Interactive comment on "Should radar precipitation depend on incident air temperature? A new estimation algorithm for cold climates" *by* Kuganesan Sivasubramaniam et al.

Anonymous Referee #2

Overview

In their paper entitled *"Should radar precipitation depend on incident air temperature? A new estimation algorithm for cold climates"*, Sivasubramaniam and colleagues present an attempt to improve radar-based quantitative estimation of precipitation by using observed near surface air temperature as a co-variate. The authors claim that the inclusion of air temperature in a non-parametric statistical estimation technique reduces the RMSE of precipitation estimates, particularly for cold climates.

The paper is fairly well written and fits the scope of HESS, although a discussion of

hydrological implications is clearly missing. However, I have serious doubts about the methodological approach, the significance of the results, and the conclusions drawn by the authors, some of which I will explain in the following.

Understanding the cumulative and non-uniform effects of temperature

The authors state, on p. 2, ll. 9 ff. that *"the scope of the present study is limited to radar precipitation estimation uncertainty during conversion from reflectivity to rain rate, with a focus on cold regions experiencing a mixture of solid and liquid precipitation."* I doubt that. Temperature not only affects the R(Z) relationship e.g. subject to the precipitation phase. Near surface air temperature either indirectly affects or correlates with different processes in the formation of precipitation in the atmosphere, but also along the radar observation and processing chain. That might lead to systematic effects on estimated precipitation intensities which accumulate over the entire estimation chain with R(Z) transformation only being the very last step. Near surface air temperature is e. g. indicative of different vertical reflectivity gradients (thus affecting observed reflectivity as a function of distance from radar), and also vertical air density gradients (affecting atmospheric refractivity and thus beam propagation / altitude). Higher precipitation intensities and thus path-integrated attenuation tend to increase with higher temperatures. The study misses is a systematic framework that takes into account different temperature effects that cause systematic precipitation estimation errors. The fact that the VPR is addressed in the met.no radar data processing chain does not mean that VPR effects (or the effects of correction) are not systematically present in the data anymore. A way to better understand these effects is to use polarimetric radar observations. That way, snow and rain can be discriminated (where radar actually measures them – not in a gauge on the ground!), so the quantitative effects could be investigated in order to understand the contribution of R(Z) uncertainty. The authors provide a general reasoning about some temperature-related effects (p. 3. ll. 25-30), but these effects are never really picked up again in the rest of the paper.

Effects of systematic undercatch of snowfall

Still, one might argue: I am not so much interested in understanding the processes behind the phenomenon. I just want to produce a better precipitation estimate (i.e. decrease the systematic error). But is that really achieved? On p. 3, ll. 1-15, the authors illustrate the motivation of their study by showing different regression slopes between radar-based precipitation (predictor) and precipitation as observed by the gauge (independent variable), for assumed snowfall and rainfall conditions. That example clearly reveals a fundamental issue: Measurement of snowfall by precipitation gauges has consistently been shown to exhibit a systematic undercatch that is significantly more pronounced than the undercatch of rainfall (see e.g. Gross et al. 2017, Wolff et al. 2015). So which effect do the authors observe: a *"bias in the radar precipitation estimation for snow"* (p. 3, l. 11), as assumed out by the authors, or a bias in the snow observation by gauges. Maybe a mix of both? In my opinion, it is almost impossible to reach any substantial conclusions based on the data and methods presented in the study.

Transparency of the cross-validation framework

The methods section does not elaborate on the leave-one-out-cross-validation (LOOCV) setup. Only in the results section, on p. 10, ll. 17-25, the application of LOOCV is pointed out. Still, the exact setup of the LOOCV remains unclear and leaves the reader with substantial doubts. If one gauge is left out to test the prediction, on which basis are the partial weights inferred for that prediction? From the nearest neighbour? From a weighted average of neighbours? For Fig. 3, as a result of what is described on p. 10, ll. 21-25, we do not know that. For Fig. 4, where an average partial weight of the entire study area is used, we do not know whether LOOCV has been applied at all, and on what basis. Only for Fig. 5, the authors state that for each gauge, the partial weights had been derived from the five neighbouring gauges. Apart from that, p. 11, l. 16 – p. 12, ll. 1-2, casts serious doubts on the integrity of the LOOCV

setup: *"As mentioned in section 3, this study uses the gauge precipitation as the observed response for the regression estimation. So insufficient data points can also be the reason for lesser improvement in these locations because nonparametric k-nearest neighbour prediction (a data based model) depends highly on the availability of sufficient data."* Does that mean that only the partial weights are independently computed from the validation target, while the observations of the target are still used in order to carry out the k-nearest neighbour regression? That could not be considered a valid LOOCV approach. The issue needs to be clarified in the minutest details, including a full disclosure of the data and the code used for the analysis. In this context, I have to emphasize that I cannot say anything about the nonparametric statistical techniques used in this study, which are outlined in section 3.1. I assume these techniques are fine, but I do not feel qualified to assess their applicability and the implications for the validity of the LOOCV, at least not from the present manuscript.

Other issues (some of them major)

- The entire section 2 (background) is far too extensive and provides a lot of information that is not pertinent to the study, and that does not play a role for discussing the results.

- How is Eq. 1 an equation?

- Section 3.2: Would be helpful to additionally use an evaluation criterion that measures the systematic error (e.g. mean error)

- p. 8, ll. 27-29: It should be carefully analyzed whether simply choosing the nearest neighbour provides the best correspondence between radar-based QPE and rain gauge observations. Particularly at hourly intervals, the consistency depends on the neighbourhood definition due to representativeness issues.

- p. 8, ll. 33-25: Including such intensities as low as 0.05 – 0.5 mm might lead to the fact that insignificant precipitation dominates the results in terms of relative changes of the RMSE, as hourly rainfall follows a gamma-like distribution. Confining the analysis to significant precipitation should address that issue.

- Most figures in section 5: I find the continuous colorbars very difficult to interpret. Please use a discrete colorbar instead.

- p. 10, ll. 4-6: How do you know the number of snow data pairs?

- p. 10, ll. 9-10: *"This outcome is a result of sampling uncertainty due to which a minimum of 0.2 for the partial weight for radar rain rate has been used in the results."* That explanation is not satisfactory at all. More generally, any gauge with a very high partial weight for temperature should be considered with great caution – why should temperature be a better predictor than radar rainfall?? Presumably only in case radar rainfall at that particular location is affected by serious artefacts. Here, a systematic analysis of the relation between radar and gauges on a per-gauge basis is required, together with a spatial analysis of systematic errors in the radar rainfall (e.g. due to partial beam blockage, residual clutter, ...).

- Table 1: Why not show a histogram instead?

- p. 10. ll. 27-28: *"Further, it can be noted that all the gauge locations with an associated partial weight of air temperature (betaT > 0) shows an improvement in radar precipitation estimation."* What does that mean in the context of a cross-validation where you do not know the partial weight at the target location?

- p. 10, ll. 1-2: *"raingauge locations with a minimum partial weight for radar rain rate (betaR = 0.2 betaT = 0.8), did not show improvement in RMSE."* How can you compare the two settings – one only based on radar rainfall and the other with betaR arbitrarily set to 0.2?

- For all investigated changes in RMSE, please investigate the statistical significance of that change (e.g. by using bootstrapping). E.g. on p. 11, ll. 7-8, it is stated that *"over 80 percent of the gauge locations in the study area show more than 3 percent improvement in RMSE"*. Which portion of these changes is statistically significant?

- p. 11, ll. 3-11: What is the implication of the fact that using an average partial weight over the entire study area produces better results? Was that result also achieved from an LOOCV? It is important to understand how the partial weights were assigned in the analysis that assigns a specific partial weight to each station, on order to understand the implications of this experiment.

- p. 11, ll. 3-11: One implication of this paragraph is that the spatial pattern of partial weights is meaningless (unless proved otherwise), since a simple average provides better results. As a consequence, I recommend to drop the maps in section 5, and show e.g. histograms instead, or any other visualisation that allows the reader to better understand the quantitative implications of the results.

- p. 12, ll. 11-13: Meaning remains unclear.

- p. 13, ll. 1-5: *"[...] resulted in a maximum of 20 percent total improvement [...]"* – where has that been shown? *"[...] The nonparametric k-nn predictive model with radar rain rate as a single predictor improves the prediction."* – that has not been shown in the results, either – no benchmark based on a direct rainfall estimation from reflectivity is shown.

- p. 13, l. 12 ff.: That paragraph basically shows that the predictive model does not contribute to an understanding of temperature effects. It rather feeds the suspicion that the effect of temperature merely balances different observational biases of the precipitation gauges with regard to rain and snow.

References

Grossi, G., Lendvai, A., Peretti, G., and Ranzi, R. (2017): Snow Precipitation Measured by Gauges: Systematic Error Estimation and Data Series Correction in the Central Italian Alps, Water, 9, 461; doi: 10.3390/w9070461.

Wolff, M. A., Isaksen, K., Petersen-Øverleir, A., Ødemark, K., Reitan, T., and Brækkan, R. (2015): Derivation of a new continuous adjustment function for correcting wind-induced loss of solid precipitation: results of a Norwegian field study, Hydrol. Earth Syst. Sci., 19, 951-967, https://doi.org/10.5194/hess-19-951-2015.