EGU24-19257, updated on 19 Jul 2024
https://doi.org/10.5194/egusphere-egu24-19257
EGU General Assembly 2024

Dynamic Locally Binned Density Loss

Jan Prosi[1,3], Sebastian Otte[2], and Martin V. Butz[1]
[1]Cognitive Modeling, Department of Computer Science, University of Tübingen, Germany
[2]Adaptive AI Research Group, Institute for Robotics and Cognitive Systems, University of Lübeck
[3]Max Planck Institute for Intelligent Systems, Tübingen, Germany

In the field of precipitation nowcasting recent deep learning models now outperform traditional approaches such as optical flow [1,2]. Despite their principled effectiveness, these models and their respective training setups suffer from particular shortcomings. For instance, they often rely on pixel-wise losses, which lead to blurred predictions by which the model expresses its uncertainty [2]. Additionally, these losses can negatively impact training dynamics by overly penalizing small spatial or temporal discrepancies between predictions and actual observations, i.e., the double penalty problem [3]. Generative methods such as discriminative losses or diffusion models do not suffer from the blurring effect as much [1, 4]. However, training these methods is complicated because training success is highly sensitive to the network architecture as well as to the learning setup and its parameterization [5].

Previous research has shown that spatial verification methods such as the fractions skill score offer an easy-to-implement alternative to solve the problem of pixel-wise losses [6, 7]. However, the fact that each pixel within the neighborhood of a spatial kernel is weighted equally poses a limiting factor to their performance and potential. Inspired by theories of cognitive modeling and in relation to the fractions skill score loss, we introduce a dynamic locally binned density (DLBD) loss: Forecasting target is not the actual precipitation in a grid cell but a target distribution, which encodes the density of binned precipitation values in a locally weighted area of grid cells. The loss is then determined via the cross-entropy of the predicted and the target distribution. We show that our novel prediction loss avoids the double penalty problem. It thus diminishes the negative impact of small spatial offsets. Moreover, it enables the learning model to gradually shift focus towards progressively more accurate predictions.

We achieve best performance by simultaneously training on multiple concurrent forecasting targets that cover different local extents. We schedule the weighting of the loss terms such that the focus shifts from larger to smaller neighborhoods over the course of training. This way, the DL model first learns density dynamics and basic precipitation shifts. Later, it focuses on minimizing small spatial deviations, tuning into the local dynamics towards the end of training. Our DLBD loss is easy-to-implement and shows great performance improvements. We thus believe that DLBD losses can also be used by other forecasting architectures where the current forecasting loss precludes smooth loss landscapes.

1: Leinonen et al. 2023: Latent diffusion models for generative precipitation nowcasting with accurate uncertainty quantification
2: Espeholt et al. 2022: Deep learning for twelve hour precipitation forecasts
3: Grilleland et al. 2009: Intercomparison of spatial forecast verification methods.
4: Ravuri et al. 2021: Skilful precipitation nowcasting using deep generative models of radar
5: Mescheder et al. 2018: Which training methods for GANs do actually converge?
6: Roberts et al. 2008: Scale-selective verification of rainfall accumulations from high resolution forecasts of convective events.
7: Lagerquist et al. 2022: Can we integrate spatial verification methods into neural-network loss functions for atmospheric science?