


Cloud Nowcasting Using Deep Neural Nets

Eric Gritmit, Ph.D.

Head of Science – North America

Vaisala Xweather

Solar Power Forecast Limitations



Baseline: A forecast that considers local observations and NWP only is not very skillful

Spatial Awareness Could Be Helpful



- NWP-resolved cloudiness reduces the solar power forecast partially, but not enough.
- Past observations (clear-sky adjusted) help in this case, but for the wrong reasons.
- No direct satellite ingredient included in this forecast, but clearly one is needed!

Satellite-Based Cloud Nowcasting

AMS Glossary of Meteorology Definitions:

Nowcast = 0-3 hours

Very short-range forecast = 0-6 hours





Figure 11. Techniques suitable for different forecasting's horizon and spatial resolution. (L. Ramírez & Vindel. 2016)

Cloud Mask Extrapolation

- Create a binary cloud mask; segment satellite image into cloudy and non-cloudy pixels.
- Assign or compute cloud-motion vectors (CMV) from successive satellite images, atmospheric motion vector retrievals, or NWP model estimates.
- Advect the binary cloud mask forward in time using the CMVs as a forecast.
- This kind of technique is the basis for EUMETSAT's "kinematic extrapolation" product known as EXIM.



Optical Flow

- A method for identifying the **pattern of apparent motion** of objects and edges with a series of subsequent images.
- Gradient-based estimation of cloud motion vectors (CMV) computed from consecutive images (e.g., Lucas-Kanade, Gunnar-Farneback, or Horn-Schunck methods).
- The resulting CMV field is used to warp the latest image, resulting in a forecast image.
- Can only predict the movement of features found in current image, not longer-term evolutions.
- Physical processes such as orographic lift, convergence, and convective initiation are not well captured.





FIGURE 2 The optical flow analysis for satellite images.

From Guo et al. (2023) IET Computer Vision

NCAR Sun4Cast System

- Sun4Cast leverages several observation-based nowcasting technologies, each with its own “sweet spot”.
- Includes CIRACast and MADCast components, which take different approaches to satellite-based cloud advection.
- These technologies are blended via the Nowcasting Expert System Integrator (NESI).
- Improvement over smart persistence varies by month/year, lead time and sky condition.
- Over lead times <6h, improvement ranged from 36-56% in all-sky conditions for each component.
- Satellite-based optical flow was not a comparison baseline in this study.



From Figure 5, Haupt et al. (2018) BAMS

Deep Neural Nets with Satellite Data

- Over the last 5+ years, a number of studies have tested deep learning on satellite imagery for the purpose of better solar energy forecasting.
- Common deep learning model architectures, such as convolutional neural networks (CNNs) and encoder-decoders (U-Nets) were applied to this problem first.

Cloud Cover Nowcasting with Deep Learning (Berthomier et al. 2020) [Meteo France AI Lab]

- EUMETSAT Met-11 cloud cover analysis
- All selected models improved over persistence
- U-net surpassed AROME, EXIM models






Fig. 11. Average MSE of the U-Net model, persistence, EXIM and AROME for each time step of the forecast. The dashed line represent the MSE computed on binarized values.

DeePSat: A Deep Learning Model for Prediction of Satellite Images for Nowcasting Purposes (Ionescu et al. 2021) [Romania]

- EUMETSAT Met-11, 5 channels
- Improved over a baseline CNN, but not as much as recurrent approaches




Deep Neural Nets with Satellite Data (cont.)

- Later approaches tested model architectures which combine encoders with recurrent units to capture the spatial covariances and time series evolution together.
- Motivation increased because these techniques were shown to significantly outperform both optical flow and state-of-the-art NWP in precipitation nowcasting.

- **SunCast: Solar Irradiance Nowcasting from Geosynchronous Satellite Data** (Kumareson et al. 2022) [UC Berkeley]:

- GOES-16 downward solar radiation (DSR)
- ConvLSTM outperformed HRRR model by 13% in RMSE (daytime 10-15:00) for 4-week test



Grouping	HRRR RMSE	Model RMSE
Overall	124.9	108.6
Low DSR (0-300)	165.3	135.3
Medium DSR (300-600)	170.7	131.7
High DSR (600+)	103.5	98.3

Cloud Nowcasting with Structure-Preserving Convolutional Gated Recurrent Units (Kellerhals et al. 2022) [U. Amsterdam]

- ConvGRU w/ structure preserving loss function
- Beats optical flow MAE by 9-12% at 1-3h leads; in turn, optical flow beats persistence by 11-17%




Table 2. Average percentage differences in accuracy metrics of ConvGRU predictions against the optical flow ensemble baseline, grouped by lead time (τ).

τ	Model	% ΔR^2	% ΔMAE	% $\Delta SSIM$
4	MSE	4.27	-1.78	1.42
	MAE	5.95	0.12	3.69
	Huber	6.33	-2.46	3.12
	SSIM	8.11	-0.06	8.63
	SSIM + MAE	7.63	-9.27	7.28
8	MSE	13.43	-7.42	8.86
	MAE	15.17	-6.50	10.53
	Huber	16.05	-2.17	9.92
	SSIM	15.95	-7.65	13.61
	SSIM + MAE	16.43	-11.95	13.03
12	MSE	19.23	-6.38	11.59
	MAE	20.96	-7.23	12.85
	Huber	22.29	-0.72	12.31
	SSIM	21.10	-9.48	15.53
	SSIM + MAE	22.21	-11.18	15.07

Xweather CloudCast

Vaisala's specialist deep neural network (DNN) for
satellite-based cloud nowcasting



Our Approach

Multi-Modality:

- Use multiple channel satellite data as inputs and targets.
- Utilizing both visible and infrared radiances benefits the other and assists with the night-to-daytime transition.


Image-to-Image Sequences:

- Use a stack of recent images to predict the next few images.
- Retain the native projection and resolution of the source images. No re-mapping or interpolation at this stage.

Timing Requirements:

- Rapid updates created every 5 min. Use limited-area scan windows. Target 4-5 min as the maximum inference time.
- Create forecasts for the next 3 hours to cover the needs of intra-hour solar power nowcasting with a buffer for failover.

Training data (in the native projection of GOES-16 satellite)





DNN's input tile
(mega-tile
context region)

DNN's target tile

Training Data (GOES-16 Example)

- Primary: GOES-16 CONUS window (5-min)
 - Channel 02 (visible, 0.64 μm)
 - Channel 07 (near-infrared, 3.9 μm)
 - Channel 14 (infrared, 11.2 μm)
 - Level-2 ACMC (binary cloud mask)
- Auxiliary:
 - Elevation (90-m; NASA SRTM)
 - Solar zenith and azimuth angles
 - Forecast lead time (5, 10, ..., 180 minutes)




First Attempts


Candidate model architecture evaluation:

- Encoder-ConvLSTM-Attention (similar to Google's MetNet)
 - A leading option from recent precipitation nowcasting advancements, with performance superior to persistence and NOAA's 3-km HRRR
 - Modified to reduce the number of trainable parameters
 - Each target tile only produced a smooth field near the mean value
 - Scene background seemed to confuse the algorithm further
- CoaT-GRU (similar to U. Amsterdam ConvGRU)
 - Co-Scale Conv-Attentional Image Transformer (CoaT) is an efficient image transformer that performs well in classification tasks
 - Gated recurrent units (GRUs) are well suited for sequence-to-sequence prediction tasks and have a smaller number of trainable parameters compared to LSTMs
 - "Off-the-shelf", quick implementation, faster inference time
 - Early performance results immediately showed more realism

Encoder-ConvLSTM-Attention




CoaT-GRU



Chosen Model Architecture


CoaT-GRU




- An input data sequence is processed (normalized GOES image modalities).
- The conditional time and the pre-processed data sequence is condensed from 3D to 2D with a ConvGRU layer.
- 2D tensors are processed through a CoAT spatial encoder-decoder module compressing by a factor of 4.
- The conditional time and the CoAT outputs are passed through a number of ConvGRU cells for temporal refinement.
- A set of regression heads expand the predictions into the desired output layers (predicted GOES image modalities).

Baseline Comparison Methods

Recurrent All-Pairs Field Transforms (RAFT; Teed and Deng 2020)



Spectral Prognosis (S-PROG; Seed 2003, Pulkkinen et al. 2019)








- The extra effort and computational costs to train and deploy a deep neural network model should be justified by significantly out-performing baselines.



- **Benchmarks Considered:**

- (weak) Persistence: *nothing changes*
- (moderate) Lagrangian persistence: *the background flow is estimated at T=0 and it is held constant (no evolution)*
- (strong) Optical flow: *the motion field is estimated at T=0 and its evolution is modeled*
- From computer vision: Recurrent All-Pairs Field Transforms (RAFT)
- From radar/precipitation forecasting: Spectral Prognosis (S-PROG)

Evolution of Results – Infrared C14




Final Model Results – CONUS Wide




- Validation for 0-3 h CloudCast predictions of GOES-16 (East) satellite images over whole CONUS area for infrared radiance (channel 14) shown at left and visible radiance (channel 2) at right.
- 24% and 50% improvement in MAE at 30 min lead time compared to best optical flow baseline method.
- Smaller percentage improvements at shorter lead times and larger skill at longer lead times.

Final Model Results – By Climate Region




Solar Forecast Arbiter Climate Zones
<https://forecastarbiter.epri.com/climatezones/>




- For infrared radiance (channel 14) predictions at 30 min ahead, we see a 41% improvement over optical flow in the SW USA (climate region 3) and 20%-30% elsewhere.

Final Model Results – By Region & Lead Time




GOES-16 (Continental US) Example

Optical Flow (SPROG)



CloudCast (CoaT-GRU)




January 1, 2023 00:00 UTC example case:

- Channel 14 (infrared) movies for optical flow baseline (left) and deep neural network (right)
- Optical flow preserves some structure, while CloudCast shows limited performance against smoothing out the predicted fields in longer lead times, a common problem for nongenerative data-driven nowcasting models.

Meteosat-10 (Europe) Example


Optical Flow (SPROG)

2023-02-01 08:00 T = 15 min



CloudCast (CoaT-GRU)

2023-02-01 08:00 T = 15 min



February 1, 2023 08:00 UTC example case:

- Channel 02 (visible) movies for optical flow baseline (left) and deep neural network (right)
- Optical flow preserves some detailed cloud structure (possibly erroneously), but cannot handle the night-to-day transition and the changing solar angles.

Current Progress



Using modern MLOps practices, CloudCast is being deployed into our operational environment.

- + GOES-16 (CONUS-East)
- + GOES-18 (CONUS-West)
- + Meteosat-10 (Europe-RSS)




Incorporating as a new input data source for intra-hour solar power forecasts for Vaisala Xweather customers.

- + On-site irradiance & power obs
- + Multiple NWP models (ECMWF, UKMET, GFS, HRRR, ...)
- + **Satellite-based cloud nowcast**

See Pascal Storck's presentation in Session 3B (Wednesday) for further details.



Plan for Meteosat 3rd generation upgrade later in 2024 and early 2025.



Extension to other GEO satellites as needed.

Future Improvements

- Deep learning continues to rapidly evolve with disruptive new ideas coming to the earth sciences every few months (or even weeks!).
- Classes of generative AI models and new training methods are being evaluated by many researchers in both public and private sectors.
- In particular, generative adversarial networks (GANs) and diffusion models have been popular lately, but will they yield another step change in accuracy?
- Vaisala Xweather will continue to monitor, collaborate, and innovate to bring the best technologies for renewable energy forecasting.


Omnivision forecasting: Combining satellite and sky images for improved deterministic and probabilistic intra-hour solar energy predictions

Quentin Paletta^{a,b,*}, Guillaume Arbod^b, Joan Lasenby^a

^a Department of Engineering, University of Cambridge, UK

^b Lab CRIGEN, Engie, France

GRAPHICAL ABSTRACT




Conditioning Diffusion Models

Goal: guide generation toward a particular data distribution, by estimating

- an image given a class label or text embedding $p(\mathbf{x}|\mathbf{y})$
- the next image given the previous image $p(\mathbf{x}_t|\mathbf{x}_{t-1}) \leftarrow \text{this work}$

Soln: update our model to input the noisy state $\hat{\mathbf{x}}_t = \mathbf{x} + \mathbf{n}$ and the previous $\mathbf{x}_{t-1}$

$$\mathbb{E}_{\sigma, \mathbf{x}_{(t-1), p}} \mathbb{E}_{\mathbf{n}} [\lambda(\sigma) \|D(\hat{\mathbf{x}}_t, \mathbf{x}_{t-1}; \sigma) - \mathbf{x}_t\|_2^2]$$



Reverse diffusion with the input condition, individual sampling steps ($t_0 \rightarrow t_{40}$), the next time step estimate, and target output



The truth about renewable energy forecasting. Part two: "The Bad."

<https://www.xweather.com/blog/article/truth-about-renewable-energy-forecasting-part-2>