dr Piotr Mirowski
DeepMind

Precipitation nowcasting, the high-resolution forecasting of precipitation up to two hours ahead, supports the real-world socioeconomic needs of many sectors reliant on weather-dependent decision-making. State-of-the-art operational nowcasting methods typically advect precipitation fields with radar-based wind estimates, and struggle to capture important nonlinear events such as convective initiations. Recently introduced deep learning methods use radar to directly predict future rain rates, free of physical constraints. While they accurately predict low-intensity rainfall, their operational utility is limited because their lack of constraints produces blurry nowcasts at longer lead times, yielding poor performance on rarer medium-to-heavy rain events. Here we present a deep generative model for the probabilistic nowcasting of precipitation from radar that addresses these challenges. Using statistical, economic and cognitive measures, we show that our method provides improved forecast quality, forecast consistency and forecast value. Our model produces realistic and spatiotemporally consistent predictions over regions up to 1,536 km × 1,280 km and with lead times from 5–90 min ahead. Using a systematic evaluation by more than 50 expert meteorologists, we show that our generative model ranked first for its accuracy and usefulness in 89% of cases against two competitive methods. When verified quantitatively, these nowcasts are skillful without resorting to blurring. We show that generative nowcasting can provide probabilistic predictions that improve forecast value and support operational utility, and at resolutions and lead times where alternative methods struggle.