Automated Quality Control Scheme for GPM Satellite Precipitation Products

Enter authors here: Jackson Tan[1,2], George J. Huffman[1], and Yi Song[1]

[1]NASA Goddard Space Flight Center, Greenbelt, MD, USA

[2]University of Maryland, Baltimore County, Baltimore, MD, USA

Corresponding author: Jackson Tan (jackson.tan@nasa.gov)

Key Points:

- A machine learning scheme called SPEEDe can detect artifacts in satellite precipitation fields based on an aggregate grid-wide score.
- SPEEDe identifies bad data by modifying anomalous inputs substantially to produce realistic-looking outputs, leading to a large score.
- SPEEDe is being applied operationally in GPM data production for quality control.

Abstract

The constellation approach underpinning precipitation products such as the Integrated Multi-satellitE Retrievals for GPM (IMERG) is key to achieving high resolution, but the use of data from multiple sources can unintentionally incorporate instrumental artifacts. Here, we introduce a machine learning–based anomaly detection scheme called SPEEDe, which processes a two-dimensional precipitation field into a re-estimated precipitation field that can be compared with the input. Large differences identify IMERG fields with bad orbit data, separating most of the bad cases from the good cases. When modified to process the passive microwave inputs, SPEEDe can pick out orbits with bad data, enabling quality control on these IMERG inputs. SPEEDe works by producing a locally realistic-looking precipitation field when given unphysical data, which results in a larger-than-normal difference between the input and the output. SPEEDe is implemented as an automated quality control for GPM precipitation products.

Plain Language Summary

A machine learning–based scheme, SPEEDe, has been developed to identify bad data in satellite precipitation products. This scheme is able to pick out precipitation maps that have unrealistic patterns better than a conventional approach based on the distribution of values. It can be run as the precipitation data are being produced, thus allowing for automated quality control. Avoiding manual intervention is important in computing quality-controlled precipitation data in near-real time for applications such as flood monitoring that require highly reliable data with minimal delay.

1 Introduction

A wide range of scientific studies and societal applications relies on high resolution global satellite precipitation products such as the Integrated Multi-satellitE Retrievals for GPM (IMERG) product by the U.S. Global Precipitation Measurement (GPM) Science Team (Hou et al., 2014; George J. Huffman et al., 2020; Kirschbaum et al., 2017; Skofronick-Jackson et al., 2017, 2018). IMERG uses retrievals from passive microwave (PMW) sensors on board the GPM constellation of low-Earth orbit satellites as the backbone of its precipitation estimates. To achieve sufficiently frequent sampling, the GPM constellation incorporates satellites of opportunity, which are satellites operated by other agencies that carry sensors capable of observing precipitation. However, the diverse sources of satellite data result in varying data quality.

Sporadic instances of poor data quality have been a longstanding and highly frustrating issue in IMERG and the Goddard Profiling (GPROF) algorithm (Kummerow et al., 2015; Randel et al., 2020) that provides the PMW precipitation retrievals that are ingested into IMERG. The situation is particularly prominent in the new IMERG V07A (George J. Huffman et al., 2023) due to a confluence of factors that exacerbated the frequency of bad data. Though the frequency is still low at about 0.2% of the time (as identified by the scheme developed here), it can have an appreciable impact. For example, Figure 1a shows a half-hour snapshot of the precipitation field from an unusually egregious case. It is obvious that the precipitation field contains artifacts, whose pattern indicates bad data from a polar-orbiting PMW sensor. In this particular case, the issue was tracked down to the SSMIS sensor onboard the DMSP-F16 satellite. These bad data affect not only the half-hour during which it occurs, but also the entire month due to its "contamination" of the gauge adjustment ratio and in some cases neighboring months due to the

rolling calibration of the infrared precipitation and the Kalman filter statistics employed in IMERG.

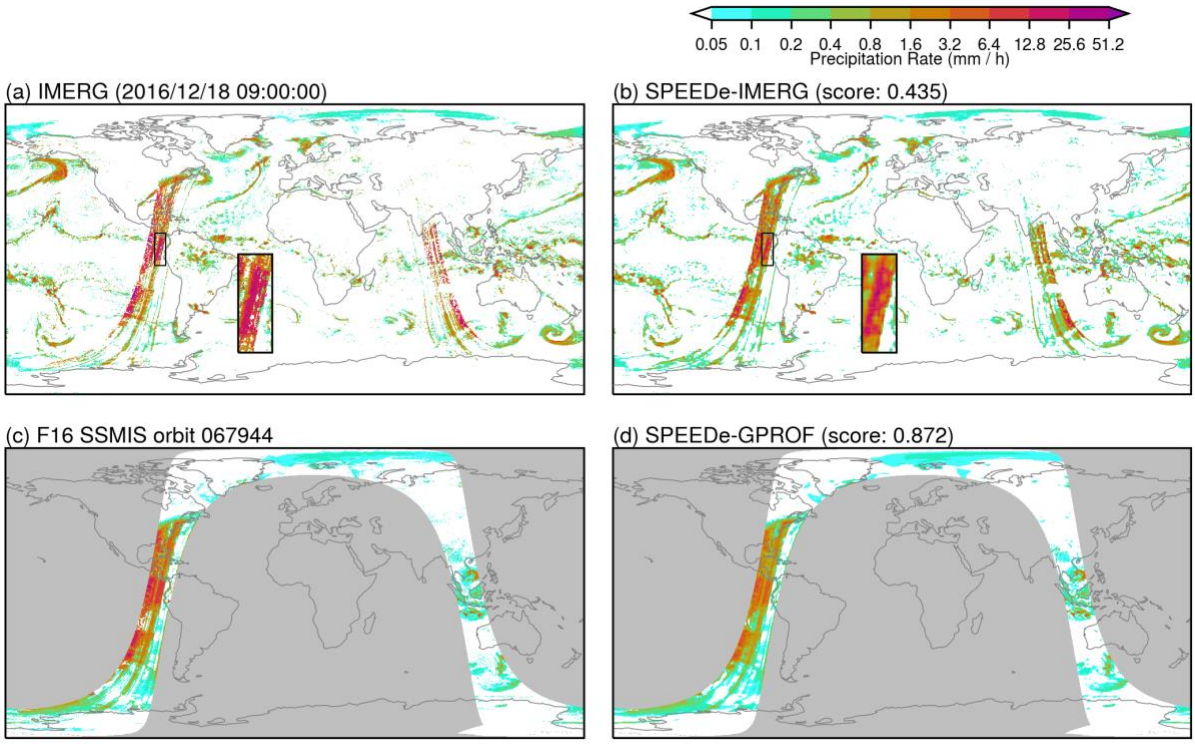


Figure 1. (a) A half-hour IMERG V07A snapshot on 2016/12/18 09:00–09:29 showing the presence of a bad orbit. (b) The SPEEDe-IMERG output corresponding to the snapshot in (a). The black boxes and the inset panels highlight a region for further discussion in the text. (c) A bad F16 SSMIS precipitation field associated with the bad IMERG case. (d) The SPEEDe-GPROF output corresponding to the snapshot in (c).

Such bad data or unphysical features—which can manifest as streaks, speckles, excessive patches, or unrealistic gradients of precipitation—have a myriad of causes, from instrument malfunction to geolocation error to calibration issues, making them hard to detect from their brightness temperature fields. The most reliable method at present for detecting such bad data is to visually inspect every precipitation field, but this is laborious and wholly impractical, and an automated detection scheme is highly desirable. Since such artifacts are visually obvious, we use a machine learning–based image recognition algorithm, specifically a convolutional autoencoder scheme, for detecting anomalies in IMERG and GPROF. A convolutional autoencoder takes a two-dimensional input, reduces its size using convolution and max-pooling layers, then reverses the process using deconvolution and upscaling layers, giving a two-dimensional output with the same size as the input. We then compare the output to the input—effectively assuming the input as "truth"—with a loss function that quantifies the difference between the two. For anomaly detection, the convolutional autoencoder is trained only with data confirmed as "good" (in the sense of an absence of satellite artifacts) so that the model parameters encode information about "good" precipitation patterns. Therefore, when bad data are encountered, such as in Figure 1a, the convolutional autoencoder will produce an output that differs from the input substantially,

resulting in a large absolute difference. We call this convolutional autoencoder scheme Satellite Precipitation Estimate Error Detector (SPEEDe).

2 Data

2.1 GPROF

GPROF is the primary algorithm for retrieving Level 2 precipitation rate estimates at the resolution of individual sensor footprints from PMW brightness temperatures in GPM (Kummerow et al., 2001, 2011; Randel et al., 2020). GPROF is a physically-based parametric algorithm that estimates the surface precipitation rate using a conditional Bayesian retrieval of precipitation by comparing the observed brightness temperature to a priori databases according to the surface class, surface temperature, and total column water vapor (Kummerow et al., 2015). These a priori databases generally match brightness temperatures with associated surface precipitation rates and hydrometeor profiles from the Ku-band Dual-frequency Precipitation Radar and the Combined Radar-Radiometer Algorithm product (Grecu & Olson, 2020), with exceptions made for snow-covered surfaces.

GPROF is applied to most of the PMW sensors in the GPM constellation, which contains satellites operated by U.S. and international agencies with sensors capable of observing precipitation (Table S1). Except for GPM Microwave Imager (GMI) and TRMM Microwave Imager (TMI), data from the various sensors are received by different processing centers before being transmitted in near-real-time to the GPM Precipitation Processing System (PPS). Such a constellation approach allows the GPM mission, and thus IMERG, to leverage existing assets to achieve a reasonable revisit rate over the entire record. For more information about the GPM Constellation, see https://gpm.nasa.gov/missions/GPM/constellation.

2.2 IMERG

IMERG is the gridded satellite precipitation product from the U.S. GPM Science Team that merges precipitation observations from the GPM constellation (George J. Huffman et al., 2020, 2023). (The counterpart to IMERG from the Japanese GPM Science Team is the Global Satellite Mapping of Precipitation, or GSMaP.) IMERG grids the GPROF precipitation estimates separately for each sensor to 0.1° in half-hour intervals. The gridded PMW estimates are then intercalibrated in a two-step process: first by a seasonally varying climatological adjustment to TMI/GMI so the sensors have a consistent distribution, then second by applying calibration factors computed from TMI/GMI matched to the Ku-band Combined Radar and Radiometer product, which contains further climatological adjustments over oceans based on the Global Precipitation Climatology Project V3.2 monthly product. These PMW estimates from different sensors are then merged onto a common grid. The gridded PMW estimates are propagated forward and backward in time using motion vectors computed from several precipitation-related variables in numerical models (Tan et al., 2019). Meanwhile, PMW-calibrated infrared precipitation is computed using a scheme adapted from Nguyen et al. (2020) and applied to the Climate Prediction Center (CPC) Merged IR brightness temperature (Janowiak et al., 2017). The forward-propagated PMW precipitation, backward-propagated PMW precipitation, and infrared precipitation fields are then merged via a Kalman filter approach following Joyce and Xie (2011), after which the Scheme for the Histogram Adjustment of Ranked Precipitation Estimate in the Neighborhood (Tan et al., 2021) is applied to the Kalman filter output to mitigate the deleterious effects of averaging on the precipitation distribution (Rajagopal et al., 2021). The

resulting "best" satellite-only, globally complete precipitation field is then adjusted with the Full and Monitoring gauge analysis products from the Global Precipitation Climatology Centre (Schneider et al., 2015).

IMERG has three Runs—Early, Late, and Final—to accommodate different user requirements for latency and accuracy. The Early and Late Runs are computed in near-real time, while the Final Run is at a 3.5-month latency and is intended for research applications. SPEEDe is equally applicable for all three, so here we just focus on the IMERG V07A Final Run.

2.3 Data Preparation for SPEEDe

Any missing values in the precipitation fields are filled with zeros. Next, following data scaling recommendations for machine learning models, the precipitation rates, p, are "normalized" by applying a log-transformation of $\log_{10}(p + 1)$. This reduces the skewness of the precipitation rate distribution to a reasonable range (mostly below a value of 2, which is equivalent to 99 mm / h) while keeping it nonnegative. This log-transformation is perhaps unnecessary, since the typical justification is to equalize the gradients for multiple inputs that have disparate ranges, while SPEEDe only has precipitation rates as a single input. Nevertheless, preliminary testing (see Text S1) suggests that a log-transformation marginally improves the detection skill, or at least is not detrimental, so we opted to retain it.

3 Convolutional Autoencoder Architecture

SPEEDe takes the precipitation field as an input, reduces its size twice, then increases its size twice, before producing an output array with the same size as the input. The reduction in size is achieved using a 2D convolutional layer (with zero-padding) followed by a max-pooling layer, while the increase in size is achieved through the use of a 2D upsampling layer followed by a 2D transposed convolutional (or deconvolutional) layer. Figure 2 shows the architecture of SPEEDe as applied to IMERG, implemented in Python version 3.10.10 with the TensorFlow module version 2.11.0 (TensorFlow Developers, 2023). With a fixed 3600 × 1800 input array corresponding to the number of longitude and latitude grid boxes in IMERG, the first set of convolution and max-pooling operations reduces the input by a factor of 5 with 64 filters and the second set of operations further reduces the layers by a factor of 3 with 32 filters. These parameters were chosen after some preliminary testing (see Text S1). Then, the process is reversed to produce an output of 3600 × 1800. The activation function for each layer is the rectified linear unit, including for the output layer since the log-transformed precipitation rate is nonnegative. The loss function is the mean squared error, from which we define the SPEEDe score simply as the mean squared error multiplied by 100; coupled with the log-transformation, this means that,

$$\text{SPEEDe score} = \frac{1}{3600 \times 1800} \sum_{i=1}^{3600} \sum_{j=1}^{1800} \left(\log_{10}\left(p_{i,j}^{\text{in}} + 1\right) - \log_{10}\left(p_{i,j}^{\text{out}} + 1\right)\right)^2 \times 100 . \quad (1)$$

A higher score indicates a higher likelihood that the array contains artifacts. The optimizer is the Adam algorithm with the default learning rate of 0.001 (Kingma & Ba, 2014). A random selection of 500 IMERG fields from the entire record were manually confirmed to be "good" and used as inputs for training and validation with an 80%-20% split. The training was conducted over 57 epochs as determined by an early stopping algorithm (loss delta = 0.00001, patience = 5

epochs) to prevent overfitting and the model from the epoch with lowest validation loss is saved. The validation loss of the trained model is 0.00157.

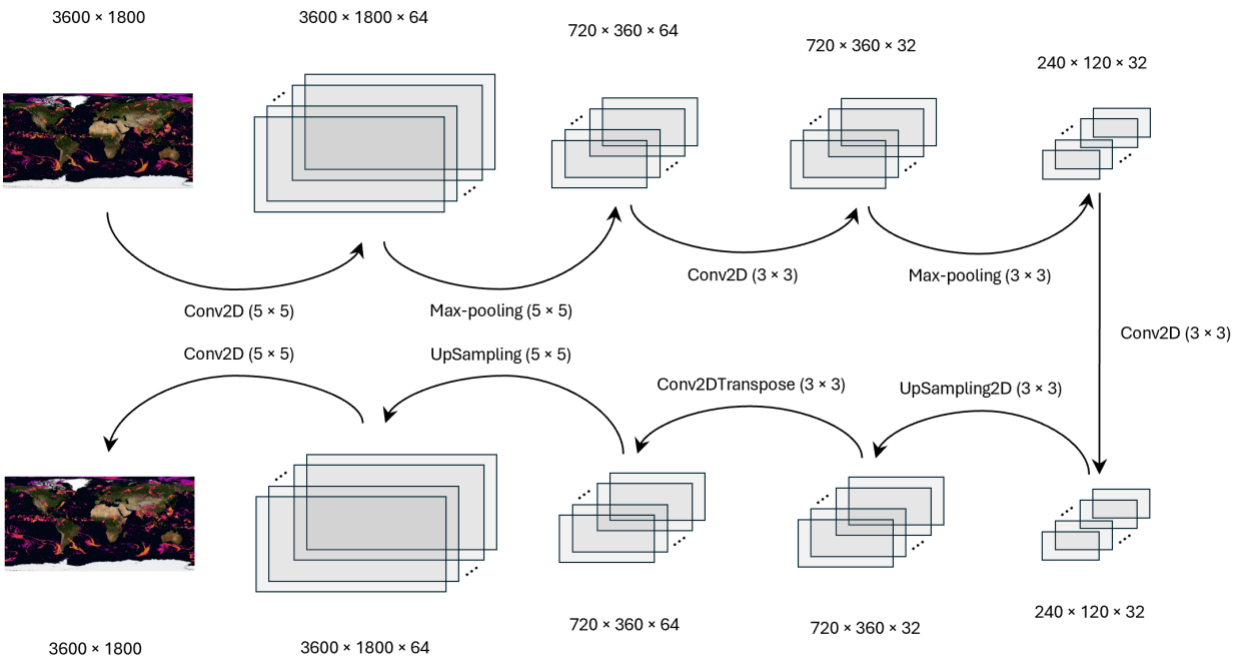


Figure 2. Architecture of SPEEDe-IMERG. The numbers in the outer rows show the array dimensions at each layer, while the text in the inner rows shows the operation and corresponding kernel size.

The architecture and parameters of SPEEDe as summarized in Figure 2 are the best of several variants explored, as listed in Text S1 and Table S2.

In order to apply SPEEDe to GPROF, some modifications to Figure 2 are required. First, instead of an array with dimensions of longitude and latitude, the precipitation array has dimensions of along-track (or number of scans in each orbit granule) and cross-track (or number of pixels in each scan). As well, since different sensors have different along-track and cross-track dimension sizes in their orbit granules, each sensor will need to have its own SPEEDe-GPROF model. Moreover, different orbit granules can have slightly different along-track sizes. Thus, for simplicity, we modified SPEEDe to reduce/increase the arrays by a factor of 10—5 × 5 followed by 2 × 2—and either truncate the inputs or pad the inputs with zeros to the appropriate assumed size for that sensor. For SSMIS, the assumed size is 3200 × 180; for MHS, the assumed size is 2280 × 90. The SPEEDe score (Eq. (1)) is modified accordingly. Second, SPEEDe was trained for each sensor using 1000 randomly selected orbits from either 2007 (TRMM-era sensors) or 2015 (GPM-era sensors). These orbits were considered "good" because bad orbits have been screened and removed using an earlier version of SPEEDe. Third, since GPROF, unlike IMERG, can occasionally have a substantial fraction of missing values in a particular ~90-min orbit, the SPEEDe score is computed over only non-missing pixels.

4 Applying SPEEDe to IMERG

4.1 Case Studies

Figure S1 shows four examples of good half-hourly precipitation fields from IMERG V07A and the corresponding SPEEDe-IMERG output. SPEEDe-IMERG is able to preserve the broad precipitation patterns while removing the fine-scale structures and noise, with an SPEEDe-IMERG score ranging from 0.127 to 0.169. Figure 1b shows the SPEEDe-IMERG output using the previous bad case as input (Figure 1a). The SPEEDe-IMERG output shows obvious influence from the bad SSMIS orbit, inheriting the swath pattern of high precipitation rates. In particular, the SPEEDe-IMERG score is 0.435, substantially higher than the range of scores in the good examples (Figure S1), thus raising the prospect of using the score as a quantitative tool for quality control. In fact, a close inspection is instructive on why the SPEEDe score is high when the input contains bad data. Looking only at the precipitation system in the region marked by the black box in Figure 1b, it is not immediately obvious that it is an artifact. In fact, viewed in isolation and without context to its spatial scale, one could conceivably interpret that pattern as a squall line or a frontal system, with a realistic gradient of light precipitation rates on the outside to intense precipitation rates at the core. We speculate that the reason SPEEDe-IMERG produced a realistic-looking precipitation field—especially with the light-to-heavy-to-light gradient—is because it is trained with "good" data. This means that, compared to the "good" case (e.g., Figure S1), SPEEDe-IMERG modified the anomalous input substantially in order to have a more realistic-looking output, resulting in a high mean squared error between the input and the output, and thus a high SPEEDe score.

4.2 Distribution of Scores

Figure 3a gives the distributions of the SPEEDe-IMERG scores for known cases of bad orbits, known cases of marginally bad orbits, as well as a random selection of 200 cases confirmed to be "good" by manual visual inspection. (Marginally bad cases are subjectively determined using the criterion that an experienced person would take more than a second to recognize a bad swath in the snapshot; see Figure S2 for examples. They usually occur in the half-hours neighboring a bad PMW orbit, as a result of the bad orbit "fading out" as it is morphed.) The most egregious cases have scores that are well above those of "good" cases; in fact, Figure 3a indicates that about 75% of the bad cases have scores higher than the maximum score of "good" cases. However, there is an appreciable overlap in scores between the "good" cases and bad cases, and more so between the "good" cases and the marginal cases, implying a limit in the effectiveness of SPEEDe-IMERG for detecting the presence of bad data. Fortunately, manual inspection of snapshots of the bad cases with low SPEEDe-IMERG scores revealed that the unphysical values tend to be modest and have limited spatial extent (Figure S2). Therefore, SPEEDe-IMERG can be effective in capturing the worst cases and many of the moderately bad cases.

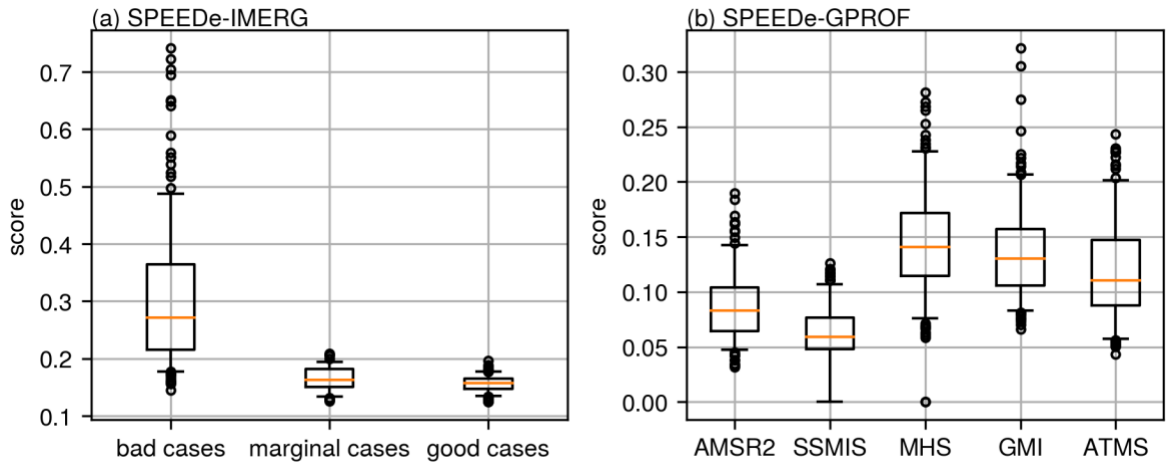


Figure 3. (a) Distributions of SPEEDe-IMERG scores for known bad cases, known marginal cases, and a selection of 200 "good" cases. (b) Distributions of SPEEDe-GPROF scores for GPM-era sensors using a random selection of 200 orbit granules in 2015. The box-whisker diagrams show the median (orange line), 25th/75th percentiles (box), 5th/95th percentiles (whiskers), and outliers beyond the whiskers (circles).

In fact, Figure 3a suggests several feasible approaches for data quality control. If an IMERG precipitation field has a SPEEDe-IMERG score above a sufficiently high threshold—e.g., 0.3—the data can be withheld for correction. This threshold will of course depend on the nature of the IMERG Run: for Final Run, this threshold should be reasonably low due to its emphasis of research quality, whereas for the Late and Early Runs, data latency may take priority with data interruptions kept to a minimum, so a higher threshold is warranted. Furthermore, a second and lower threshold can be set that would not interrupt data availability but flag it for further inspection and, if necessary, documentation.

4.3 Why Histograms of Precipitation Fields Cannot Detect Bad Orbits

There is clearly an excess of intense precipitation rates in Figure 1a, so a natural question that arises is whether the distribution of precipitation rates can indicate the presence of bad data. To explore this question, we examined how the histogram of each half-hour IMERG precipitation field differs from the climatological histogram (Text S2). Figure S3 shows the distribution of this histogram error for the same cases in Figure 3a. While the egregious cases indeed have a high mean squared error in their histograms as expected, there is an appreciable overlap between the scores of the bad cases and the scores of the good cases, particularly for marginal bad cases. The fact that good cases can have histograms that deviate substantially from climatology simply reflects the highly variable nature of precipitation, even at the global scale. It may be possible that refining how the difference in distribution is computed—e.g., seasonally and/or regionally—can induce a greater separation of the bad cases from the good cases, but experience suggests that any improvements acquired by increased complexity would come at the expense of robustness. It is clear from this experiment that identifying a simple deviation in histogram is operationally unfeasible as an anomaly detection scheme, especially considering that an overwhelming majority of the cases are good.

5 Applying SPEEDe to GPROF

While an anomaly detection scheme for IMERG is highly useful for flagging bad data in the completed product, it would be preferable to prevent bad data from being ingested into IMERG in the first place. This means extending the SPEEDe scheme to GPROF, which provides the PMW input into IMERG as discussed earlier. Figure S4 shows examples of good precipitation fields from GPROF and their corresponding SPEEDe-GPROF output. Just as with the normal case of IMERG (Figure S1), SPEEDe-GPROF generally smooths the precipitation fields, removing small-scale noise. The SPEEDe-GPROF scores are 0.050 (F16 SSMIS), 0.116 (MetOp-B MHS), 0.090 (SNPP ATMS), and 0.103 (GCOM-W1 AMSR2).

Figure 1c shows the precipitation field from the F16 SSMIS orbit granule containing bad data that was ingested into IMERG in Figure 1a, and Figure 1d shows the corresponding SPEEDe-GPROF output. First, note that the SPEEDe-GPROF score jumped to 0.872. Examining the SPEEDe-GPROF output (Figure 1d) in detail, the precipitation patterns have a more realistic-looking pattern on the local scale, depicting clearer light-to-heavy-to-light gradients just as we have observed in Figure 1b.

Figure 4 shows more examples of bad orbits and their SPEEDe-GPROF output. For the MetOp-A MHS granule (Figure 4a), Asia and Australia show unphysical precipitation patterns. SPEEDe-GPROF produced an output that has less abrupt gradients in the precipitation systems over northern and southern Australia, leading to a score of 0.392. For the Aqua AMSR-E granule (Figure 4c), there is an obvious excess of oceanic precipitation across the entire track, leading to a SPEEDe-GPROF score of 0.871. For the NOAA-17 AMSU-B granule (Figure 4e), the precipitation field contains speckles of precipitation mixed with speckles of missing values—a sign of instrument issues—with a SPEEDe-GPROF score of 0.83926. Just as with Figure 1d, the SPEEDe-GPROF outputs corresponding to these unphysical precipitation systems appear more similar to realistic precipitation features; in the case of MetOp-A MHS, the outputs resemble mesoscale convective systems. In this case the SPEEDe-GPROF score is lower than that of the bad F16 SSMIS orbit because the bad data cover a smaller area; the score both reflects the severity of the “badness” and its spatial extent.

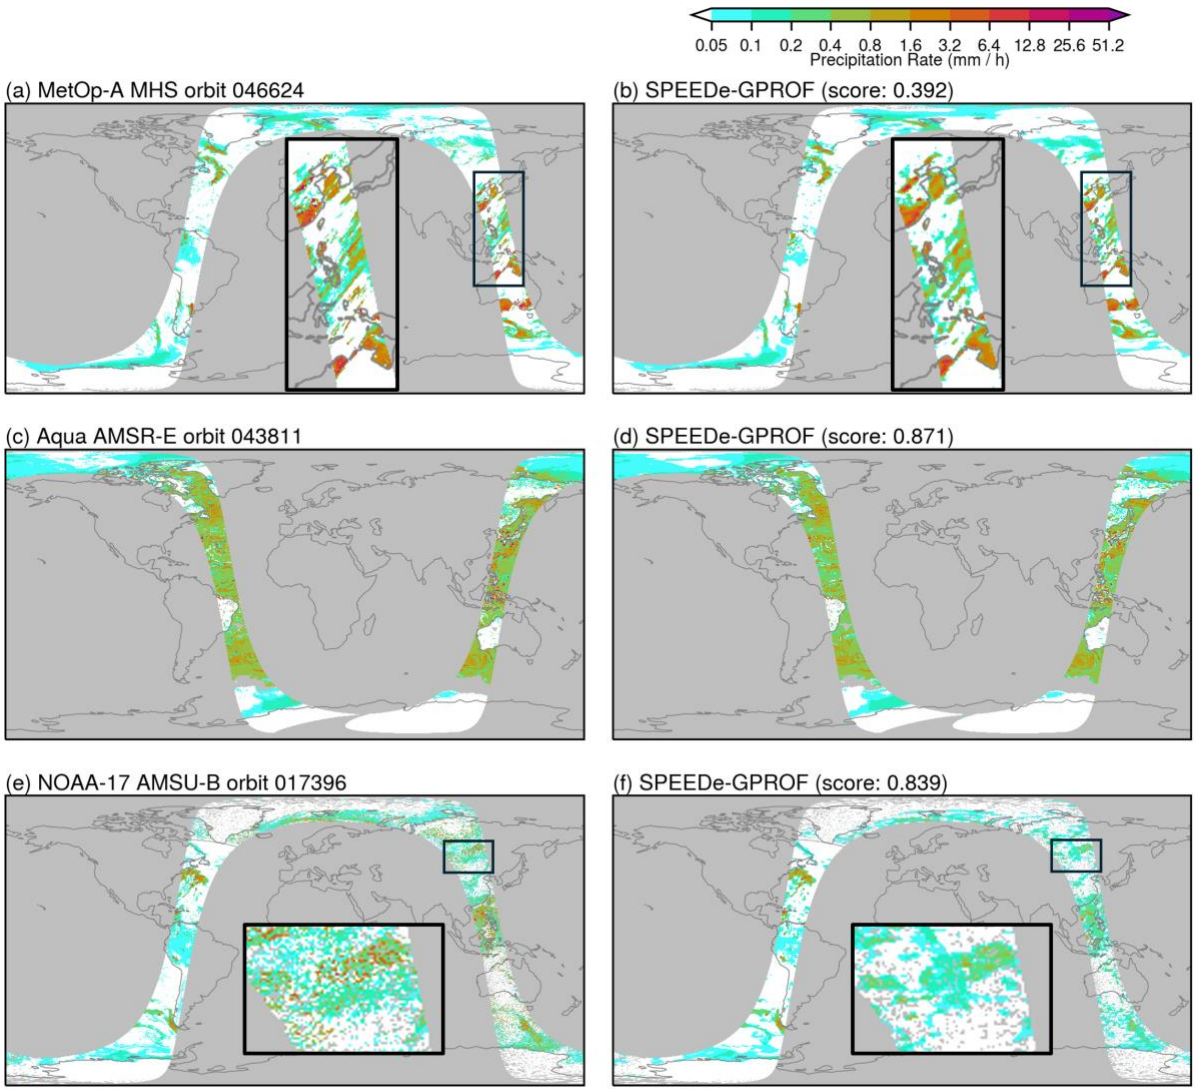


Figure 4. Examples of bad orbit granules from (a) MetOp-A MHS, (c) Aqua AMSR-E, and (e) NOAA-17 AMSU-B, with (b, d, f) their corresponding SPEEDe-GPROF outputs. The black boxes indicate inset figures that enlarge the regions with artifacts.

To provide a sense of the "good" SPEEDe-GPROF scores, Figure 3b shows the distribution of scores for the five GPM-era sensors sampled for the year of 2015. While it is perhaps not surprising that each sensor has a different range of scores. given the different array dimensions and footprint sizes, it does reveal that any score-based detection will need to account for sensor variations. As well, for orbit granules with limited coverage by non-missing values (due to, e.g., upstream screening of pixels that were flagged as bad), the score can be highly variable even if the precipitation field appears "good" (Figure S5) because the mean squared error is computed over such a small sample size. It is a matter of future research to refine SPEEDe-GPROF to mitigate these challenges.

5 Conclusions

In response to the occurrence of bad data in multi-satellite precipitation products—in particular IMERG V07A—this paper introduces a bad data identification tool called SPEEDe for quality control. Since fields with bad data are easily recognized by visual inspection, the tool adopts a machine learning approach and has a convolutional autoencoder architecture. Trained with "good" data, SPEEDe produces an output precipitation field that smooths the input precipitation field but retains realistic structures of precipitation patterns and broadly resembles the input. In particular, the output field contains gradients of light-to-heavy-to-light precipitation rates in large systems that appear physical. When fed bad data, SPEEDe produces an output that locally appears reasonable, but is systematically different than the input and thus has a higher mean squared difference between the two that can be used as a quantitative metric for identifying bad data. When implemented for IMERG, simple thresholding of the SPEEDe score is able to identify many of the bad cases and, especially, the egregious cases. SPEEDe has also been successfully adapted for GPROF orbit granules. The IMERG and GPROF products have since been updated to V07B to remove these identified bad orbit granules.

Applying SPEEDe to GPROF instead of IMERG has advantages but also introduces challenges. The advantages of SPEEDe-GPROF are (i) the ability to screen bad orbit granules before they are ingested into IMERG, (ii) a cleaner separation between bad data and normal data, and (iii) faster computational time and lower memory demand due to smaller precipitation arrays. The challenges are (i) the varying sizes of the PMW sensor data fields—thus necessitating a model for each sensor—even from the same sensor, (ii) different ranges of scores for the different sensors, and (iii) the variable nature of the scores when the orbit granule has substantial missing data. Nevertheless, SPEEDe demonstrates promise as a quality control tool for both products and is being considered for operational implementation by the respective algorithm teams. As of the time of writing, SPEEDe is being run operationally on GPROF files, with orbits whose scores exceed sensor-dependent thresholds excluded from IMERG in near-real-time processing and flagged for manual review in post-real-time processing. It is likely that SPEEDe should be used in combination with other quality control tools for maximum effectiveness.

Acknowledgments

The authors thank the IMERG and PPS teams for detailed discussions during the development of SPEEDe. All authors are supported with NASA Precipitation Measurement Missions funding. IMERG was provided by the NASA Goddard Space Flight Center's IMERG and PPS teams, which develop and compute IMERG as a contribution to GPM, and archived at PPS and the NASA Goddard Earth Science Data and Information Services Center (GES DISC).

Open Research

IMERG and GPROF are publicly available at GES DISC (G. J. Huffman et al., 2023) and PPS (https://gpm.nasa.gov/data/directory). Codes for SPEEDe and the analysis for this study are available at https://github.com/JacksonTanBS/SPEEDe.

References

Grecu, M., & Olson, W. S. (2020). Precipitation Retrievals from Satellite Combined Radar and Radiometer Observations. In V. Levizzani, C. Kidd, D. B. Kirschbaum, C. D.

Kummerow, K. Nakamura, & F. J. Turk (Eds.), *Satellite Precipitation Measurement* (Vol. 67, pp. 231–248). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-24568-9_14

Hou, A. Y., Kakar, R. K., Neeck, S., Azarbarzin, A. A., Kummerow, C. D., Kojima, M., et al. (2014). The Global Precipitation Measurement Mission. *Bulletin of the American Meteorological Society*, *95*(5), 701–722. https://doi.org/10.1175/BAMS-D-13-00164.1

Huffman, G. J., Stocker, E. F., Bolvin, D. T., Nelkin, E. J., & Tan, J. (2023). GPM IMERG Final Precipitation L3 Half Hourly 0.1 degree x 0.1 degree V07 [Data set]. NASA Goddard Earth Sciences Data and Information Services Center. https://doi.org/10.5067/GPM/IMERG/3B-HH/07

Huffman, George J., Bolvin, D. T., Braithwaite, D., Hsu, K.-L., Joyce, R. J., Kidd, C., et al. (2020). Integrated Multi-satellite Retrievals for the Global Precipitation Measurement (GPM) Mission (IMERG). In V. Levizzani, C. Kidd, D. B. Kirschbaum, C. D. Kummerow, K. Nakamura, & F. J. Turk (Eds.), *Satellite Precipitation Measurement* (Vol. 67, pp. 343–353). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-24568-9_19

Huffman, George J., Bolvin, D. T., Joyce, R., Nelkin, E. J., Tan, J., Braithwaite, D., et al. (2023). *Algorithm Theoretical Basis Document (ATBD). NASA Global Precipitation Measurement (GPM) Integrated Multi-satellitE Retrievals for GPM (IMERG) Version 07* (Algorithm Theoretical Basis Document) (p. 52). NASA. Retrieved from https://gpm.nasa.gov/resources/documents/imerg-v07-atbd

Janowiak, J., Joyce, R., & Xie, P. (2017). NCEP/CPC L3 Half Hourly 4km Global (60S - 60N) Merged IR V1 [Data set]. NASA Goddard Earth Sciences Data and Information Services Center. https://doi.org/10.5067/p4hzb9n27eku

Joyce, R. J., & Xie, P. (2011). Kalman Filter–Based CMORPH. *Journal of Hydrometeorology*, *12*(6), 1547–1563. https://doi.org/10.1175/JHM-D-11-022.1

Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization (Version 9). https://doi.org/10.48550/ARXIV.1412.6980

Kirschbaum, D. B., Huffman, G. J., Adler, R. F., Braun, S., Garrett, K., Jones, E., et al. (2017). NASA's Remotely Sensed Precipitation: A Reservoir for Applications Users. *Bulletin of the American Meteorological Society*, *98*(6), 1169–1184. https://doi.org/10.1175/BAMS-D-15-00296.1

Kummerow, C. D., Hong, Y., Olson, W. S., Yang, S., Adler, R. F., McCollum, J., et al. (2001). The Evolution of the Goddard Profiling Algorithm (GPROF) for Rainfall Estimation from Passive Microwave Sensors. *Journal of Applied Meteorology*, *40*(11), 1801–1820. https://doi.org/10.1175/1520-0450(2001)040<1801:TEOTGP>2.0.CO;2

Kummerow, C. D., Ringerud, S., Crook, J., Randel, D., & Berg, W. (2011). An Observationally Generated A Priori Database for Microwave Rainfall Retrievals. *Journal of Atmospheric and Oceanic Technology*, *28*(2), 113–130. https://doi.org/10.1175/2010JTECHA1468.1

Kummerow, C. D., Randel, D. L., Kulie, M., Wang, N.-Y., Ferraro, R., Joseph Munchak, S., & Petkovic, V. (2015). The Evolution of the Goddard Profiling Algorithm to a Fully Parametric Scheme. *Journal of Atmospheric and Oceanic Technology*, *32*(12), 2265–2280. https://doi.org/10.1175/JTECH-D-15-0039.1

Nguyen, P., Ombadi, M., Gorooh, V. A., Shearer, E. J., Sadeghi, M., Sorooshian, S., et al. (2020). PERSIANN Dynamic Infrared–Rain Rate (PDIR-Now): A Near-Real-Time,

Quasi-Global Satellite Precipitation Dataset. *Journal of Hydrometeorology*, *21*(12), 2893–2906. https://doi.org/10.1175/JHM-D-20-0177.1

Rajagopal, M., Zipser, E., Huffman, G., Russell, J., & Tan, J. (2021). Comparisons of IMERG Version 06 Precipitation At and Between Passive Microwave Overpasses in the Tropics. *Journal of Hydrometeorology*, *22*(8), 2117–2130. https://doi.org/10.1175/JHM-D-20-0226.1

Randel, D. L., Kummerow, C. D., & Ringerud, S. (2020). The Goddard Profiling (GPROF) Precipitation Retrieval Algorithm. In V. Levizzani, C. Kidd, D. B. Kirschbaum, C. D. Kummerow, K. Nakamura, & F. J. Turk (Eds.), *Satellite Precipitation Measurement* (Vol. 67, pp. 141–152). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-24568-9_8

Schneider, U., Becker, A., Finger, P., Meyer-Christoffer, A., Rudolf, B., & Ziese, M. (2015). GPCC Full Data Reanalysis Version 7.0 at 1.0°: Monthly Land-Surface Precipitation from Rain-Gauges built on GTS-based and Historic Data: Gridded Monthly Totals [Data set]. Global Precipitation Climatology Centre (GPCC) at Deutscher Wetterdienst. https://doi.org/10.5676/dwd_gpcc/fd_m_v7_100

Skofronick-Jackson, G., Petersen, W. A., Berg, W., Kidd, C., Stocker, E. F., Kirschbaum, D. B., et al. (2017). The Global Precipitation Measurement (GPM) Mission for Science and Society. *Bulletin of the American Meteorological Society*, *98*(8), 1679–1695. https://doi.org/10.1175/BAMS-D-15-00306.1

Skofronick-Jackson, G., Kirschbaum, D., Petersen, W., Huffman, G., Kidd, C., Stocker, E., & Kakar, R. (2018). The Global Precipitation Measurement (GPM) mission's scientific achievements and societal contributions: reviewing four years of advanced rain and snow observations. *Quarterly Journal of the Royal Meteorological Society*, *144*, 27–48. https://doi.org/10.1002/qj.3313

Tan, J., Huffman, G. J., Bolvin, D. T., & Nelkin, E. J. (2019). IMERG V06: Changes to the Morphing Algorithm. *Journal of Atmospheric and Oceanic Technology*, *36*(12), 2471–2482. https://doi.org/10.1175/JTECH-D-19-0114.1

Tan, J., Huffman, G. J., Bolvin, D. T., Nelkin, E. J., & Rajagopal, M. (2021). SHARPEN: A Scheme to Restore the Distribution of Averaged Precipitation Fields. *Journal of Hydrometeorology*, *22*(8), 2105–2116. https://doi.org/10.1175/JHM-D-20-0225.1

TensorFlow Developers. (2023, November 14). TensorFlow (Version v2.15.0). Zenodo. https://doi.org/10.5281/ZENODO.4724125