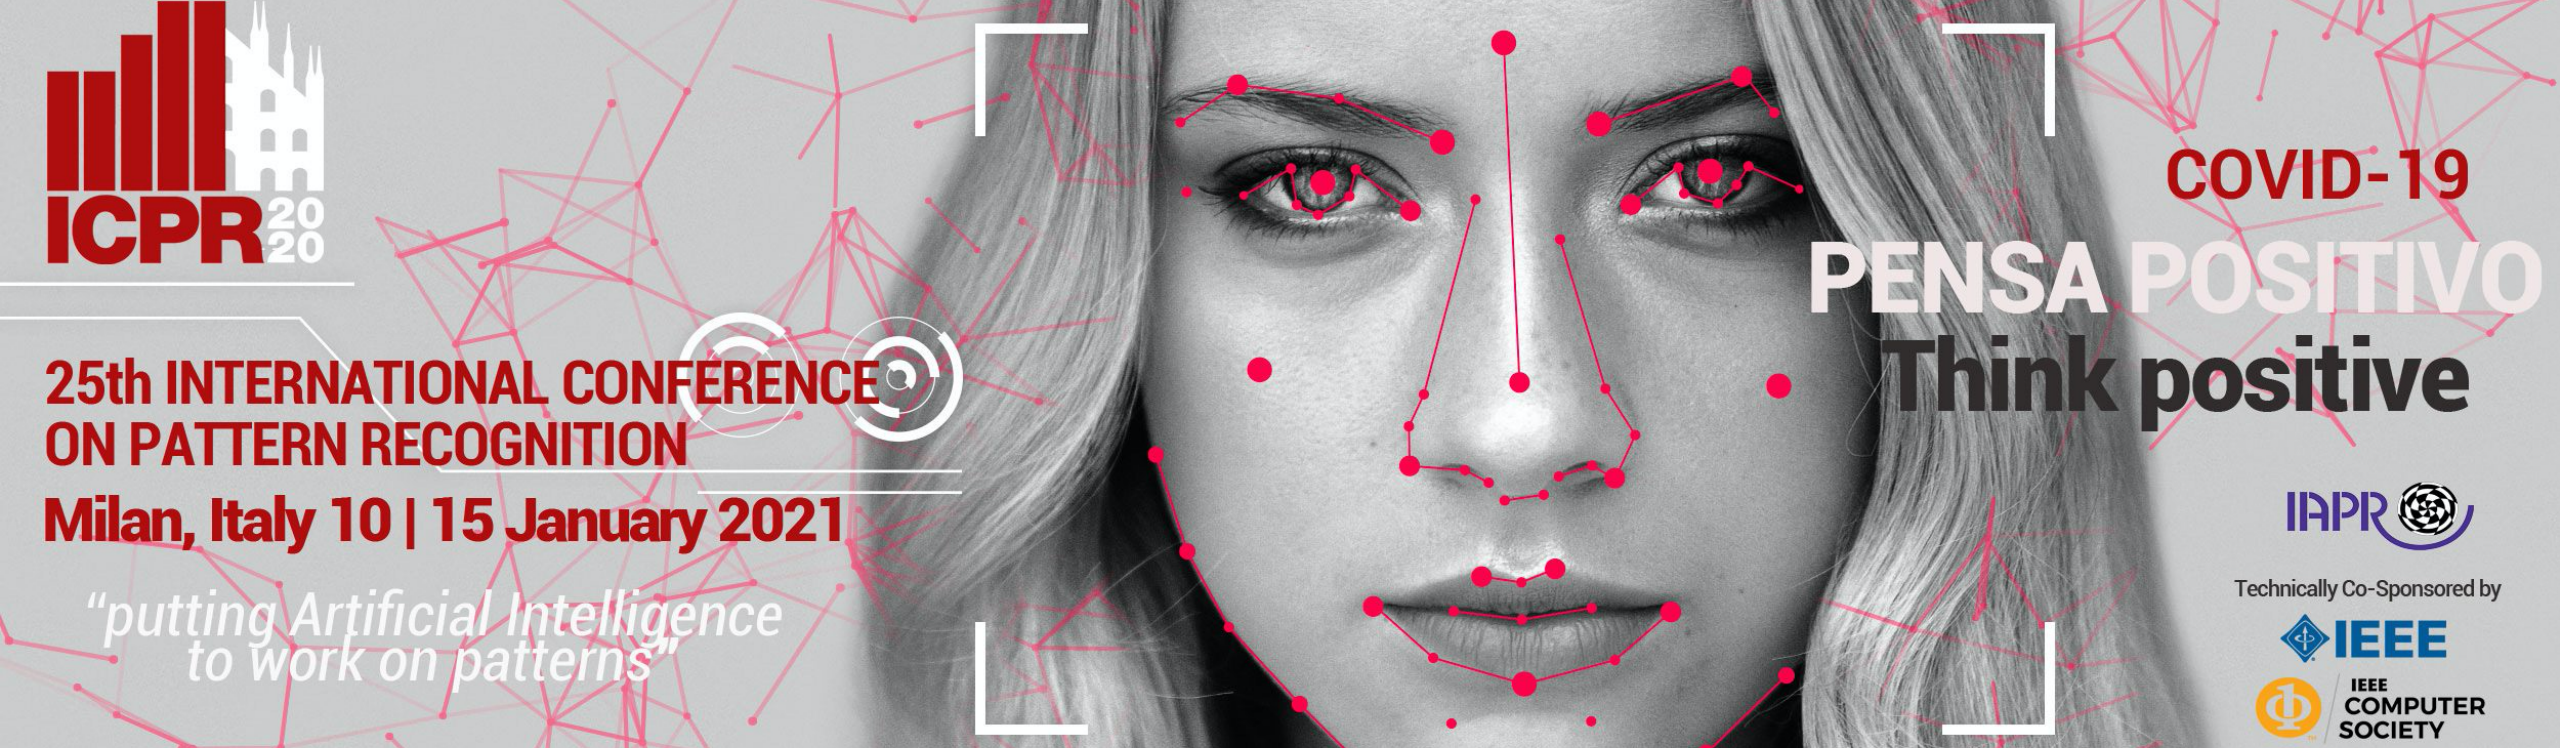

Thank you!

POLITECNICO MILANO 1863

UNIVERSITÀ DEGLI STUDI DI PADOVA