

42 weather information, and cannot process this information in time, the task of
43 precipitation nowcasting is challenging. We take advantage of a novel machine learning
44 approach to learn what possible precipitation conditions are, given current precipitation
45 condition observed from radar. Our results offer accurate precipitation prediction. More
46 importantly, this method assigns high uncertainty to predictions where predictions are
47 more biased. This accurate estimate of prediction uncertainty is crucial for weather
48 related decision makings.

50 **1 Introduction**

51 Precipitation nowcasting is the task of predicting upcoming precipitation (e.g, 0-2
52 hours) at high spatiotemporal resolutions. Reliable precipitation nowcasting, especially
53 for storm cases, is crucial for risk and crisis preparation, water resource management,
54 and many other societal sectors (Zhang et al., 2023).

55 Numerical weather prediction provides the most reliable short-to-medium range
56 (6 hours to 2 weeks) forecasts. It makes predictions by first inferring the initial weather
57 state, followed by calculating the state evolution, using numerical solvers of
58 atmospheric fluid dynamics, and associated parameterization schemes that account for
59 unresolved processes. Despite its theoretical soundness, numerical weather prediction
60 offers poor precipitation nowcasting, due to difficulty in assimilating hydrometeor
61 observations, limited spatiotemporal resolution, and high computation cost.

62 Empirical methods can make flexible use of detailed initial hydrometeor
63 observations, such as those from radar and satellite. Vanilla forecasts can therefore be
64 achieved by simply propagating the initial observations along time, such as the optical
65 flow approach (Cheung & Yeung.,2012; Pulkkinen et al., 2019; Sakaino., 2013). More
66 advanced approaches try to better simulate the dynamical processes by “learning” from
67 data. These data-driven models are highly parameterized functions, for which the
68 functional design is guided by inductive biases of the considered process, and the
69 parameters are optimized by fitting the data to the model, guided by a learning objective
70 function.

71 The design of learning objective functions is vital for data-driven prediction. A
72 popular option is to minimize the mean squared error between predictions and
73 observations. This objective function is based on the assumption that plausible
74 predictions subject to a conditional Gaussian distribution, where the mean vector is a
75 learnable function of the initial state, and covariance matrix is independent of the initial
76 state:

$$P_{\theta}(y|x) = N(y; \mu_{\theta}(x), \sigma^2), \quad (1)$$

77 here $\mu_{\theta}(x)$ serves as the deterministic forecast. This formulation comes with two
78 shortcomings. Firstly, it prohibits the exploration of the spatial structure of predictions,
79 making it difficult to leverage data-informed prior knowledge for achieving structurally
80 reasonable predictions. Secondly, it assumes a deterministic outcome, despite the
81 absence of a full-profile and strictly accurate initial state estimate. As a result,
82 deterministic models tend to yield poorly structured, blurry estimates, missing extreme
83 cases and uncertainty quantification. These deficiencies are evident in models such as
84 ConvLSTM, ConvGRU and Unet (Shi et al., 2015, 2017; Ayzel et al., 2019, 2020).

85 To fully explore the spatial structure of data and provide predictions along with
86 uncertainty information, it is imperative to free our predictive model from a pre-defined
87 distributional form. Instead, it is preferable to deploy generative models to learn
88 empirical distribution that maximize the likelihood of the observations:

$$\hat{y} \sim P_{\theta}(y|x), \text{ where } \theta = \operatorname{argmax}_{\theta} P(y|x; \theta). \quad (2)$$

89 A landmarking work along this direction is the Deep Generative Models of Radar
90 (DGMR, Ravuri et al., 2021), which achieves state-of-the-art performance regarding
91 the forecast skill and value. We believe the key contribution of DGMR is that, it marks
92 a pioneering attempt to bridge probabilistic forecast and generative modeling: a
93 probabilistic forecast should encapsulate all plausible outcomes (requirement of
94 calibration), thereafter maximize the sharpness of its predictive distribution
95 (requirement of sharpness, Gneiting et al., 2007). DGMR employs a spatial and a
96 temporal discriminator neural network to guarantee that observation stays within the
97 predictive distribution. Meanwhile, it implicitly enhances the sharpness of its predictive
98 distribution by having the ensemble mean stay close to observation. There are two
99 potential drawbacks here. First, the two objectives in DGMR can be in conflict, making
100 it tricky to maximize the sharpness of the predictive distribution while guaranteeing the
101 model is well calibrated. Second, due to unneglectable optimization errors, generative
102 adversarial net (GAN) tends to miss plausible modes in approximating complicated
103 distributions, resulting in biased probabilistic forecast (Prafulla Dhariwal & Alex
104 Nichol, 2021; Ali Razavi et al., 2019).

105 To address these challenges, we introduce diffusion models (Sohl-Dickstein et al.,
106 2015; Song & Ermon, 2020b; Ho et al., 2020) for precipitation nowcasting. Unlike
107 GANs, probabilistic diffusion models are likelihood-based generative models, that is,
108 they are trained to directly maximize the probability assigned to the observed samples.
109 This enables a full coverage of the target distributions (Ali Razavi et al., 2019; Dhariwal
110 & Nichol, 2021). Moreover, their iterative generation nature allows us to flexibly
111 control the resulting distribution using initial state information. As a result, we can
112 gradually enhance the sharpness of the predictive distribution, while guaranteeing the
113 predictive distribution encapsulates all plausible outcomes.

114 Diffusion models have proven successful in various research domains, tackling
115 complex tasks like image synthesis (Dhariwal & Nichol, 2021), audio synthesis (Kong
116 et al., 2020), and video generation (Voleti et al., 2022; Höppe et al., 2022; Ho et al.,
117 2022). Their desirable properties make them an effective tool for achieving reliable
118 probabilistic forecasts with informative forecast uncertainty estimates. In this study, we
119 propose an advanced diffusion model of nowcasting and verify with the subset of well-
120 established UK MetOffice radar dataset.

122 2 Methods

123 2.1 Probabilistic modeling the Precipitation nowcasting

124 Consider a sequence of precipitation field data $\mathbf{R} = [r_1, r_2, \dots, r_M]$, the nowcasting
125 task is to predict future precipitation field trajectories (N fields) based on a given past
126 trajectory of observations (M fields). Here, we formulate this problem as a probabilistic
127 machine learning task. Using an extensive dataset of sequences of precipitation field

128 data, we learn conditional probability model of $P_\theta(\mathbf{R}_{M+1:M+N}|\mathbf{R}_{1:M})$, thus

$$\hat{r}_{M+1}, \dots, \hat{r}_{M+N} \sim P_\theta(\mathbf{R}_{M+1:M+N}|\mathbf{R}_{1:M}). \quad (3)$$

129 This learning process is facilitated by a conditional diffusion model. A common
130 strategy for approximating this target distribution is learning a mapping between the
131 target and a tractable latent distribution, such as a standard Gaussian. Then we can
132 deduce the target distribution via a procedure termed ancestral sampling, describe by

$$P_\theta(\mathbf{R}_{M+1:M+N}|\mathbf{R}_{1:M}) = \int P(\mathbf{R}_{M+1:M+N}|Z, \mathbf{R}_{1:M}, \theta)P(Z|\mathbf{R}_{1:M})dZ \quad (4)$$

133 In the following sections, we demonstrate how this is accomplished in diffusion
134 models. Mathematical details are given in Supporting Information.

135

136 2.2 Basic diffusion

137 Diffusion model approximates a target distribution by sequentially reversing a
138 stochastic process, using a series of neural network models. Let $P(X_0)$ be the target
139 distribution. We define the following discrete time Gaussian process:

$$q(X_t|X_{t-1}) = N(X_t; \sqrt{1 - \beta_t}X_{t-1}, \beta_t \mathbf{I}) \quad (4)$$

140 Here, $X_{t=[1,T]}$ are latent variables. $0 < \beta_t < 1$ is diffusion coefficient. Given large
141 enough T , $q(X_T|X_0)$ is close to standard Gaussian. Therefore, the forward Gaussian
142 process maps any target distribution $P(X_0)$ to standard Gaussian. To approximate $P(X_0)$,
143 starting from standard Gaussian, we sequentially reverse the Gaussian process using
144 the following variational distributions:

$$P_\theta(X_{t-1}|X_t) = N(X_{t-1}; \mu_\theta(X_t), \Sigma_\theta(X_t)) \quad (5)$$

145 a common objective function for learning these variational distributions is the following
146 evidence lower bound L_{VLB} defined over $X_{1:T}$,

$$L_{VLB} = \mathbb{E}_q[D_{KL}(q(X_T|X_0)||p_\theta(X_T)) + \sum_{t=2}^T D_{KL}(q(X_{t-1}|X_t, X_0)||p_\theta(X_{t-1}|X_t)) - \log p_\theta(X_0|X_1)] \quad (6)$$

147 under certain simplification, this evidence lower bound can be simplified to a
148 remarkably short expression in terms of fisher divergence:

$$L_{simple} = \mathbb{E}_{t \sim [1,T], X_0 \sim q(X_0), \epsilon \sim N(0, I)} [|\nabla \log P(X_t) - \epsilon_\theta(X_t, t)|^2] \quad (7)$$

149 Here $\epsilon_\theta(X_t, t)$ is a neural network parameterization of $\nabla \log P(X_t)$, which is called
150 score function. By learning the score function of the true data distribution, we can
151 generate samples by starting at $X_T \sim N(0, I)$, and iteratively following the score function
152 until a mode (X_0) is reached.

153

154 2.3 Conditional diffusion

155 Our objective is to approximate the conditional distribution of $P(X_t|y)$. Begin with
156 the score-based formulation of a diffusion model, the goal is to learn $\nabla \log P(X_t|y)$, by
157 Bayes rules, we can get the equivalent:

$$\nabla \log P(X_t|y) = \nabla \log \left(\frac{P(y|X_t)P(X_t)}{P(y)} \right) \quad (8)$$

$$= \nabla \log P(X_t) + \nabla \log P(y|X_t) - \nabla \log P(y) \quad (9)$$

$$= \underbrace{\nabla \log P(X_t)}_{\text{unconditional score}} + \underbrace{\nabla \log P(y|X_t)}_{\text{conditional score}} \quad (10)$$

158 To better control the conditional information, a hyperparameter γ is introduced to
 159 scale the gradient of the conditioning information. The score function can then be
 160 summarized as:

$$\nabla \log P(X_t|y) = \nabla \log P(X_t) + \gamma \nabla \log P(y|X_t). \quad (11)$$

161 Intuitively speaking, the $\gamma = 0$ the diffusion model can ignore the conditional
 162 information entirely, while a large γ value would cause the model to heavily incorporate
 163 the conditional information during sampling. In order to implement effective control
 164 over the conditional information, we use classifier-free guidance (Ho & Salimans, 2021).
 165 To get the score function under Classifier-Free Guidance, we can rearrange:

$$\nabla \log P(y|X_t) = \nabla \log P(X_t|y) - \nabla \log P(X_t). \quad (12)$$

166 Substituting equation (12) into equation (11) then we get:

$$\nabla \log P(X_t|y) = \nabla \log P(X_t) + \gamma(\nabla \log P(X_t|y) - \nabla \log P(X_t)). \quad (13)$$

$$= \underbrace{(1 - \gamma)\nabla \log P(X_t)}_{\text{unconditional score}} + \underbrace{\gamma\nabla \log P(X_t|y)}_{\text{conditional score}} \quad (14)$$

167 In this paper, we model the conditional distribution of precipitation frames in the
 168 future given the past precipitation frames $\mathbf{R} = [p_1, p_2, \dots, p_M]$, we learn two sets of neural
 169 networks, $\epsilon_\theta(X_t, t)$ and $\epsilon_\theta(X_t, t, R)$, to approximate the unconditional and conditional
 170 score functions $\nabla \log P(X_t)$ and $\nabla \log P(X_t|y)$, our conditional diffusion loss function is:

$$L_{\text{condition}} = \mathbb{E}_{t \sim [1, T], X_0 \sim q(X_0), \epsilon \sim N(0, I)} \left[\|\nabla \log P(X_t|y) - \epsilon_\theta(X_t, t, R)\|^2 \right] \quad (15)$$

171

172 **3 Data**

173 We utilized the publicly available UK MetOffice radar network dataset, which was
 174 obtained from DeepMind (Ravuri et al., 2021). The dataset provides radar echo data
 175 with a temporal resolution of 5 minutes and a spatial resolution of 1 km for the entire
 176 UK region from 2015 to 2019. Each data point in the dataset consists of 24 time steps
 177 and covers an area of 256 km x 256 km.

178 Due to computational resource limitations, we employed a subset comprising
 179 11,000 radar samples, partitioned into three subsets: training (8,000 samples),
 180 validation (2,000 samples), and testing (1,000 samples). The principal objective of this
 181 investigation is to assess the efficacy and reliability of diffusion-based and GAN-based
 182 models for precipitation nowcasting. To optimize resource usage, we exclusively
 183 evaluated these models for 30-minute precipitation predictions. Consequently, we
 184 performed random 80x80 sub-sample extractions from the original 256x256-sized data
 185 to speed up training.

186

187 **4 Model Evaluation**

188 **4.1 Baseline models**

189 Generative models of radar (DGMR) holds the current state of the art in
 190 precipitation nowcasting. We utilized Google-Colab to load the pre-trained DGMR
 191 model and evaluate its performance using the first 30 minutes of forecasted results
 192 (Ravuri et al., 2021). UNet serves as the baseline for deterministic forecasting using
 193 deep learning (Ayzel et al., 2020). PySTEPS is a widely used precipitation nowcasting
 194 system based on ensembles (Pulkkinen et al., 2019). We adopt PySTEPS as a non-

195 machine learning baseline. More details of the baseline can be found in the support
196 information.

197

198 **4.2 Evaluation strategy**

199 We employ various metrics to assess the performance of both the baseline and
200 diffusion models on the test set. We evaluate the deterministic skill of the ensemble
201 mean using the mean absolute error (MAE), and we provide versions of MAE that
202 consider extreme value prediction accuracy under different precipitation intensities.
203 The accuracy of spatial prediction is evaluated using the Critical Success Index (CSI)
204 at different precipitation thresholds. We use the Pearson correlation coefficient to
205 evaluate the spatial pattern of predictions at different resolutions. Furthermore, the
206 calibration and sharpness of the ensemble together is evaluated using Continuous
207 Ranked Probability Score (CRPS) at different spatial scales. As a measure of the
208 reliability of the ensemble, we examine the spread-skill ratio (Spread/RMSE). For
209 details of these metrics, see support information.

210

211 **5 Results and discussion**

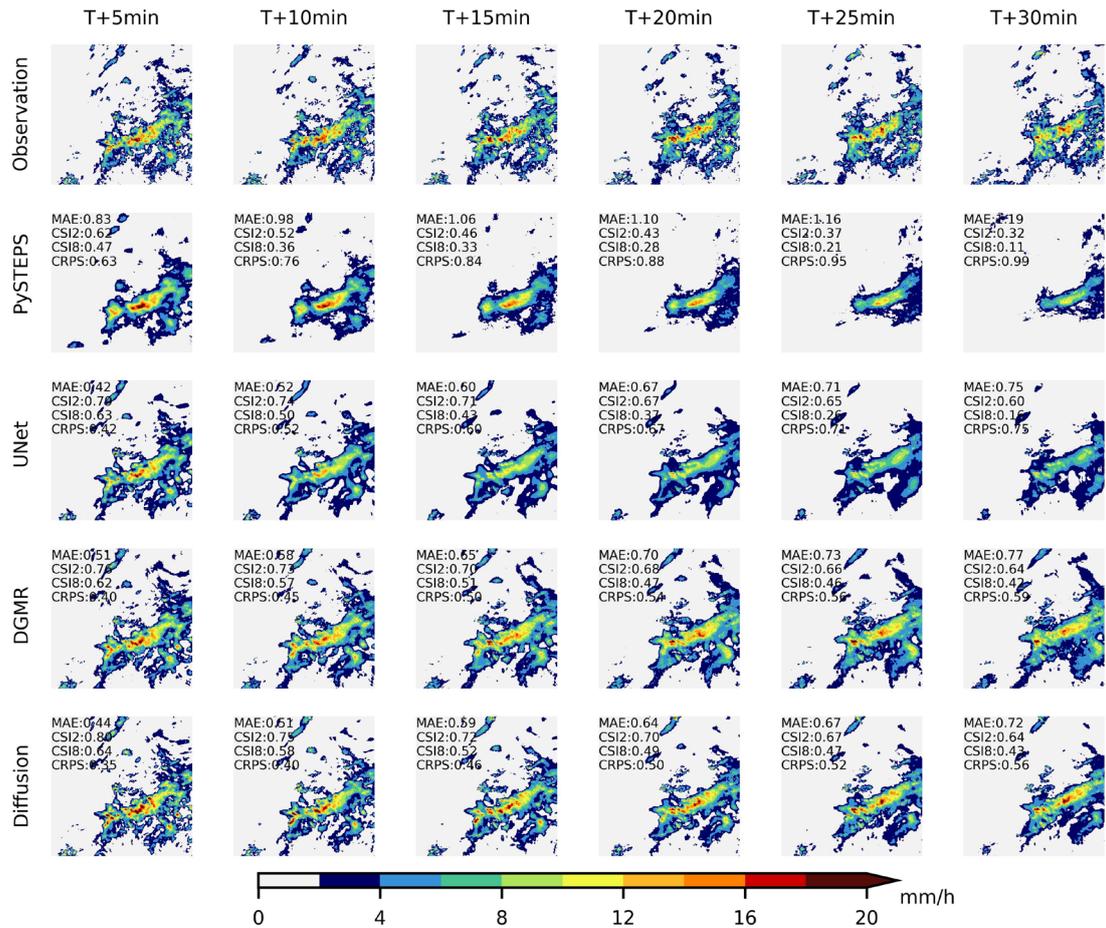
212 **5.1 Model performance for heavy precipitation forecasts**

213 We employ a case study of heavy precipitation to compare the performance of our
214 model with the three baseline models. Figure 1 shows the ground truth and predicted
215 precipitation fields. In this case, our model has consistently demonstrated superior
216 performance across various evaluation metrics.

217 PySTEPS tends to underestimate the temporal changes in precipitation intensity,
218 and falls short in adequately capturing the entire precipitation field. As lead time
219 increases, the UNet model provides only coarse estimates of the precipitation field,
220 resulting in highly blurred predictions that lack accuracy in predicting precipitation
221 intensity and small-scale spatial features.

222 GAN-based models (DGMR) can indeed address blurred predictions. However, it
223 is more difficult to capture the precipitation pattern, results in poor probabilistic
224 forecasting performance, which is evident on larger CRPS, higher ensemble-averaged
225 MAE and worse CSI compared to diffusion models.

226 By comparison, our model ensures accurate and comprehensive coverage of
227 precipitation fields and shows an enhanced ability in predicting precipitation intensity
228 and small-scale spatial features, making its predictions more informative and valuable.



229

230 **Figure 1.** The performance of different baselines in heavy precipitation scenarios. The
 231 predictions for 6 time steps from T+5min to T+30min were evaluated. CSI at thresholds
 232 2 (mm/h) and 8 (mm/h), MAE and CRPS for an ensemble of 8 samples displayed in the
 233 top left corner of each time step prediction.

234

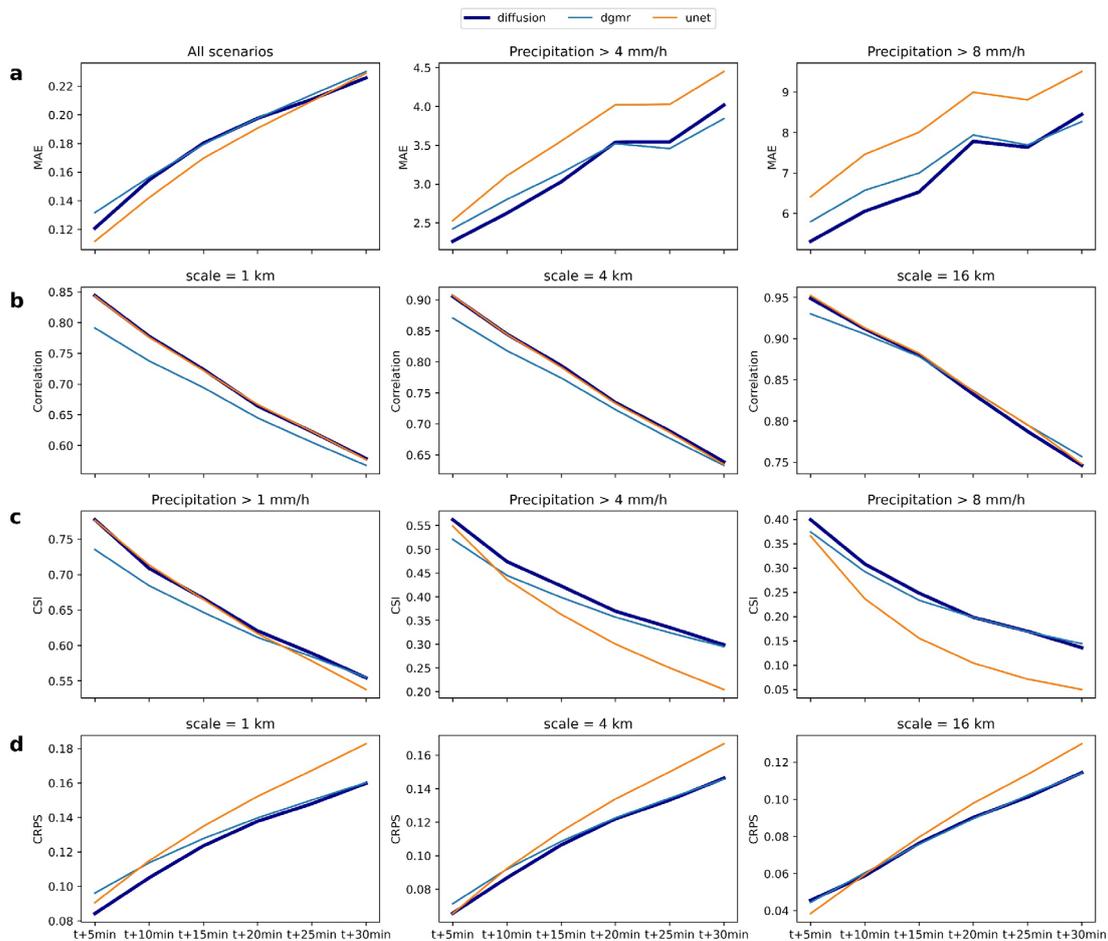
235 5.2 Forecast skill evaluation

236 Machine learning methods are superior to PySTEPS indicated by all metrics except
 237 for CSI (thresholds at 8 mm/h) where PySTEPS outperforms UNet. For the sake of
 238 clarity, Figure 2 will not display the metrics for PySTEPS, the complete forecast skill
 239 evaluation can be found in the support information.

240 Figure 2a (all scenarios) shows that the performance of UNet is slightly better than
 241 that of DGMR and diffusion on MAE. It is because that deterministic models are
 242 optimized for the mean of all precipitation scenarios, and therefore, the ensemble mean
 243 is expected to exhibit slightly lower performance in deterministic metrics like
 244 correlation and MAE compared to UNet. Unet's performance noticeably declines for
 245 heavy precipitation due to its tendency to generate blurred precipitation forecasts and
 246 our model (diffusion) performing better on heavy precipitation. Figure 2b evaluates the
 247 spatial correlation at different resolutions. Our model performs similarly to UNET and
 248 outperforms DGMR at resolutions of 1 km and 4 km. Figure 2c proves the superiority
 249 of our model over other baseline models in terms of location accuracy, as measured

250 across varying CSI threshold values. Unet deteriorates significantly with increased lead
 251 time and CSI threshold due to its inherent theoretical constraints in addressing this
 252 challenge. At both grid scales (1km) and 4km spatial resolutions, our model surpasses
 253 other baseline models in terms of CRPS (Figure 2d). With a spatial resolution of 16km,
 254 our model performance aligns with that of DGMR.

255 Despite being trained against a limited dataset, our model shows significant
 256 competitiveness. Within its 30-minute training period, our model consistently surpasses
 257 DGMR and other baselines in CSI and CRPS metrics. On average, across all forecasted
 258 time steps, our model exhibits an improvement of 3.7% in the CRPS (at a resolution of
 259 1km) and an enhancement of 2.6%, 5.2%, 3.5% in CSI at an intensity threshold of 1.0,
 260 4.0, and 8.0 mm/h, compared to the DGMR.



261
 262 **Figure 2.** Evaluation metrics for the test-dataset. The probability forecast is generated
 263 using 8 ensemble members, while the Unet model is used for a single deterministic
 264 forecast. **a**, shows the MAE under different precipitation intensity conditions. MAE
 265 across all precipitation conditions (left); MAE considering observed precipitation
 266 greater than 4 mm/h (middle), MAE considering observed precipitation greater than 8
 267 mm/h (right). **b**, correlation at different resolution. **c**, CSI for precipitation thresholds
 268 at 1 mm/h (left), 4 mm/h (middle) and 8 mm/h (right). **d**, CRPS score at different spital
 269 resolution. Grid resolution (1km) (left), average pooled 4km resolution (middle),
 270 average pooled 16km resolution (right). For MAE and CRPS, lower is better. For CSI
 271 and correlation, closer to 1 is better.

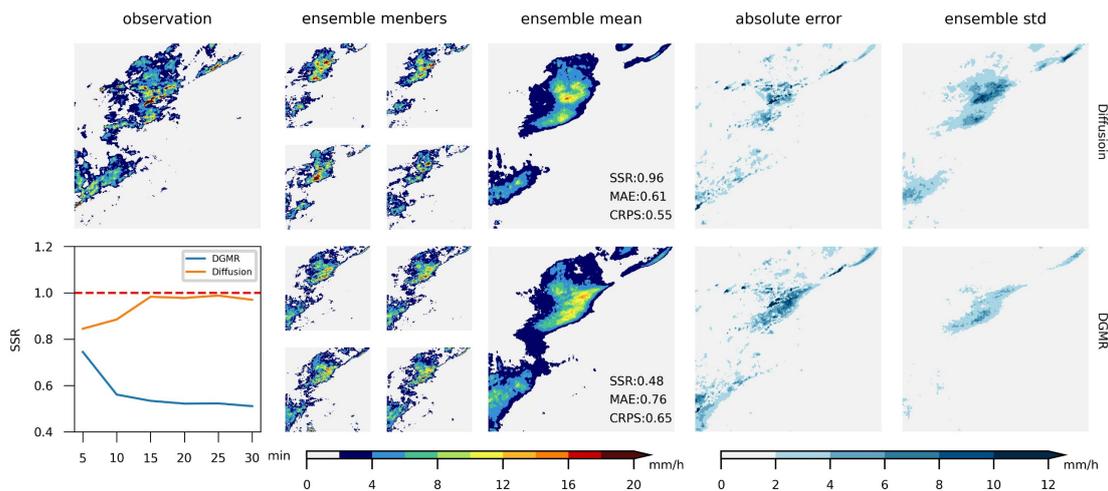
272

273 5.3 Reliability quantification

274 The incorporation of reliability estimation is crucial for decision-making processes
275 and risk assessment. We assess forecast reliability for the diffusion model and DGMR,
276 presenting ensemble members, ensemble mean, standard deviation, and absolute error
277 maps in comparison with observations at the thirtieth minute (Figure 3). Reliability
278 evaluations for alternative scenarios are available in the support information. The
279 standard deviation represents the spread of the ensemble predictions, serving as a proxy
280 of uncertainty within the precipitation forecasts. The spatial map of absolute error
281 provides insights into the areas where the model may struggle to predict. Therefore, it
282 is desirable for the model to provide a higher level of uncertainty in regions where its
283 performance is poor. A balance between calibration and spread must be achieved.

284 Figure 3 illustrates that DGMR achieve a smaller standard deviation compared to
285 diffusion model, which is also reflected in a high degree of similarity among ensemble
286 members. DGMR enhanced ensemble sharpness, but fell short in terms of calibration,
287 evident in the larger mean absolute error. It failed to establish a spatial consistency
288 between forecast skill and forecast spread. For example, DGMR's predictions fail to
289 reflect uncertainty at the left boundary. This means DGMR may generate overconfident
290 predictions. Reliability is quantified using the spread-skill ratio (SSR), where an ideal
291 ensemble model yields an SSR of 1.0. Here, the diffusion model attains an SSR of 0.96,
292 surpassing DGMR's 0.48, establishing its superior reliability. Additionally, diffusion
293 model exhibits superior probabilistic and deterministic forecast skills.

294 We also calculated SSR over the test dataset, displayed at the bottom left of
295 Figure 3. For DGMR, the SSR values are 0.745, 0.561, 0.534, 0.522, 0.523, and 0.511,
296 with an average of 0.56. In contrast, diffusion yields SSR values 0.845, 0.885, 0.983,
297 0.978, 0.988, and 0.970, with an average of 0.94. Diffusion model achieves a 68% gain
298 in the spread-skill ratio, underscoring its ability to provide more reliable forecasts.



299

300 **Figure 3.** The example of ensemble forecasts provided by DGMR and Diffusion at the
301 thirtieth minute. From left to right: four randomly selected ensemble members, the
302 ensemble mean, the absolute error map comparing the ensemble mean to observations,
303 and the ensemble standard deviation. The bottom-left panel displays the reliability

304 quantification SSR (Spread-Skill Ratio) calculated using the entire test dataset for
305 forecasts.

306

307 **6 Conclusions**

308 Predicting when and where precipitation is likely to occur with high accuracy in
309 the short term remains a difficult task. Such forecasts are essentially probabilistic: as
310 we do not have comprehensive initial weather state estimate, and cannot fully resolve
311 the weather dynamics, we should provide a range of possible outcomes along with their
312 likelihood estimates, instead of a single deterministic prediction.

313 Data-driven methods have proven highly advantageous for precipitation
314 nowcasting, due to their flexibility in utilizing detailed initial hydrometeor observations,
315 and their capability to approximate meteorological dynamics effectively. State-of-the-
316 art data-driven precipitation nowcasting approaches take advantage of deep generative
317 models to yield probabilistic forecast. However, these methods, mostly based on
318 generative adversarial nets (Goodfellow et al. 2014), are often faced with severe
319 approximation/optimization errors, rendering their predictions and associated
320 uncertainty estimates unreliable.

321 In this study, we present a probabilistic diffusion model-based methodology for
322 precipitation nowcasting. The model learns predictive distributions by explicitly
323 maximizing the data likelihood. It achieves advantageous sample fidelity, distribution
324 diversity, and control flexibility by applying a principled, iterative way for generative
325 modeling tasks.

326 Our diffusion model provides significantly improved probabilistic forecasts and
327 consistently outperforms benchmark models over a thirty-minute forecast period, as
328 indicated by well-established probabilistic CRPS and SSR skill scores. In terms of
329 deterministic metrics, including MAE CSI and correlation, our model performs on par
330 with the deterministic model UNet and probabilistic model DGMR but particularly
331 excels Unet for heavy rainfall forecasts. More importantly, the diffusion model provides
332 a more informative assessment of the uncertainty associated with its forecasts, making
333 its prediction more reliable.

334 However, there remain some challenges to be addressed for our probabilistic
335 nowcasting model. Its high computational resource requirement restricts the input size
336 and limits our prediction horizon to 30 minutes. Nevertheless, this constraint may
337 potentially be addressed by employing a latent diffusion model (Robin et al., 2021).
338 Furthermore, we could explore the use of 3D convolutions and the development of
339 temporal attention modules to improve temporal continuity.

340 In conclusion, despite these constraints, our model has demonstrated superior
341 predictive accuracy and reliability. These qualities make our model a promising tool for
342 precipitation nowcasting, capable of delivering more accurate and reliable forecasts.

343

344 **Data availability**

345 All data used in this study are available from Ravuri et al. (2021).

346

347 **Acknowledgements**

348 This research was supported by the Third Xinjiang Scientific Expedition Program
349 (Grant No. 2021xjkk0806), the National Natural Science Foundation of China
350 (42271032, U2243226, 42288101), NSFC-DFG mobility (M-0468), Chinese Academy
351 of Science Light of the West Interdisciplinary Research Grant (grant no. xbzg-zdsys-
352 202104) and National Key R&D Program of China 2021YFA0718000.

353

354 **Reference**

355 Ayzel, G., Heistermann, M., and Winterrath, T (2019). Optical flow models as an open
356 benchmark for radar-based precipitation nowcasting (rainymotion v0.1), *Geosci.*
357 *Model Dev.*, 12, 1387–1402, <https://doi.org/10.5194/gmd-12-1387-2019>

358 Ayzel, G., Scheffer, T., and Heistermann, M (2020). RainNet v1.0: a convolutional neural
359 network for radar-based precipitation nowcasting, *Geosci. Model Dev.*, 13, 2631–2644,
360 <https://doi.org/10.5194/gmd-13-2631-2020>

361 Ayzel, G., Scheffer, T., and Heistermann, M. (2020). A convolutional neural network for radar-
362 based precipitation nowcasting, *Geosci. Model Dev.*, 13, 2631–2644,
363 <https://doi.org/10.5194/gmd-13-2631-2020>

364 Cheung, P., & Yeung, H. Y. (2012). Application of optical-flow technique to significant
365 convection nowcast for terminal areas in Hong Kong.

366 Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and
367 sharpness. *Journal of the Royal Statistical Society Series B: Statistical Methodology*,
368 69(2), 243–268. <https://doi.org/10.1111/j.1467-9868.2007.00587.x>

369 Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio,
370 Y. (2014). Generative adversarial nets. *Advances in neural information processing*
371 *systems*, 27. <https://doi.org/10.48550/arXiv.1406.2661>

372 H. Sakaino, (2013). Spatio-Temporal Image Pattern Prediction Method Based on a Physical
373 Model With Time-Varying Optical Flow, in *IEEE Transactions on Geoscience and*
374 *Remote Sensing*, vol. 51, no. 5, pp. 3023–3036,
375 <http://doi.org/10.1109/TGRS.2012.2212201>

376 Jonathan Ho and Tim Salimans (2021). Classifier-free diffusion guidance. In *NeurIPS 2021*
377 *Workshop on Deep Generative Models and Downstream Applications*

378 Jonathan Ho, Ajay Jain, and Pieter Abbeel (2020). Denoising diffusion probabilistic models.
379 *Advances in Neural Information Processing Systems* 33, 6840–6851.
380 <https://doi.org/10.48550/arXiv.2006.11239>

381 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi and David
382 J. Fleet (2022). Video Diffusion Models. <https://doi.org/10.48550/arXiv.2204.03458>

383 Pulkkinen, S., Nerini, D., Pérez Hortal, A. A., Velasco-Forero, C., Seed, A., Germann, U., &
384 Foresti, L (2019). Pysteps: an open-source Python library for probabilistic precipitation
385 nowcasting (v1.0), *Geosci. Model Dev.*, 12, 4185–4219, [https://doi.org/10.5194/gmd-](https://doi.org/10.5194/gmd-12-4185-2019)
386 [12-4185-2019](https://doi.org/10.5194/gmd-12-4185-2019).

387 Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M.,
388 Athanassiadou, M., Kashem, S., Madge, S., Prudden, R., Mandhane, A., Clark, A.,
389 Brock, A., Simonyan, K., Hadsell, R., Robinson, N., Clancy, E., Arribas, A., and
390 Mohamed, S (2021). Skilful precipitation nowcasting using deep generative models of
391 radar, *Nature*, 597, 672–677, <https://doi.org/10.1038/s41586-021-03854-z>

392 Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image
393 synthesis with latent diffusion models. In Proceedings of the IEEE/CVF conference on
394 computer vision and pattern recognition (pp. 10684-10695).
395 <https://doi.org/10.48550/arXiv.2112.10752>

396 Shi, Xingjian, Chen, Zhourong, Wang, Hao, Yeung, Dit-Yan, Wong, Wai Kin & WOO, Wang-
397 chun. (2015). Convolutional LSTM Network: A Machine Learning Approach for
398 Precipitation Nowcasting.

399 Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., and Ganguli, S (2015). Deep
400 unsupervised learning using nonequilibrium thermodynamics.
401 <https://doi.org/10.48550/arXiv.1503.03585>

402 Song, J., Meng, C., and Ermon, S (2021). Denoising Diffusion Implicit Models. In International
403 Conference on Learning Representations. <https://doi.org/10.48550/arXiv.2010.02502>

404 Song, Y. and Ermon, S (2020). Generative modeling by estimating gradients of the data
405 distribution. <https://doi.org/10.48550/arXiv.1907.05600>

406 Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen and Andrea Dittadi (2022).
407 Diffusion Models for Video Prediction and Infilling.
408 <https://doi.org/10.48550/arXiv.2206.07696>

409 Vikram Voleti, Alexia Jolicoeur-Martineau and Christopher Pal (2022). MCVD: Masked
410 Conditional Video Diffusion for Prediction, Generation, and Interpolation.
411 <https://doi.org/10.48550/arXiv.2205.09853>

412 Xingjian Shi, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and
413 Wang-chun Woo. (2017). Deep learning for precipitation nowcasting: a benchmark and
414 a new model. In Advances in Neural Information Processing Systems vol. 30, 5622–
415 5632.

416 Zhang, Y., Long, M., Chen, K. et al. Skilful nowcasting of extreme precipitation with
417 NowcastNet. Nature 619, 526–532 (2023). <https://doi.org/10.1038/s41586-023-06184-4>

418 Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao and Bryan Catanzaro (2020). DiffWave: A
419 Versatile Diffusion Model for Audio Synthesis.
420 <https://doi.org/10.48550/arXiv.2009.09761>