
Spatiotemporal Data Fusion for Precipitation Nowcasting

Vladimir Ivashkin
Yandex, Moscow, Russia
vlivashkin@yandex-team.ru

Vadim Lebedev
Yandex, Moscow, Russia
cygnus@yandex-team.ru

Abstract

Precipitation nowcasting using neural networks and ground-based radars has become one of the key components of modern weather prediction services, but it is limited to the regions covered by ground-based radars. Truly global precipitation nowcasting requires fusion of radar and satellite observations. We propose the data fusion pipeline based on computer vision techniques, including novel inpainting algorithm with soft masking.

1 Introduction

Precipitation nowcasting is a problem of precipitation prediction on a very short term (up to two hours). In this time range, extrapolation techniques working with detailed observational data can compete with and even win over meteorological models. The sources of data for this problem include ground-based radars and geostationary satellites. Radars provide more detailed and accurate observations, so using radars is more desirable when they are available. Nowcasting based on radar data was investigated in detail in literature [1, 2].

However, radar coverage is not complete across the globe. There are huge areas where only the satellite observations available, so any global solution for precipitation nowcasting have to incorporate satellite data and through the following steps

1. Detect precipitation on satellite observations
2. Perform fusion of radar and satellite data
3. Train model and perform prediction

The detection is described in [3, 4] and approaches proposed for prediction step include extrapolation with optical flow [5] and convolutional LSTM networks [1]. In this paper, we focus on the fusion step of this pipeline that was not explored before.

The fusion of radar and satellite observations is a challenging task, as two data domains differ both spatially (each radar detects precipitation in the range of 200 km around its location) and temporally (radars and satellites provide observations in different moments of time, with different cadence). We address temporal fusion by simple interpolation algorithm using optical flow, and spacial fusion (image blending) by alpha-blending and special modification of inpainting algorithm with a soft mask to correct artifacts.

In the simplest case of spatial fusion, the satellite data is replaced by radar data when the latter is available. As both radar and satellite have errors, this procedure produces sharp circular seams. In case of early fusion (fusion before prediction), the fixed in space seams will confuse the nowcasting model, which is supposed to estimate and predict the movement of precipitation fields. In case of late fusion, the seams will be left in the output map and confuse the users.

Inpainting networks allow setting mask explicitly. Starting from neural network with partial convolutions [6], we improve inpaint quality by developing a new mechanism of mask updates.

In summary, we make the following contributions:

- We develop a global nowcasting using radar and satellite data.
- We develop a soft mask inpainting which outperform binary mask inpainting.
- We show that it will better variant for certain types of tasks like large area recovery and image blending.

2 Related Work

Image blending is a well-developed topic of classical computer vision. Smooth transition between composited images can be guided by the pyramid image representation [7]. Optimal seam methods [8, 9] enhance the visual quality of the composited image by cutting at optimal places, and Poisson image editing [10] corrects different exposure or tint of the composited images. Recently, GANs were applied to the problem [11], allowing to make non-local photorealistic changes.

While our task essentially is the one of image blending, several factors stop us from adopting these methods: in our case, seam position is predefined by the radar range, there is no need for exposure correction, and it is desirable to make only the local changes.

Image inpainting, the process of reconstructing lost or corrupted regions of the image, can also be used for blending: compose two images together, declare the region near the seam to be corrupted, then inpaint. Multiple approaches for image inpainting were developed in the pre-CNN era of computer vision (i.e., reconstruction of isophotes [12] or Patch-Match algorithm [13]), but the modern approaches rely on neural networks, their ability to learn the shapes and textures of the objects in the training set and reproduce them to cover corrupted parts of the image. These approaches commonly use GANs with rectangular binary masks to cover the 'corrupted' parts of the image [14, 15]. Alternatively, [6] propose a CNN architecture (inspired by UNet [16] model) with the partial convolution layers and combination of losses which outperform [14, 15].

The main idea behind partial convolution from [6] is the following. Let $\mathbf{W}$ be the convolutional weights and b the corresponding bias. $\mathbf{X}$ are the pixels values for the current convolution window and $\mathbf{M}$ is the corresponding binary mask. The partial convolution at every location is expressed as:

$$x' = \begin{cases} \frac{1}{\text{sum}(\mathbf{M})} \mathbf{W}^T (\mathbf{X} \odot \mathbf{M}) + b, & \text{if } \text{sum}(\mathbf{M}) > 0 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where $\odot$ is an element-wise multiplication. After each layer with the partial convolution, the mask values of the next layer m' are updated:

$$m' = \begin{cases} 1, & \text{if } \text{sum}(\mathbf{M}) > 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

3 Method

3.1 Temporal fusion

The first problem we encounter during mixing satellite and radar data is a different time resolution. While most radars provide images once every 10 minutes, the time resolution for the satellite imagery (Meteosat-8) is 15 minutes.

Most of the temporal change in the data is associated with precipitation fields moving in the wind, so the optical flow based algorithm comes up as a natural solution to produce intermediate frames required to link radar and satellite data. We use a simplified dense optical-flow based frame rate conversion algorithm [17] to generate synthetic satellite data at 10-minute intervals.

3.2 Soft-mask inpainting

We extend partial convolutions [6] by allowing non-binary values in the mask, interpreted as the confidence levels of the corresponding pixels in a source image. With this modification, the inpainting algorithm incorporates some information from the corrupted part of the image and successfully fills larger regions.

In case of smooth masking, Equation 2 changes the following way:

$$m' = \begin{cases} 1, & \text{if } \max(\mathbf{M}) = 1 \\ m, & \text{otherwise} \end{cases} \quad (3)$$

where m is a correspondent element of the mask from the previous layer.

3.3 Soft-mask preprocessing for blending

Since our inpainting model uses information from the area of blending, we preprocess the data to allow the network using both sources. During inference, we use simple alpha-blending to compose two images I_1 and I_2 with smooth transition between them

$$I_{\text{inference}} = \alpha \odot I_1 + (1 - \alpha) \odot I_2 \quad (4)$$

where values of the map α define this transition, linearly changing from 0 to 1 in the blending area. Note that ground truth can't be obtained for this scheme, so we need another approach for training.

During training, we use the satellite data I as the ground truth, and apply noise N to the simulated blending area to generate the training data I_{training}

$$I_{\text{training}} = |1 - 2\alpha| \odot I + (1 - |1 - 2\alpha|) \odot N \quad (5)$$

In this case, α is equal to zero in the middle of blending area, equal to one outside the blending area, and changes linearly between these two points.

4 Experiments

4.1 Comparison with binary mask inpainting

In this section, we perform experiments on a different dataset to demonstrate the validity of our approach with soft masking.

We use CelebA-HQ [18, 19] dataset for training and testing. For CelebA-HQ, we randomly partition into 27K images for training and 3K images for testing. All images scaled and cropped to 256x256 pixels. For the training process, we add noise to the images in these datasets using Equation 5 and try to predict source images. We use quick-draw [20] as a shape of the mask.

We compare binary inpainting network and two soft-mask networks: one with a constant mask filled with value 0.5 (semi-transparent mask) and another with mask $M = \alpha$, where α is taken directly from Equation 5. We compare these networks using PSNR and SSIM metrics and present the results in Table 1, demonstrating improvement of both variants of soft masking compared to binary masking.

Table 1: PSNR and SSIM metrics for inpainting networks trained on Celeba-HQ

Inpainting	PSNR	SSIM
Binary	32.6	0.9765
Semi-transparent	35.3	0.9798
α	35.9	0.9822

The best performance is achieved in the case of $M = \alpha$, but in context of the task this is data leak. However, the network with semi-transparent mask have almost the same metrics, demonstrating that

the main improvement is caused by utilization of information from the masked region of the original image, and not from the leak of parameters of data synthesis through α . Thus, we can use soft-mask inpainting with any reasonable α instead of exact values and still have an improvement in quality compared to binary masking.

4.2 Radars and satellite data fusion

Figure 1 shows results of the same model on our target task, satellite and radar data fusion. It can be seen that our approach successfully hides the border between two types of data.




Figure 1: Results of data fusion in a spatial domain. The first panel demonstrates radar data (notice the partial coverage), the second panel shows satellite observations, third - results of the composition without any form of blending, with a sharp edge between two types of data. In the fourth column, the edge is smoothed by alpha blending, and in the last column, it is further cleaned by our inpainting procedure

5 Conclusion

We presented a pipeline for spatiotemporal fusion of precipitation data. Our pipeline uses simple techniques, such as optical flow based interpolation and alpha-blending, as well as a novel approach of image inpainting with soft masking. We have demonstrated an advantage of soft masking in the case then useful information can be recovered from the corrupted region. The proposed pipeline allows expanding precipitation nowcasting from the list of selected locations covered with radars (USA, Hong Kong, Europe), to the rest of the world.

References

- [1] Xingjian Shi, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, 2015.
- [2] Xingjian Shi, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Deep learning for precipitation nowcasting: A benchmark and a new model. In *Advances in Neural Information Processing Systems*, 2017.
- [3] Hanna Meyer, Meike Kühnlein, Tim Appelhans, and Thomas Nauss. Comparison of four machine learning algorithms for their applicability in satellite-based optical rainfall retrievals. *Atmospheric Research*, 2016.
- [4] Yumeng Tao, Xiaogang Gao, Alexander Ihler, Soroosh Sorooshian, and Kuolin Hsu. Precipitation identification with bispectral satellite information using deep learning approaches. *Journal of Hydrometeorology*, 2017.
- [5] Neill EH Bowler, Clive E Pierce, and Alan Seed. Development of a precipitation nowcasting algorithm based upon optical flow techniques. *Journal of Hydrology*, 2004.
- [6] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. *arXiv preprint arXiv:1804.07723*, 2018.
- [7] Edward H Adelson, Charles H Anderson, James R Bergen, Peter J Burt, and Joan M Ogden. Pyramid methods in image processing. *RCA engineer*, 1984.
- [8] David L Milgram. Computer methods for creating photomosaics. *IEEE Transactions on Computers*, 1975.
- [9] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001.
- [10] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Transactions on graphics (TOG)*, 2003.
- [11] Huikai Wu, Shuai Zheng, Junge Zhang, and Kaiqi Huang. GP-GAN: Towards realistic high-resolution image blending. *arXiv preprint arXiv:1703.07195*, 2017.
- [12] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000.
- [13] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 2009.
- [14] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (TOG)*, 2017.
- [15] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. *arXiv preprint*, 2018.
- [16] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015.
- [17] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International Journal of Computer Vision*, 2011.
- [18] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality. *Stability, and Variation. arXiv preprint*, 2017.
- [20] David Ha and Douglas Eck. A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*, 2017.