

# Photographic Visualization of Weather Forecasts with Generative Adversarial Networks

ECMWF Machine Learning Workshop

March 31<sup>st</sup>, 2022

Christian Sigg (MeteoSwiss), Flavia Cavallaro (Comerge),  
Tobias Günther (FAU) and Martin R. Oswald (ETH Zürich and UvA)

Contact: [christian.sigg@meteoswiss.ch](mailto:christian.sigg@meteoswiss.ch)


# Outline

- Motivation: Why Photographic Images?
- Baseline and Evaluation Criteria
- Method: Conditional GANs
- Results
- Conclusions and Future Work



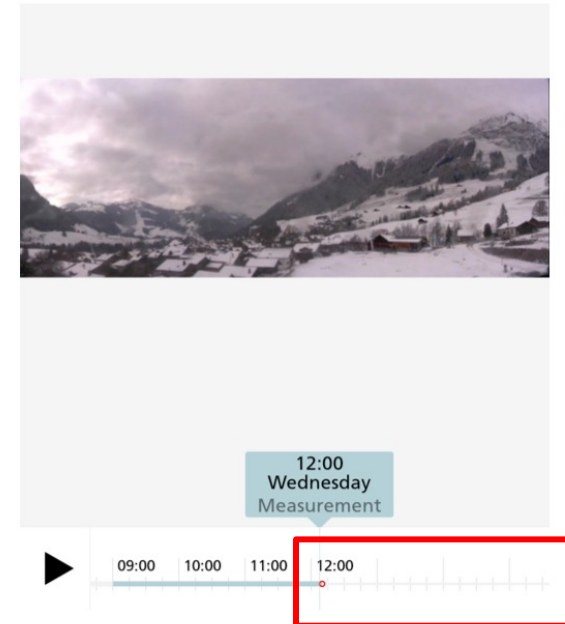
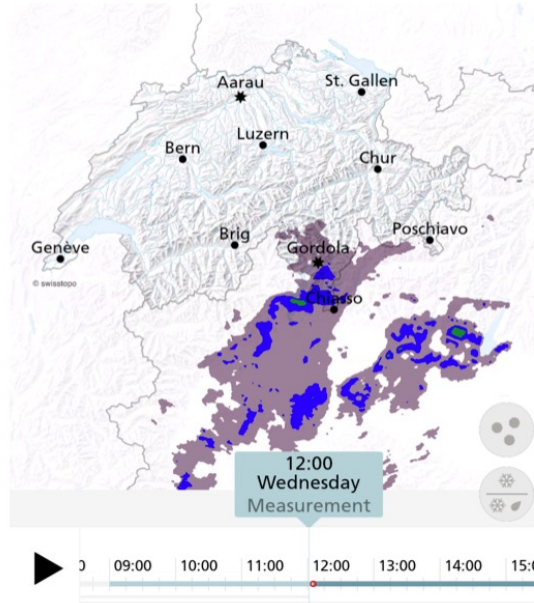
# Outdoor Weather Cameras

An information-dense yet accessible visualization of past and present weather:



# Visualization of Weather Forecasts

## Screenshots of the MeteoSwiss smartphone app




Also use photographic images to visualize future weather conditions!

# Outline

- Motivation: Why Photographic Images?
- Baseline and Evaluation Criteria**
- Method: Conditional GANs
- Results
- Conclusions and Future Work



# Baseline: Analog Retrieval



$I_t$  Image sequence taken at Flüela, 10 to 16 UTC on July 2<sup>nd</sup>, 2020

$\hat{I}_t^{ind}$  Retrieval of best matching individual images from annotated archive

$\hat{I}_t^{seq}$  Retrieval of best matching sequence

# Proposed Evaluation Criteria



- I. Images should look real, no obvious artifacts
- II. Match future atmospheric, ground and illumination conditions
- III. Seamless transition from observation to forecast
- IV. Visual continuity between consecutive images

# Evaluation of Analog Retrieval



	<b>I. Realism</b>	<b>II. Matching conditions</b>	<b>III. Seamless transition</b>	<b>IV. Visual continuity</b>
<b>Analog images</b>	😊	😐	😞	😞
<b>Analog sequence</b>	😊	😞	😞	😊

High information density of images → retrieving analogs is not feasible 🤔



# Outline

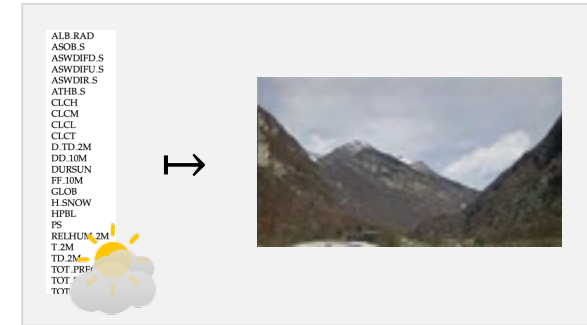
- Motivation: Why Photographic Images?
- Baseline and Evaluation Criteria
- Method: Conditional GANs**
- Results
- Conclusions and Future Work



# Image Synthesis: A Regression Problem

Generate photographic image  $\hat{I}_t$ , given forecast  $w_t$  of future weather conditions

$$G: w_t \mapsto \hat{I}_t$$



Generator  $G(w; \theta)$  is a neural network,  $\theta$  trained by minimizing expected loss

$$\operatorname{argmin}_{\theta} \mathbb{E}_{w_t, I_t} [L(G(w; \theta), I_t)]$$

# Choice of Loss Function $L$

$$\operatorname{argmin}_{\theta} \mathbb{E}_{w_t, I_t} [L(G(w_t; \theta), I_t)]$$

Forecast  $w_t$  does not determine exact shapes and locations of clouds  
→ Pixel-wise loss function is not appropriate, results in uniform sky:



$\hat{I}_t$  for  $L_1$  loss



$I_t$

Goal: User should not be able to tell whether  $I_t$  or  $\hat{I}_t$  is the real image, even if they are not identical.

# Generative Adversarial Networks Goodfellow et al., 2014

**Discriminator**  $D: I \mapsto [0, 1]$  mimics user, learns loss function through adversarial training

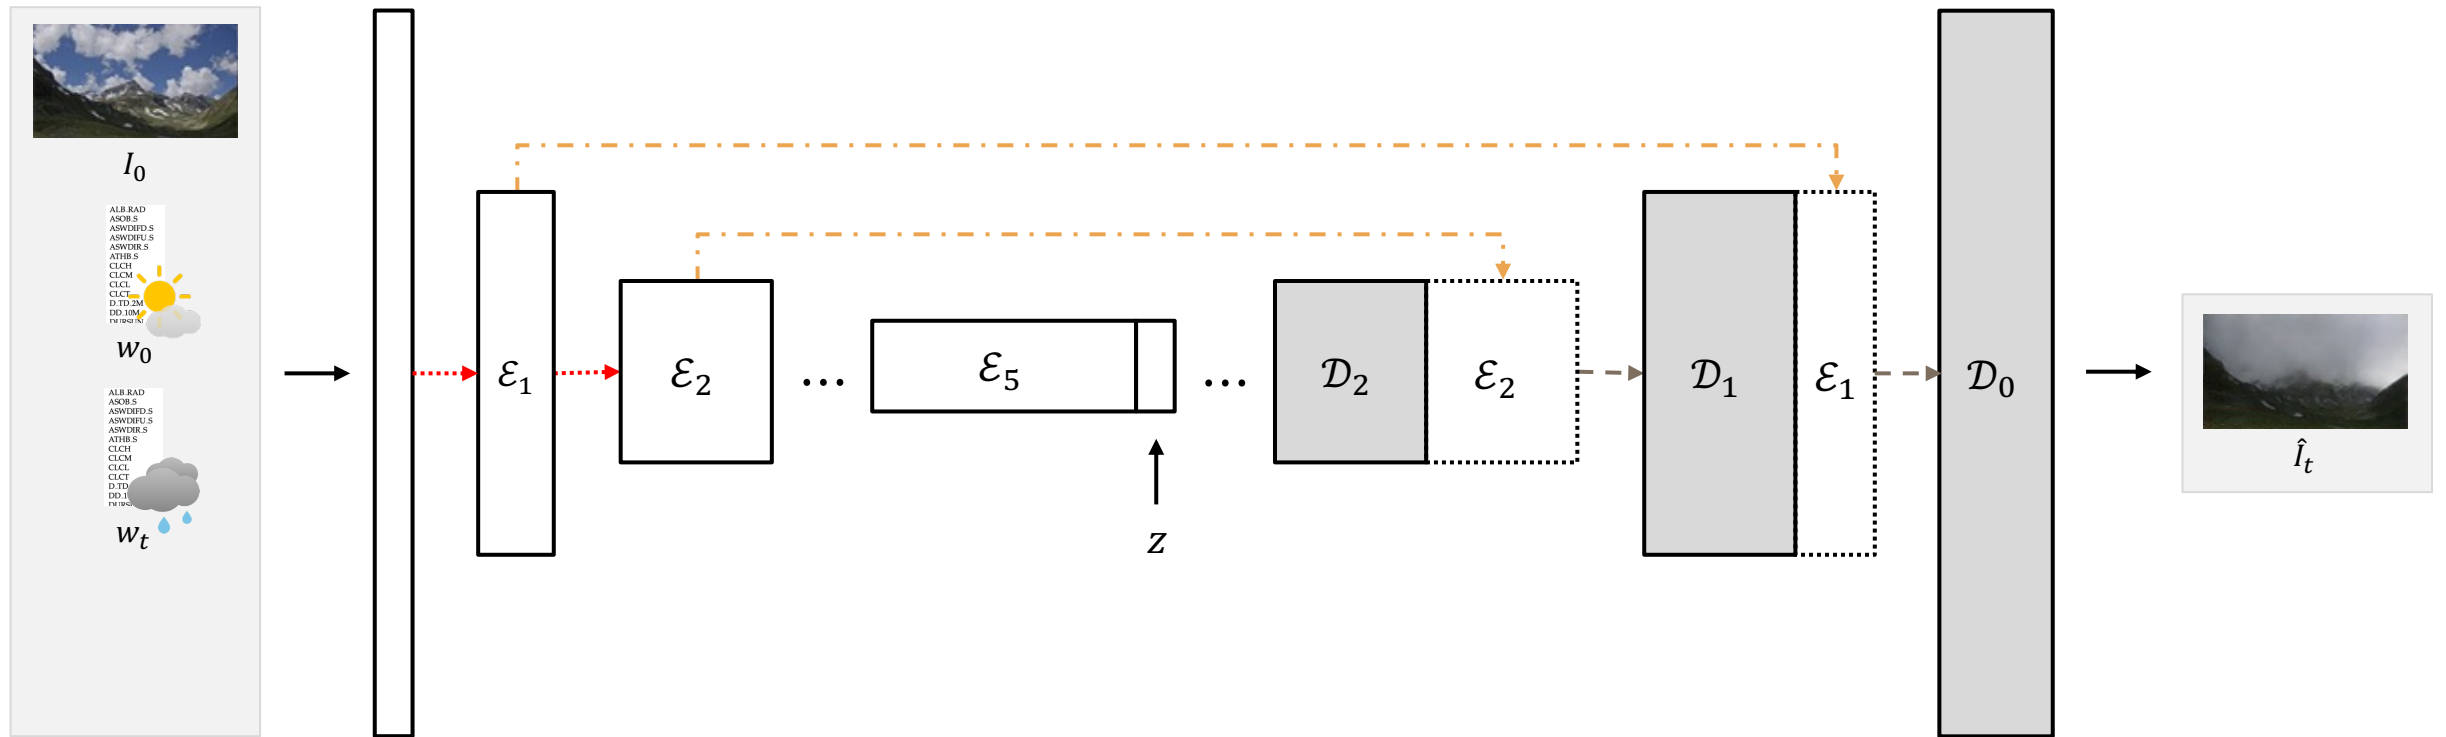
**Generator**  $G: z \mapsto I$ , creates image  $I$  from random input  $z \sim \mathcal{N}(0, 1)$

$$\min_{\theta} \max_{\eta} \underbrace{\mathbb{E}_I[\log D(I; \eta)]}_{\text{authenticate real images}} + \underbrace{\mathbb{E}_z[\log\{1 - D(G(z; \theta); \eta)\}]}_{\text{fool discriminator}}$$

---

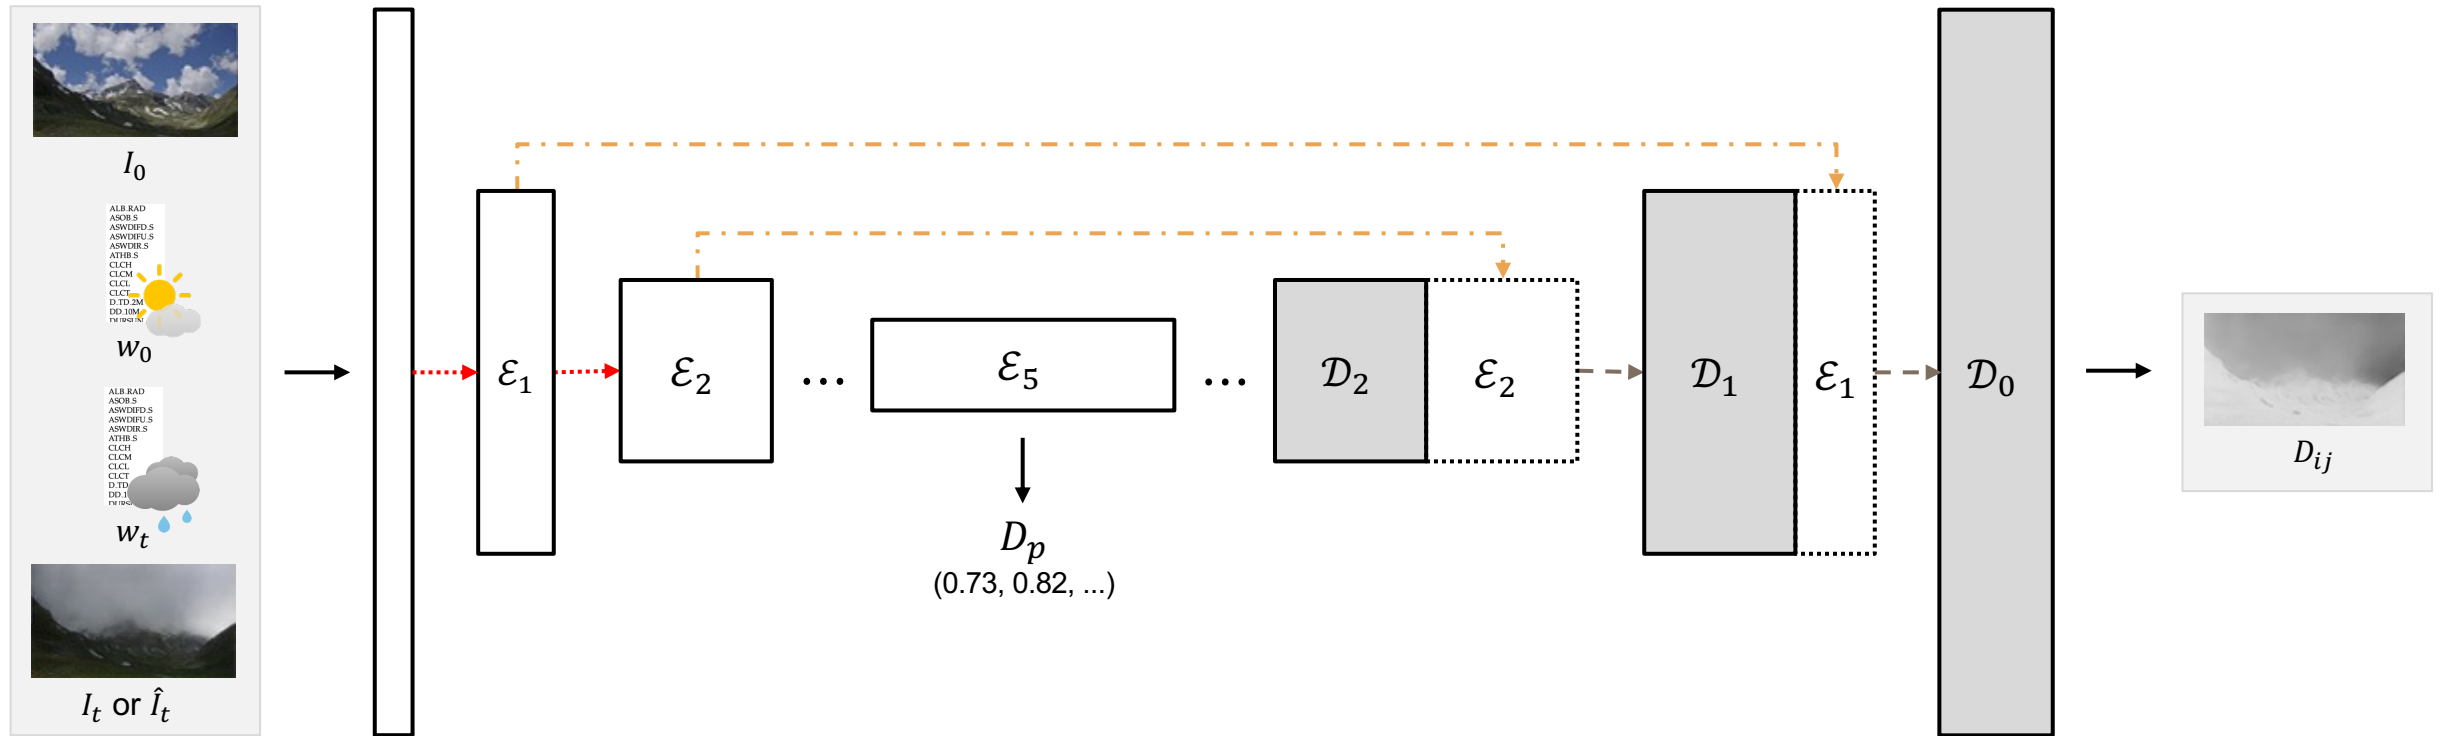
*spot fake images*

# Generator Architecture



- Conditional Generator Mirza and Osindero, 2014 transforms current image  $I_0$
- Encoder-decoder with skip connections Ronneberger et al., 2015
- Spectral normalization applied to all convolution layers Miyato et al., 2018

# Discriminator Architecture



- Conditional discriminator  $D(I|I_0, w_0, w_t)$
- Two output heads: patch-level  $D_p$  and pixel-level  $D_{ij}$  Schonfeld et al., 2020

# Outline

- Motivation: Why Photographic Images?
- Baseline and Evaluation Criteria
- Method: Conditional GANs
- Results**
- Conclusions and Future Work



# Evaluation Data



*Cevio (elevation 421 m)*



*Etziken (elevation 524 m)*



*Flüela (elevation 2177 m)*

**Descriptor  $w$ :** time of day, day of year, 31 COSMO-1 hourly output fields

**Training:** all pairs  $(I_0, w_0)$  and  $(I_t, w_t)$ ,  $t \in [0, 10, 20, \dots, 360 \text{ min}]$  of 2019

**Test:** Jan to Aug of year 2020 (until decommissioning of COSMO-1 at MCH)

Downscaled to 64 x 128 pixels to speed up training and conserve GPU memory



# I. Realism

What is your first impression of the image?



*generated*



*real*



*generated*



*real*



# I. Realism

Results of study with 5 professional users of MCH camera feeds:

Actual	Judgment	
	Real	Generated
Real	57	18
Generated	43	32

*Cevio: 59 % accuracy*

Actual	Judgment	
	Real	Generated
Real	52	23
Generated	32	43


*Etziken: 63 % accuracy*

Actual	Judgment	
	Real	Generated
Real	57	18
Generated	49	26

*Flüela: 55 % accuracy*

User accuracy is not much better than random guessing 😊

## II. Matching Weather Conditions



**Atmosphere:** cloud cover, cloud type, visibility

**Ground:** dry, wet, frost, snow

**Illumination:** time of day, diffuse or direct

## II. Matching Weather Conditions

Camera	Matching conditions					
	Atmosphere			Ground	Illumination	
	Cloud cover	Cloud type	Visibility		Time of day	Diffuse/direct
Cevio	73	78	98	100	100	88
Etziken	74	83	96	100	100	85
Flüela	63	68	60	99	95	75

Example: Mismatch in cloud cover



$\hat{I}_t$



$I_t$


but forecast  $w_t$  predicted 100 % cloud area fraction in medium troposphere!

## II. Matching Weather Conditions

Camera	Matching conditions						
	Atmosphere			Ground	Illumination		Viz. failures
	Cloud cover	Cloud type	Visibility		Time of day	Diffuse/direct	
Cevio	73	78	98	100	100	88	10
Etziken	74	83	96	100	100	85	6 😊
Flüela	63	68	60	99	95	75	19

**Visualization failure:** forecast  $w_t$  is accurate, but generated image  $\hat{I}_t$  is inconsistent with it

# III. Seamless Transition – IV. Visual Continuity



Possible because  $G$  is conditioned on  $I_0$ , compare to analog retrieval:



# Outline

- Motivation: Why Photographic Images?
- Baseline and Evaluation Criteria
- Method: Conditional GANs
- Results
- Conclusions and Future Work**

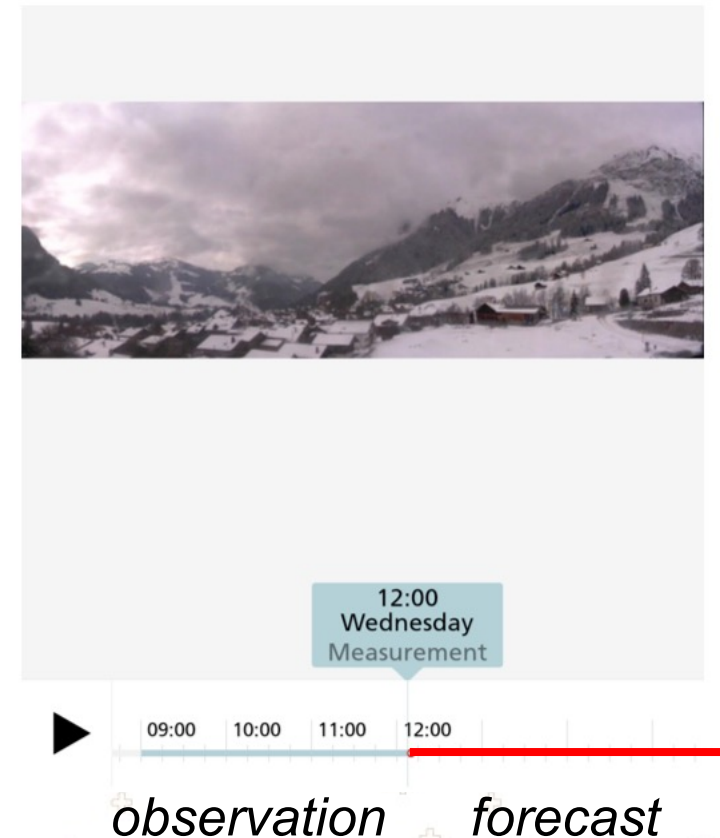


# Conclusions

- Photographic images can also visualize future weather conditions
- Look realistic, match predicted weather conditions, attain seamless transition from observation to forecast and visual continuity

## Applications:


- Communicate localized forecasts in webcam feeds, smartphone app
- Provide similar service to communities and tourism organizations





# Future Work

- Use more accurate and descriptive weather descriptors, to improve matching the forecast *and* the conditions visible in the future
- Scale image size beyond 64 x 128 pixels e.g. using Karras et al., 2018
- Improve transformations involving translations of isolated clouds:



(Including self-attention layers Zhang et al., 2019 did not help)

- Synthesize whole sequences to improve temporal evolution Wu et al., 2020

# Resources

The pre-print of our paper is available at

<https://arxiv.org/abs/2203.15601>

Tensorflow code, trained models and results are available at

<https://github.com/meteoswiss/photocast>



# Acknowledgments

We thank Rega for giving us the permission to use images from the Cevio camera in this study.

We thank Tanja Weusthoff for the preparation of the COSMO-1 forecast data.

We thank Christian Allemann, Yannick Bernard, Eliane Thürig, Deborah van Geijtenbeek and Abbès Zerdouk for evaluating the realism of individual generated images.

We thank Daniele Nerini for providing his expertise on nowcasting and post-processing of forecasts.

# Bibliography

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770-778).

Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.

Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.

Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*.

Ronneberger, O., Fischer, P., & Brox, T. (2015, October). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (pp. 234-241).

U. Schättler, G. Doms, and C. Schraff. (2021). COSMO-Model Version 6.00: A Description of the Non- hydrostatic Regional COSMO-Model - Part VI: Model Output and Data Formats for I/O.

Schonfeld, E., Schiele, B., & Khoreva, A. (2020). A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 8207-8216).

Wu, S., Xiao, X., Ding, Q., Zhao, P., Wei, Y., & Huang, J. (2020). Adversarial sparse transformer for time series forecasting. *Advances in Neural Information Processing Systems*, 33, 17105-17115.

Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., & Yoo, Y. (2019). Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 6023-6032).

Zhang, H., Goodfellow, I., Metaxas, D., & Odena, A. (2019, May). Self-attention generative adversarial networks. In *International conference on machine learning* (pp. 7354-7363). PMLR.

# Generator Objectives to be Minimized

How much  $G(I_0, z|w_0, w_t)$  struggles to fool the discriminator on the patch level

$$\mathbb{E}_{I_0, w_0, w_t} \mathbb{E}_z \left[ \sum_p \log [D_p(G(I_0, z|w_0, w_t)|I_0, w_0, w_t)] \right]$$

and on the pixel level

$$\mathbb{E}_{I_0, w_0, w_t} \mathbb{E}_z \left[ \sum_{ij} \log [D_{ij}(G(I_0, z|w_0, w_t)|I_0, w_0, w_t)] \right]$$

How similar two generated images look at the pixel level, given different random inputs  $z_1, z_2 \sim \mathcal{N}(0, 1)$

$$-\mathbb{E}_{I_0, w_0, w_t} \mathbb{E}_{z_1, z_2} \left[ \sum_{ijc} |G_{ijc}(I_0, z_1|w_0, w_t) - G_{ijc}(I_0, z_2|w_0, w_t)| \right]$$

# Discriminator Objectives to be Maximized

How well the patch head  $D_p$  authenticates real images

$$\mathbb{E}_{I_0, w_0, I_t, w_t} \left[ \sum_p \log D_p(I_t | I_0, w_0, w_t) \right]$$

and spots generated images

$$\mathbb{E}_{I_0, w_0, w_t} \mathbb{E}_z \left[ \sum_p \log [1 - D_p(G(I_0, z | w_0, w_t) | I_0, w_0, w_t)] \right]$$

How well the pixel head  $D_{ij}$  can distinguish pixels of a cut-mix composite  $C$  Yun et al., 2019

$$\mathbb{E}_C \left[ \sum_{ij} M_{ij} D_{ij}(C) + (1 - M_{ij}) \log(1 - D_{ij}(C)) \right]$$

# Artifacts Induced by Residual Learning He et al., 2015



Present image  $I_0$



Forecast visualization  $\hat{I}_t$

Clouds in  $I_0$  are still partially visible in the clear sky regions of  $\hat{I}_t$

→ Residual transformation learned by the generator does not fully cancel their appearance

# Subset of COSMO-1 Output Fields Schättler et al., 2021

Abbreviation	Unit	Name
ALB_RAD	%	Surface albedo for visible range, diffuse
ASOB_S	W/m <sup>2</sup>	Net short-wave radiation flux at surface
ASWDIFD_S	W/m <sup>2</sup>	Diffuse downward short-wave radiation at the surface
ASWDIFU_S	W/m <sup>2</sup>	Diffuse upward short-wave radiation at the surface
ASWDIR_S	W/m <sup>2</sup>	Direct downward short-wave radiation at the surface
ATHB_S	W/m <sup>2</sup>	Net long-wave radiation flux at surface
CLCH	%	Cloud area fraction in high troposphere (pressure below ca. 400 hPa)
CLCM	%	Cloud area fraction in medium troposphere (between ca. 400 and 800 hPa)
CLCL	%	Cloud area fraction in low troposphere (pressure above ca. 800 hPa)
CLCT	%	Total cloud area fraction
D.TD_2M	K	2 m dew point depression
DD_10M	°	10 m wind direction
DURSUN	s	Duration of sunshine
FF_10M	m/s	10 m wind speed
GLOB	W/m <sup>2</sup>	Downward shortwave radiation flux at surface
H.SNOW	m	Snow depth
HPBL	m	Height of the planetary boundary layer
PS	Pa	Surface pressure (not reduced)
RELHUM_2M	%	2 m relative humidity (with respect to water)
T_2M	K	2 m air temperature
TD_2M	K	2 m dew point temperature
TOT_PREC	kg/m <sup>2</sup>	Total precipitation
TOT_RAIN	kg/m <sup>2</sup>	Total precipitation in rain
TOT_SNOW	kg/m <sup>2</sup>	Total precipitation in snow
U_10M	m/s	10 m grid eastward wind
V_10M	m/s	10 m grid northward wind
VMAX_10M	m/s	Maximum 10 m wind speed