

41 the upcoming hours. It allows people to plan their activities and make decisions based
42 on expected weather conditions. As we do not always have a whole picture of current
43 weather information, and cannot process this information in time, the task of
44 precipitation nowcasting is challenging. We take advantage of a novel machine
45 learning approach to learn what possible precipitation conditions are, given current
46 precipitation condition observed from radar. Our results offer accurate precipitation
47 prediction. More importantly, this method assigns high uncertainty to predictions
48 where predictions are more biased. This accurate estimate of prediction uncertainty is
49 crucial for weather related decision makings.

50

51 **1 Introduction**

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

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

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

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

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

78 here $\mu_{\theta}(x)$ serves as the deterministic forecast. This formulation comes with two
79 shortcomings. Firstly, it prohibits the exploration of the spatial structure of predictions,
80 making it difficult to leverage data-informed prior knowledge for achieving
81 structurally reasonable predictions. Secondly, it assumes a deterministic outcome,
82 despite the absence of a full-profile and strictly accurate initial state estimate. As a

83 result, deterministic models tend to yield poorly structured, blurry estimates, missing
84 extreme cases and uncertainty quantification. These deficiencies are evident in models
85 such as ConvLSTM, ConvGRU and Unet (Shi et al., 2015, 2017; Ayzel et al., 2019,
86 2020).

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

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

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

107 To address these challenges, we introduce diffusion models (Sohl-Dickstein et al.,
108 2015; Song & Ermon, 2020b; Ho et al., 2020) for precipitation nowcasting. Unlike
109 GANs, probabilistic diffusion models are likelihood-based generative models, that is,
110 they are trained to directly maximize the probability assigned to the observed
111 samples.

112 This enables a full coverage of the target distributions (Ali Razavi et al., 2019;
113 Dhariwal & Nichol, 2021). Moreover, their iterative generation nature allows us to
114 flexibly control the resulting distribution using initial state information. As a result,
115 we can gradually enhance the sharpness of the predictive distribution, while
116 guaranteeing the predictive distribution encapsulates all plausible outcomes.

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

124

125 **2 Methods**

126 **2.1 Probabilistic modeling the Precipitation nowcasting**

127 Consider a sequence of precipitation field data $\mathbf{R} = [r_1, r_2, \dots, r_M]$, the nowcasting
 128 task is to predict future precipitation field trajectories (N fields) based on a given past
 129 trajectory of observations (M fields). Here, we formulate this problem as a
 130 probabilistic machine learning task. Using an extensive dataset of sequences of
 131 precipitation field data, we learn conditional probability model of
 132 $P_\theta(\mathbf{R}_{M+1:M+N}|\mathbf{R}_{1:M})$, thus

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

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

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

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

139
 140 **2.2 Basic diffusion**

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

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

144 Here, $X_{t \in [1, T]}$ are latent variables. $0 < \beta_t < 1$ is diffusion coefficient. Given large
 145 enough T, $q(X_T|X_0)$ is close to standard Gaussian. Therefore, the forward Gaussian
 146 process maps any target distribution $P(X_0)$ to standard Gaussian. To approximate
 147 $P(X_0)$, starting from standard Gaussian, we sequentially reverse the Gaussian process
 148 using the following variational distributions:

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

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

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

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

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

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

157
 158 **2.3 Conditional diffusion**

159 Our objective is to approximate the conditional distribution of $P(X_t|y)$. Begin

160 with the score-based formulation of a diffusion model, the goal is to learn
 161 $\nabla \log P(X_t|y)$, by Bayes rules, we can get the equivalent:

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

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

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

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

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

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

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

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

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

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

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

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

177

178 3 Data

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

184 Due to computational resource limitations, we employed a subset comprising
 185 11,000 radar samples, partitioned into three subsets: training (8,000 samples),
 186 validation (2,000 samples), and testing (1,000 samples). The principal objective of
 187 this investigation is to assess the efficacy and reliability of diffusion-based and
 188 GAN-based models for precipitation nowcasting. To optimize resource usage, we
 189 exclusively evaluated these models for 30-minute precipitation predictions.
 190 Consequently, we performed random 80x80 sub-sample extractions from the original
 191 256x256-sized data to speed up training.

192

193 **4 Model Evaluation**

194 **4.1 Baseline models**

195 Generative models of radar (DGMR) holds the current state of the art in
196 precipitation nowcasting. We utilized Google-Colab to load the pre-trained DGMR
197 model and evaluate its performance using the first 30 minutes of forecasted results
198 (Ravuri et al., 2021). UNet serves as the baseline for deterministic forecasting using
199 deep learning (Ayzel et al., 2020). PySTEPS is a widely used precipitation nowcasting
200 system based on ensembles (Pulkkinen et al., 2019). We adopt PySTEPS as a
201 non-machine learning baseline. More details of the baseline can be found in the
202 support information.

203

204 **4.2 Evaluation strategy**

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

216

217 **5 Results and discussion**

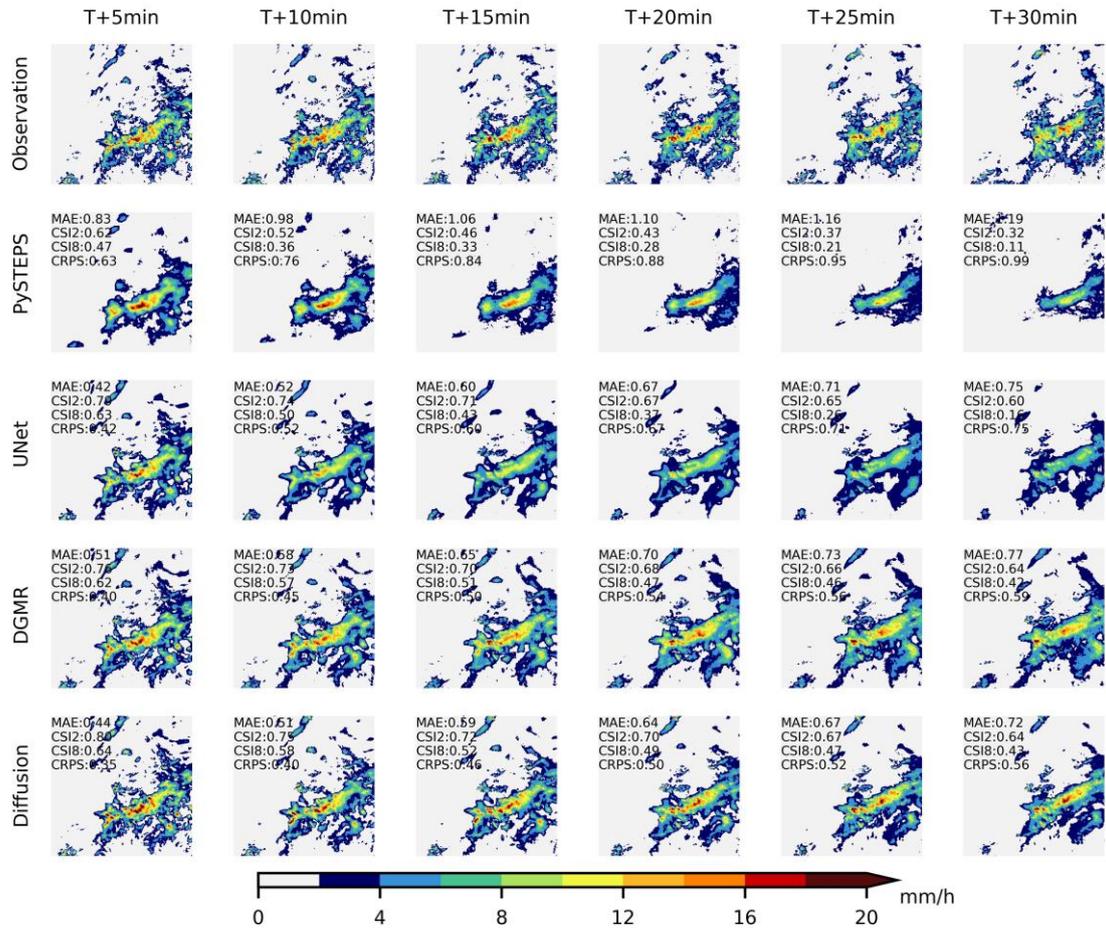
218 **5.1 Model performance for heavy precipitation forecasts**

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

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

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

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



236

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

241

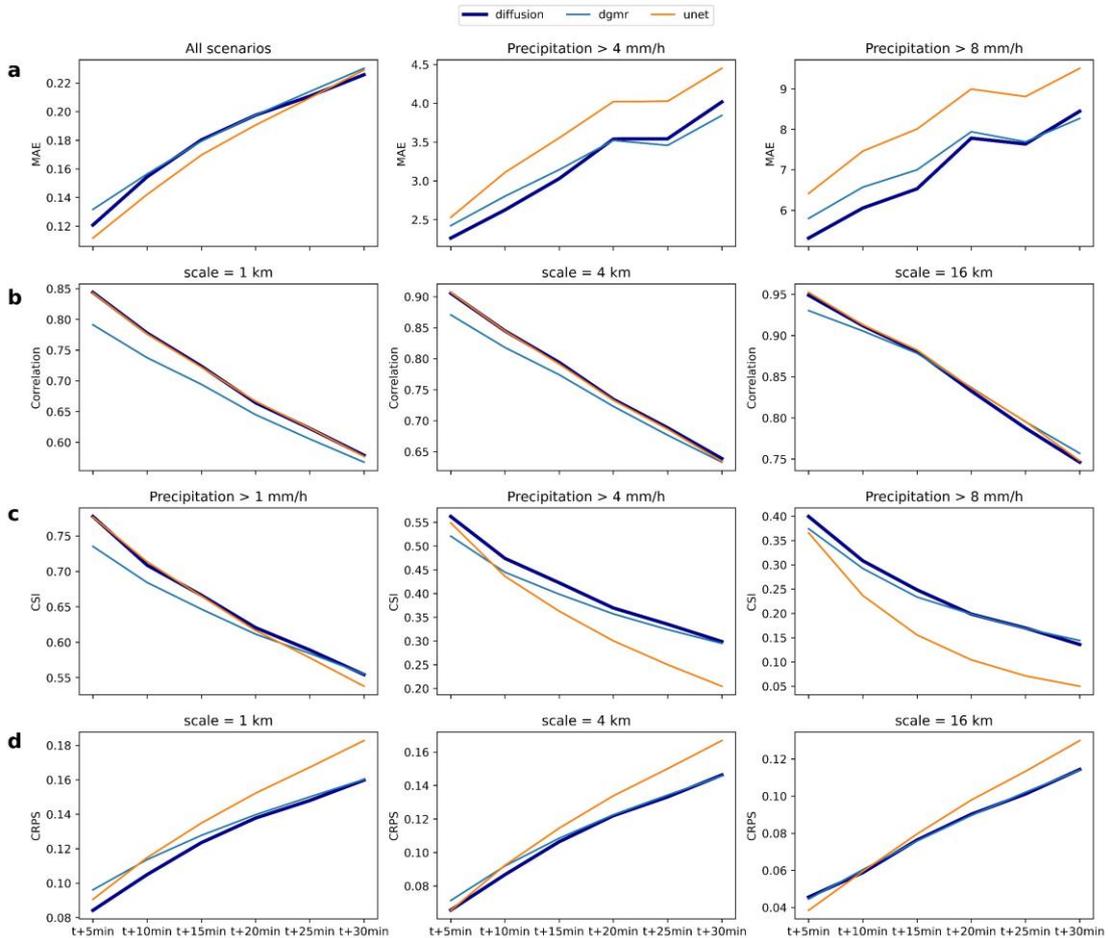
242 5.2 Forecast skill evaluation

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

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

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

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



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

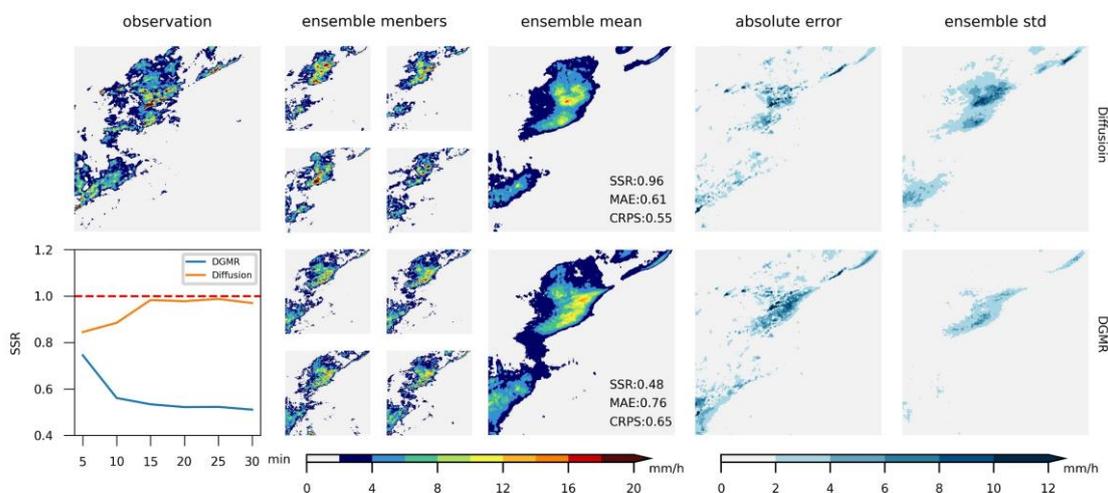
279

280 5.3 Reliability quantification

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

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

303 We also calculated SSR over the test dataset, displayed at the bottom left of
304 Figure 3. For DGMR, the SSR values are 0.745, 0.561, 0.534, 0.522, 0.523, and 0.511,
305 with an average of 0.56. In contrast, diffusion yields SSR values 0.845, 0.885, 0.983,
306 0.978, 0.988, and 0.970, with an average of 0.94. Diffusion model achieves a 68%
307 gain in the spread-skill ratio, underscoring its ability to provide more reliable
308 forecasts.



309

310 **Figure 3.** The example of ensemble forecasts provided by DGMR and Diffusion at

311 the thirtieth minute. From left to right: four randomly selected ensemble members, the
312 ensemble mean, the absolute error map comparing the ensemble mean to observations,
313 and the ensemble standard deviation. The bottom-left panel displays the reliability
314 quantification SSR (Spread-Skill Ratio) calculated using the entire test dataset for
315 forecasts.

316

317 **6 Conclusions**

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

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

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

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

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

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

354

355 **Data availability**

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

357

358 **Acknowledgements**

359 This research was supported by the Third Xinjiang Scientific Expedition Program
360 (Grant No. 2021xjkk0806), the National Natural Science Foundation of China
361 (42271032, U2243226, 42288101), NSFC-DFG mobility (M-0468), Chinese
362 Academy of Science Light of the West Interdisciplinary Research Grant (grant no.
363 xbzg-zdsys-202104) and National Key R&D Program of China 2021YFA0718000.

364

365 **Reference**

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

369 Ayzel, G., Scheffer, T., and Heistermann, M (2020). RainNet v1.0: a convolutional neural
370 network for radar-based precipitation nowcasting, Geosci. Model Dev., 13, 2631–
371 2644, <https://doi.org/10.5194/gmd-13-2631-2020>

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

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

377 Gneiting, T., Balabdaoui, F., and Raftery, A. E. (2007). Probabilistic forecasts, calibration and
378 sharpness. Journal of the Royal Statistical Society Series B: Statistical Methodology,
379 69(2), 243-268. <https://doi.org/10.1111/j.1467-9868.2007.00587.x>

380 Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville,
381 A., and Bengio, Y. (2014). Generative adversarial nets. Advances in neural
382 information processing systems, 27. <https://doi.org/10.48550/arXiv.1406.2661>

383 Ho, J. and Salimans, T. (2021). Classifier-free diffusion guidance. In NeurIPS 2021 Workshop
384 on Deep Generative Models and Downstream Applications.
385 <https://doi.org/10.48550/arXiv.2207.12598>

386 Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in
387 Neural Information Processing Systems 33, 6840-6851.
388 <https://doi.org/10.48550/arXiv.2006.11239>

389 Ho, J., Salimans, T., Gritsenko, A., Chan, W, Norouzi, M., and Fleet, D. J. (2022). Video
390 Diffusion Models. <https://doi.org/10.48550/arXiv.2204.03458>

391 Höpfe, T., Mehrjou, A., Bauer, S., Nielsen, D. and Dittadi, A. (2022). Diffusion Models for
392 Video Prediction and Infilling. <https://doi.org/10.48550/arXiv.2206.07696>

393 Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. (2020). DiffWave: A Versatile
394 Diffusion Model for Audio Synthesis. <https://doi.org/10.48550/arXiv.2009.09761>

395 Pulkkinen, S., Nerini, D., Pérez Hortal, A. A., Velasco-Forero, C., Seed, A., Germann, U., and
396 Foresti, L (2019). Pysteps: an open-source Python library for probabilistic
397 precipitation nowcasting (v1.0), Geosci. Model Dev., 12, 4185–4219,
398 <https://doi.org/10.5194/gmd-12-4185-2019>.

399 Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M.,
400 Athanassiadou, M., Kashem, S., Madge, S., Prudden, R., Mandhane, A., Clark, A.,
401 Brock, A., Simonyan, K., Hadsell, R., Robinson, N., Clancy, E., Arribas, A., and
402 Mohamed, S (2021). Skilful precipitation nowcasting using deep generative models
403 of radar, *Nature*, 597, 672–677, <https://doi.org/10.1038/s41586-021-03854-z>
404 Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution
405 image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF*
406 *conference on computer vision and pattern recognition* (pp. 10684-10695).
407 <https://doi.org/10.48550/arXiv.2112.10752>
408 Sakaino, H. (2013). Spatio-Temporal Image Pattern Prediction Method Based on a Physical
409 Model With Time-Varying Optical Flow, in *IEEE Transactions on Geoscience and*
410 *Remote Sensing*, 51 (5), 3023-3036, <http://doi.org/10.1109/TGRS.2012.2212201>
411 Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W. K. and WOO, W. (2015). Convolutional
412 LSTM Network: A Machine Learning Approach for Precipitation Nowcasting.
413 Shi, X., Gao, Z., Lausen, L., Wang, H., Yeung, D.-Y., Wong, W.-K., and Woo, W.-C.. (2017).
414 Deep learning for precipitation nowcasting: a benchmark and a new model. In
415 *Advances in Neural Information Processing Systems* vol. 30, 5622–5632.
416 Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., and Ganguli, S. (2015). Deep
417 unsupervised learning using nonequilibrium thermodynamics.
418 <https://doi.org/10.48550/arXiv.1503.03585>
419 Song, J., Meng, C., and Ermon, S (2021). Denoising Diffusion Implicit Models. In
420 *International Conference on Learning Representations*.
421 <https://doi.org/10.48550/arXiv.2010.02502>
422 Song, Y. and Ermon, S (2020). Generative modeling by estimating gradients of the data
423 distribution. <https://doi.org/10.48550/arXiv.1907.05600>
424 Voleti, V., Jolicœur-Martineau, A., and Pal, C. (2022). MCVD: Masked Conditional Video
425 Diffusion for Prediction, Generation, and Interpolation.
426 <https://doi.org/10.48550/arXiv.2205.09853>
427 Zhang, Y., Long, M., Chen, K. et al. Skilful nowcasting of extreme precipitation with
428 NowcastNet. *Nature* 619, 526–532 (2023). <https://doi.org/10.1038/s41586-023-06184-4>