Learning evapotranspiration dataset corrections from water cycle closure supervision

Anonymous Author(s)
Affiliation
Address
email

Abstract

Evapotranspiration (ET) is one of the most uncertain components of the global water cycle. Improving global ET estimates is needed to better our understanding of the global water cycle so as to forecast the consequences of climate change on the future of global water resource distribution. This work presents a methodology to derive monthly corrections of global ET datasets at 0.25 degree resolution. We use ML to generalize sparse catchment-level water cycle closure residual information to global and dense pixel-level residuals. Our model takes a probabilistic view on ET datasets and their correction that we use to regress catchment-level residuals using a sum-aggregated supervision. Using four global ET datasets, we show that our learned model has learned ET corrections that accurately generalize its water cycle-closure results to unseen catchments.

1 Introduction

In the context of Climate Change (CC), the demand on water resources is increasing as both flood and drought related damages increase. Human activities are known to impact the global water cycle. However observational uncertainties limit extreme hazard forecast capability and render human contribution to CC trend estimates very challenging in the context of high natural climate variability [1]. The main evidence of observational uncertainties and discrepancy in monitoring of the hydrosphere is that the water cycle is still not closed [2]. The water cycle is modelled through four components:, precipitation (P), evapotranspiration (ET), river discharge (R) and water storage differential (dS). Closing the water cycle refers to accurately quantifying each of these components at a given spatial and temporal resolution so that they sum to zero on all spatio-temporal locations:

$$P - ET - R + dS = 0 \tag{1}$$

Among the hydrosphere component, ET remains one of the most uncertain and elusive components of Earth's water balance: it is a difficult physical process to sense as it cannot be observed directly from space, and its field measurement via eddy-covariance method raise limited spatial representatives [3]. Improving ET is needed for an advanced closure of the water cycle at regional to local scales. Such improvement would translate into an improved capability to monitor and forecast extreme hazard and to attribute to human activity a more accurate part of global change. One the one hand previous studies have succeed in optimizing ET using the water cycle closure as a constrain [11,12]. but such attempts was limited to global or catchment scale and raise limitation in generalizing ET correction at higher scale [13]. Very recently, ML have leverage the use of hydroclimatic variables and large catchment database for inferring pixelwise correction on precipitation atlas [14]. In this work, we propose a method that corrects existing global ET datasets so as to better close the water

Submitted to Tackling Climate Change with Machine Learning: workshop at NeurIPS 2022.

cycle. Doing so presents two challenges: First, the water cycle constraint includes the R component, which is only defined at the catchment scale, while we seek evaporation corrections at the pixel scale. Second, R measurements are only sparsely available, so water cycle closure errors are only available locally in space and time, while we aim to provide dense corrections to global ET datasets. Our solution to both challenges is to use Machine Learning (ML) to generalize ET corrections from sparse catchment levels to dense pixel level. We train a pixel-wise model to regress ET corrections from globally available climatic indices, which allows us to generalize the learned corrections to a dense pixel-wise resolution. Due to the nature of the R measurements, the supervision signal is defined at the catchment level, so we train our model using a sum-aggregated supervision in which we regress the sum of model outputs over catchment pixels to the catchment-level label. Our loss is defined using a Maximum A Posteriori (MAP) formulation, in which we use prior knowledge on ET uncertainties to guide the supervision.We evaluate the ability of learned corrections to close the water cycle on unseen catchments and report consistent improvements across 4 global ET datasets. We also compare ET corrections to in-situ measurements and report improvements on 3 out of 4 datasets.

2 Dataset

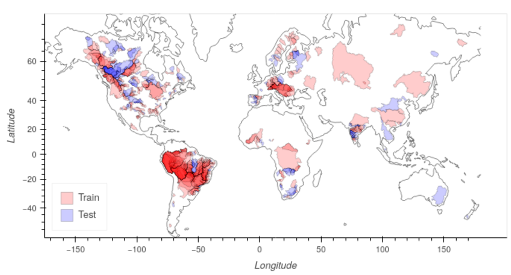

Figure 1: Illustration of our dataset's catchment locations and split.

Figure 1 illustrates the location of training and test catchments for which we have gathered data. We investigate four different global ET datasets, for each of which we learn and evaluate corrections. Each dataset estimates ET using different methodology and thus showcase different error patterns. These datasets are Global Land Evaporation Amsterdam Model [6] version va.3 and vb.3, the Penman-Monteith-Leuning (PML) estimate [7] and the reanalysis ERA5 [8]. Our dataset contains 663 catchments (C) and covers a time period (T) of 192 months ranging from January 2000 to December 2015, although many months of data are missing for most catchments. In total, our dataset consists of 71654 monthly catchment-level data points for which all components are available. Our model is defined at a spatial resolution of 0.25 degrees, and a one month time resolution. It processes $D = 7$ dimension input feature vectors, representing ET, P, dS, and four climatic indices representing vegetation cover (LAI [8], NDVI [9]), soil moisture [8], and surface water availability (P-E from [8]). Given an ET dataset E_i, for each catchment $c \in C$ and for each month $t \in T$, the ground-truth correction y is given as the difference $y(c,t) = E(c,t) - E_i(c,t)$, where E represent the best catchment-level water cycle closure corrected estimate we have. E was computed using the Optimal Interpolation (OI) method proposed in [10], which accounts for uncertainty estimates of all water cycle component and has been shown to improve the catchment-level estimates of all components [11]. In addition, we used a catchment-level simple-weighting aggregation [10] of all four ET datasets as the input ET component to the OI. For each catchment and month index pairs (c, t), our dataset thus provides an input output pair $(X(c,t), y(c,t))$. As each catchment covers many pixels, X represents a set of input feature vectors $X(c,t) = \{x_i \forall i \in N(c)\}$, with $x_i \in \mathbb{R}^D$, $N(c)$ represents the number of pixels covered by catchment c, and $D = 7$ is the input feature dimension we use.

3 Method and Experiments

Our goal is to find a function $f_\theta(x)$ that regresses pixel-wise ET corrections y to ET values E from input x, and we refer as $\hat{E} = E + y$ to the corrected evaporation values. We take a probabilistic view of ET datasets and their correction. We consider each dataset to provide us with prior knowledge on ET in the form of a Gaussian distribution centered on E: $p(\hat{E}) = \mathcal{N}(E|\sigma_E)$. Following a recent review paper [2], we use the relative uncertainty estimate $\sigma_E = \frac{7 \cdot E}{100}$ in our experiments. We can rewrite this prior in terms of y as $P(y) = \mathcal{N}(0|\sigma)$. We define a likelihood over the correction y as a Gaussian distribution whose mean we parameterize with a Multi Layer Perceptron (MLP) $h_\theta(x)$ and with standard deviation σ_y, which we calibrate on a validation set: $p(y|x) = \mathcal{N}(h_\theta(x)|\sigma_y)$. The correction $f_\theta(x)$ we bring to each dataset is then defined as a MAP given the prior P provided by the dataset and the parameterized likelihood function p:

$$f_\theta(x) = max_y p(y|x)P(y) \tag{2}$$

$$f_\theta(x) = \frac{\sigma_E^2 \times h_\theta(x) + 0 \times \sigma_y^2}{\sigma_y^2 + \sigma_E^2} \tag{3}$$

$$f_\theta(x) = \frac{h_\theta(x)}{1 + \frac{100 \times \sigma_y^2}{7 \times E}} \tag{4}$$

The rationale for this MAP formulation is that it allows to scale the correction with the original dataset value: Indeed, ET estimates are, in absolute values, less error-prone in very dry regions (where ET is close to zero) than in wet regions (where ET takes large values). The difference between using MAP and likelihood corrections is illustrated in Figure 2. Although $f_\theta(x)$ is defined at the pixel level, the supervision y is defined at the catchment level. To train the model, we thus first apply the model on the pixels of each catchments, then aggregate the model output by summation, and regress the aggregated sum of corrections to the label y. For a given catchment and month, the catchment level correction computed by our model is thus:

$$F_\theta(c,t) = \sum_{x \in X(c,t)} f_\theta(x) \tag{5}$$

so that we can write our loss function $\mathcal{L}$ and optimization problem as:

$$e_\theta(c,t) = F_\theta(c,t) - y(c,t) \tag{6}$$

$$\mathcal{L}(\theta) = \frac{1}{T} \sum_{c \in C} \sum_{t \in T} e_\theta(c,t)^2 \tag{7}$$

$$\theta* = min_{\theta \in \Theta} \mathcal{L}(\theta) \tag{8}$$

We analyse the errors of ET datasets before and after our correction to better understand the nature of the corrections we bring. To do so, we decompose residual errors into three components: a bias term B that represents the average error per catchment, a seasonality term S representing the errors of monthly-averaged difference to the bias, and an anomaly term A that random variations after elimination of the systemic bias ans seasonality components. We denote by $M = \{m_i\}$ to denote the set of 12 months, and we write $m(t) \in M$ to denote the month of a given time index $t \in T$. Given a catchment $c \in C$ at time t a residual term y can be decomposed into three component as follows: $e(c,t) = e(y,c) + e(y,c,m(t)) + e(y,c,t)$, in which:

$$b(e,c) = \frac{1}{T} \sum_{t \in T} e(c,t) \tag{9}$$

$$s(e,c,m) = \frac{M}{T} \sum_{t \in m} e(c,t) - b(y,c) \tag{10}$$

$$a(e,c,t) = e(c,t) - b(y,c) - s(y,c,m(t)) \tag{11}$$

Ignoring the cross terms, which were empirically found negligible, we can then decompose the loss into three residual error components that give us more insights on the nature of the ET residual error.

$$\mathcal{L}(\theta) = \sum_{c \in C} \left(b(e,c)^2 + \frac{1}{M} \sum_{m \in M} s(e,c,m)^2 + \frac{1}{T} \sum_{t \in T} a(e,c,t)^2 \right) \tag{12}$$

$$\mathcal{L}(\theta) = \sum_{c \in C} B(c) + S(c) + A(c) \tag{13}$$

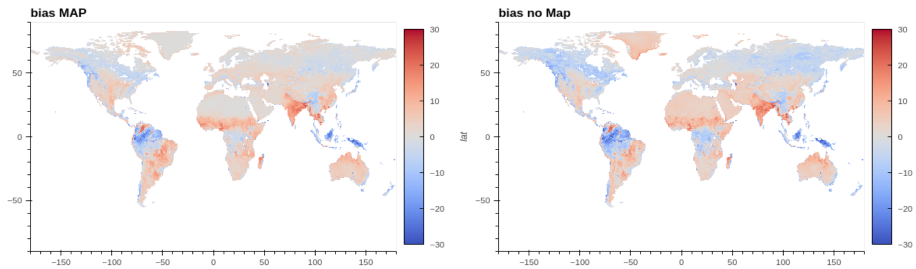

Figure 2: Illustration of MAP impact.

We trained a MLP with 4 hidden layers of width 512 on a training set of 496 catchments using the Adam [5] optimizer, and evaluated its accuracy on a test set of 166 catchments. Splits were built so that no train catchment overlap, even partially, with the test catchments. We report generalization results on the test split in terms of MSE, and decompose the error into the three components: B,S,A.
Impact of MAP modeling: Figure 2 compares the per-pixel mean corrections of the likelihood h_θ and the MAP f_θ. The MAP successfully reduces the high biases of the likelihood above the Sahara region in which ET is expected to remain close to zero for the absence of water.

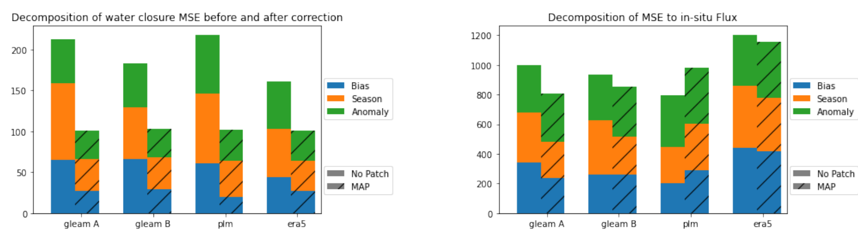

(a) Water Cycle closure MSE

(b) In-situ data MSE

Figure 3: MSE decomposition of ET estimates before (No Patch) and after (MAP) corrections. (a) Water cycle closure error computed on the test split, (b) Distance to in-situ measurements.

Water cycle closure: Figure 3(a) shows the MSE of water cycle closure of each dataset before and after applying our corrections. We successfully reduce the water closure gap on all components.**In-situ measurements**: In Figure 3(b), we show the the MSE to in-situ measurements of the FLUX dataset [15] before and after applying our learned corrections. We find corrected global ET values to better fit in-situ measurements for three out of the four datasets.

4 Conclusion

Improving global ET estimates is needed to better our understanding of the global water cycle, so as to better understand the consequences of climate change on the future of global water distribution. In this work, we proposed a methodology to learn a correction of global ET datasets. Our method uses ML to generalize sparse catchment-level water cycle closure residual information to global, dense, pixel-level residuals. To do so, we modeled a pixel-level model that we trained to regress catchment-level residuals using a sum-aggregated supervision. Using four global ET datasets, quantitative experiments have shown the ability of our model to generalize to unseen catchments and to reach relative agreement with in-situ measurements.

References

[1] (Hegerl et al. 2015,https://journals.ametsoc.org/view/journals/bams/96/7/bams-d-13-00212.1.xml).

[2] (Dorigo et al 2021, https://journals.ametsoc.org/view/journals/bams/102/10/BAMS-D-19-0316.1.xml).

[3] (Miralles et al. 2011, https://doi.org/10.1111/nyas.13912.; Fisher et al. 2017, https://doi.org/10.1002/2016WR020175.)

[4] Kingma, Diederik P., and Jimmy Ba. "Adam: A Method for Stochastic Optimization." ICLR (Poster). 2015.

[6] Martens, B., Miralles, D.G., Lievens, H., van der Schalie, R., de Jeu, R.A.M., Férnandez-Prieto, D., Beck, H.E., Dorigo, W.A., Verhoest, N.E.C., 2016. GLEAM v3: satellite-based land evaporation and root-zone soil moisture. Geoscientific Model Development Discussions , 1–36URL: http://www.geosci-model-dev-discuss.net/gmd-2016-162/, 937 doi:10.5194/gmd-2016-162.

[7] Zhang, Y., Peña-Arancibia, J., McVicar, T. et al. Multi-decadal trends in global terrestrial evapotranspiration and its components. Sci Rep 6, 19124 (2016). https://doi.org/10.1038/srep19124

[8] Hersbach, H.; Bell, B.; Berrisford, P.; Horányi, A.; Sabater, J.M.; Nicolas, J.; Radu, R.; Schepers, D.; Simmons, A.; Soci, C.; et al. Global reanalysis: Goodbye ERA-Interim, hello ERA5. In ECMWF Newsletter No. 159; Lentze, G., Ed.; European Centre for Medium-Range Weather Forecasts (ECMWF): Reading, UK, 2019; pp. 17–24.

[9] The MOD13Q1 product, DOI: 10.5067/MODIS/MOD13Q1.006

[10] Aires, F.(2014). Combining Datasets of Satellite-Retrieved Products. Part I: Methodology and Water Budget Closure. J. Hydrometeorol., 15 (4), 1677–1691.doi: 10.1175/JHM-D-13-0148.1

[11] Pellet, V., Aires, F., Munier, S., Fernández Prieto, D., Jordá, G., Dorigo, W. A., Polcher, J., and Brocca, L.: Integrating multiple satellite observations into a coherent dataset to monitor the full water cycle – application to the Mediterranean region, Hydrol. Earth Syst. Sci., 23, 465–491, https://doi.org/10.5194/hess-23-465-2019, 2019.

[12] Pan, M., Sahoo, A. K., Troy, T. J., Vinukollu, R. K., Sheffield, J., Wood, F, E. (2012). Multisource estimation of long-term terrestrial water budget for major global river basins. J. Clim., 25 (9), 3191–3206. doi:10.1175/JCLI-D-11-00300.1

[13]Munier, S., Aires, F. (2018). A new global method of satellite dataset merging and quality characterization constrained by the terrestrial water budget. Remote Sens. Environ., 205 (October 2017), 119–130. doi: 10.1016/j.rse.2017.11.008

[14] Beck, H. E., Wood, E. F., McVicar, T. R., Zambrano-Bigiarini, M., Alvarez-Garreton, C., Baez-Villanueva, O. M., Sheffield, J., Karger, D. N. (2020). Bias Correction of Global High-Resolution Precipitation Climatologies Using Streamflow Observations from 9372 Catchments, Journal of Climate, 33(4), 1299-1315. Retrieved Sep 13, 2022, from https://journals.ametsoc.org/view/journals/clim/33/4/jcli-d-19-0332.1.xml

[15] Pastorello, G., Trotta, C., Canfora, E. et al. The FLUXNET2015 dataset and the ONEFlux processing pipeline for eddy covariance data. Sci Data 7, 225 (2020). https://doi.org/10.1038/s41597-020-0534-3