

Supporting Information for “Reliable precipitation nowcasting using probabilistic diffusion model”

Congyi Nai^{1,3}, Baoxiang Pan², Jiarui Hai⁴, Xi Chen², Qihong Tang^{1,3}, Guangheng Ni⁴,
Qingyun Duan⁵, Bo Lu⁶, Ziniu Xiao², Xingcai Liu^{1,3,*}

¹ Key Laboratory of Water Cycle and Related Land Surface Processes, Institute of Geographic
Sciences and Natural Resources Research, Chinese Academy of Sciences, Beijing, China

² Institute of Atmospheric physics, Chinese Academy of Sciences, Beijing, China

³ University of Chinese Academy of Sciences, Beijing, China

⁴ State Key Laboratory of Hydro-science and Engineering, Department of Hydraulic
Engineering, Tsinghua University, Beijing 100084, China

⁵The National Key Laboratory of Water Disaster Prevention, Hohai University, Nanjing,
China

⁶Laboratory for Climate Studies and CMA-NJU Joint Laboratory for Climate Prediction
Studies, National Climate Center, China Meteorological Administration, Beijing, China

*Corresponding author. E-mail address: xingcailiu@igsrr.ac.cn

1 Details of diffusion model

1.1 Basic diffusion

Let x_0 be a sample from the data distribution $q(x_0)$, and defines a sequence of increasingly noisy versions of x which we call the latent variables x_t ($t = 1 \dots T$) through the forward diffusion process, described by

$$q(x_t|x_{t-1}) = N(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I) \quad (1)$$

Then, the form of $q(x_t|x_0)$ can be recursively derived through repeated applications of the reparameterization trick, suppose we have $\{\epsilon_t, \bar{\epsilon}_t\}_{t=0}^T \sim N(0, I)$, Then, for an arbitrary sample $x_t \sim q(x_t|x_0)$, we can rewrite it as:

$$\begin{aligned} x_t &= \sqrt{1 - \beta_t}x_{t-1} + \sqrt{\beta_t}\epsilon_{t-1} \\ &= \sqrt{\alpha_t}(\sqrt{\alpha_{t-1}}x_{t-2} + \sqrt{1 - \alpha_{t-1}}\epsilon_{t-2}) + \sqrt{1 - \alpha_t}\epsilon_{t-1} \\ &= \sqrt{\alpha_t\alpha_{t-1}}x_{t-2} + \sqrt{\alpha_t - \alpha_t\alpha_{t-1}}\epsilon_{t-2} + \sqrt{1 - \alpha_t}\epsilon_{t-1} \\ &= \sqrt{\alpha_t\alpha_{t-1}}x_{t-2} + \sqrt{1 - \alpha_t\alpha_{t-1}}\bar{\epsilon}_{t-2} \\ &= \dots \\ &= \sqrt{\prod_{i=1}^t \alpha_i} x_0 + \sqrt{1 - \prod_{i=1}^t \alpha_i} \bar{\epsilon}_0 \\ &= \sqrt{\bar{\alpha}_t} X_0 + \sqrt{1 - \bar{\alpha}_t} \bar{\epsilon}_0, \text{ where } \bar{\alpha}_t = \prod_{i=1}^t \alpha_i \end{aligned} \quad (2)$$

In equation 3, we have leveraged the property that the sum of two independent Gaussian random variables retains a Gaussian distribution, with the mean being the sum of the two individual means and the variance being the sum of their variances. $\sqrt{\alpha_t - \alpha_t\alpha_{t-1}}\epsilon_{t-2}$ is a sample from Gaussian $N(0, (\alpha_t - \alpha_t\alpha_{t-1})I)$, $\sqrt{1 - \alpha_t}\epsilon_{t-1}$ is a sample from Gaussian $N(0, (1 - \alpha_t)I)$, we can then treat their sum as a random variable sampled from Gaussian $N(0, (1 - \alpha_t + \alpha_t - \alpha_t\alpha_{t-1})I)$. Hence, the X_t can be sampled directly from X_0 , the transition kernel is

$$q(x_t|x_0) = N(x_t; \sqrt{\bar{\alpha}_t}x_0, \sqrt{1 - \bar{\alpha}_t}I), \text{ where } \alpha_t = 1 - \beta_t, \bar{\alpha}_t = \prod_{i=1}^t \alpha_i. \quad (3)$$

38 Given X_0 and a Gaussian vector $\epsilon \sim N(0, I)$ and applying the transformation
 $X_t = \sqrt{\bar{\alpha}_t}X_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon.$ (4)

39 When the $\bar{\alpha}_T \rightarrow 0$, X_T is well approximated by Gaussian distribution. During the
40 forward process, noise is gradually added to the data until it loses its original spatial
41 structure characteristics and becomes pure noise. If we can solve the reverse process
42 $P(X_{t-1}|X_t)$, we can sample $X_T \sim N(0, I)$ then a sequence of neural networks is employed to
43 gradually reduce the noise in a series of steps $X_T, X_{T-1} \dots X_0$. These properties suggest
44 learning a learnable Markov chain model $P_\theta(X_{t-1}|X_t)$ to approximate the true reverse
45 process:

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

46 Therefore, in a diffusion model, we are only interested in learning conditionals
47 $P_\theta(X_{t-1}|X_t)$, the diffusion model can be optimized by maximizing the variational lower
48 bound (VLB) of the log-likelihood of the data X_0 ,

$$\begin{aligned} 49 \mathbb{E}_{q(X_0)}(-\log P_\theta(X_0)) &\leq \mathbb{E}_{q(X_0)}[-\log P_\theta(X_0) + D_{KL}(q(X_{1:T}|X_0)||P_\theta(X_{1:T}|X_0))] \\ 50 &= \mathbb{E}_{q(X_0)}[-\log P_\theta(X_0) + \int q(X_{1:T}|X_0) \log \frac{q(X_{1:T}|X_0)}{P_\theta(X_{0:T})/P_\theta(X_0)} dX_{1:T}] \\ 51 &= \mathbb{E}_{q(X_0)}[-\log P_\theta(X_0) + \int q(X_{1:T}|X_0) \log \frac{q(X_{1:T}|X_0)}{P_\theta(X_{0:T})} dX_{1:T} + \log P_\theta(X_0)] \\ &= \mathbb{E}_{q(X_{0:T})} \log \frac{q(X_{1:T}|X_0)}{P_\theta(X_{0:T})} = L_{VLB} \end{aligned} \quad (6)$$

52 We can rewrite variational lower bound (VLB) as,

$$\begin{aligned} 53 L_{VLB} &= \mathbb{E}_{q(x_0^T)} [\log \frac{q(x_{1:T}|x_0)}{p_\theta(x_{0:T})}] \\ 54 &= \mathbb{E}_q [\log \frac{\prod_{t=1}^T q(x_t|x_{t-1})}{p_\theta(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t)}] \\ 55 &= \mathbb{E}_q [-\log p_\theta(x_T) + \sum_{t=1}^T \log \frac{q(x_t|x_{t-1})}{p_\theta(x_{t-1}|x_t)}] \\ 56 &= \mathbb{E}_q [-\log p_\theta(x_T) + \sum_{t=2}^T \log \frac{q(x_t|x_{t-1})}{p_\theta(x_{t-1}|x_t)} + \log \frac{q(x_1|x_0)}{p_\theta(x_0|x_1)}] \\ 57 &= \mathbb{E}_q [-\log p_\theta(x_T) + \sum_{t=2}^T \log \left(\frac{q(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} \cdot \frac{q(x_t|x_0)}{q(x_{t-1}|x_0)} \right) + \log \frac{q(x_1|x_0)}{p_\theta(x_0|x_1)}] \\ 58 &= \mathbb{E}_q [-\log p_\theta(x_T) + \sum_{t=2}^T \log \frac{q(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} + \sum_{t=2}^T \log \frac{q(x_t|x_0)}{q(x_{t-1}|x_0)} + \log \frac{q(x_1|x_0)}{p_\theta(x_0|x_1)}] \\ 59 &= \mathbb{E}_q [-\log p_\theta(x_T) + \sum_{t=2}^T \log \frac{q(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} + \log \frac{q(x_T|x_0)}{q(x_1|x_0)} + \log \frac{q(x_1|x_0)}{p_\theta(x_0|x_1)}] \\ 60 &= \mathbb{E}_q [\log \frac{q(x_T|x_0)}{p_\theta(x_T)} + \sum_{t=2}^T \log \frac{q(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} - \log p_\theta(X_0|X_1)] \\ &= \mathbb{E}_q [D_{KL}(q(X_T|X_0)||p_\theta(X_T)) + \sum_{t=2}^T D_{KL}(q(X_{t-1}|X_t, X_0)||p_\theta(X_{t-1}|X_t) - \log p_\theta(X_0|X_1))] \end{aligned} \quad (7)$$

61 This formulation also has an elegant interpretation, which is revealed when
62 inspecting each individual term:

- 63 1. $L_0 = \mathbb{E}_q[\log p_\theta(X_0|X_1)]$ can be interpreted as a reconstruction term.
- 64 2. $L_T = \mathbb{E}_q[D_{KL}(q(X_T|X_0)||p_\theta(X_T))]$ represents how close the distribution of the final
65 noisified input is to the standard Gaussian prior, is equal to zero under our

66 assumptions.

67 3. $L_t = \mathbb{E}_q[\sum_{t=2}^T D_{KL}(q(X_{t-1}|X_t, X_0)||p_\theta(X_{t-1}|X_t))]$ is a denoising matching term. The
 68 $q(X_{t-1}|X_t, X_0)$ acts as a ground-truth signal and $p_\theta(X_{t-1}|X_t)$ is our desired denoising
 69 transition step. This term is therefore minimized when the two denoising steps
 70 match as closely as possible. It is the primary optimization objective.

71 If we have knowledge of X_0 , we can obtain $q(X_{t-1}|X_t, X_0)$ through the Bayes'
 72 theorem,

$$\begin{aligned}
 73 \quad q(X_{t-1}|X_t, X_0) &= q(X_t|X_{t-1}, X_0) \frac{q(X_{t-1}|X_0)}{q(X_t|X_0)} \\
 74 \quad &\propto \exp\left(-\frac{1}{2}\left(\frac{(X_t - \sqrt{\alpha_t}X_{t-1})^2}{\beta_t} + \frac{(X_{t-1} - \sqrt{\bar{\alpha}_{t-1}}X_0)^2}{1 - \bar{\alpha}_{t-1}} - \frac{(X_t - \sqrt{\bar{\alpha}_t}X_0)^2}{1 - \bar{\alpha}_t}\right)\right) \\
 75 \quad &= \exp\left(-\frac{1}{2}\left(\left(\frac{\alpha_t}{\beta_t} + \frac{1}{1 - \bar{\alpha}_{t-1}}\right)X_{t-1}^2 + \left(\frac{2\sqrt{\alpha_t}}{\beta_t}X_t + \frac{2\sqrt{\bar{\alpha}_{t-1}}}{1 - \bar{\alpha}_{t-1}}X_0\right)X_{t-1} + C(X_t, X_0)\right)\right) \\
 &= N(X_{t-1}; \tilde{\mu}(X_t, X_0), \tilde{\beta}(t)I) \tag{8}
 \end{aligned}$$

76 Recall equation 8 and equation 5, we can obtain,

$$77 \quad \tilde{\mu}_\theta(X_t, X_0) = \frac{1}{\sqrt{\alpha_t}}\left(X_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(X_t, t)\right) \tag{9}$$

77 Let us consider the $L_t = D_{KL}(q(X_{t-1}|X_t, X_0)||p_\theta(X_{t-1}|X_t))$, given equation 6 and
 78 equation 8, we can get the loss function,

$$\begin{aligned}
 79 \quad L_t &= \mathbb{E}_{x_0, \epsilon} \left[\frac{1}{2\|\Sigma_\theta(x_t, t)\|_2^2} \|\tilde{\mu}_t(x_t, x_0) - \mu_\theta(x_t, t)\|_2^2 \right] \\
 80 \quad &= \mathbb{E}_{x_0, \epsilon} \left[\frac{1}{2\|\Sigma_\theta\|_2^2} \left\| \frac{1}{\sqrt{\alpha_t}}\left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_t\right) - \frac{1}{\sqrt{\alpha_t}}\left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(x_t, t)\right) \right\|^2 \right] \\
 81 \quad &= \mathbb{E}_{x_0, \epsilon} \left[\frac{(1 - \alpha_t)^2}{2\alpha_t(1 - \bar{\alpha}_t)\|\Sigma_\theta\|_2^2} \|\epsilon_t - \epsilon_\theta(x_t, t)\|_2^2 \right] \\
 &= \mathbb{E}_{x_0, \epsilon} \left[\frac{(1 - \alpha_t)^2}{2\alpha_t(1 - \bar{\alpha}_t)\|\Sigma_\theta\|_2^2} \|\epsilon_t - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_t, t)\|_2^2 \right] \tag{10}
 \end{aligned}$$

82 Ho et al. (2020) propose to reweight various terms in L_{VLB} for better sample
 83 quality, to compute this objective, we generate samples $X_t \sim q(X_t|X_0)$, then train a model
 84 ϵ_θ to predict the added noise using a standard mean-squared error loss:

$$L_{simple} = \mathbb{E}_{t \sim [1, T], X_0 \sim q(X_0), \epsilon \sim N(0, I)} [\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_t, t)\|^2]. \tag{11}$$

85 2.3 Conditional diffusion

86 So far, we have focused on modeling the data distribution $p(x)$. However, we are
 87 often also interested in the conditional distribution of $P(X_t|y)$, as it enables us to better
 88 investigate how different conditional information influences the generation of variable
 89 X . Begin with the score-based formulation of a diffusion model, the goal is to learn
 90 $\nabla \log P(X_t|y)$, by Bayes rules, we can get the equivalent:

$$91 \quad \nabla \log P(X_t|y) = \nabla \log \left(\frac{P(y|X_t)P(X_t)}{P(y)} \right) \tag{12}$$

$$= \nabla \log P(X_t) + \nabla \log P(y|X_t) - \nabla \log P(y) \tag{13}$$

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

To better control the conditional information, a hyperparameter γ is introduced to

92 scale the gradient of the conditioning information. The score function can then be
 93 summarized as:

$$\nabla \log P(X_t|y) = \nabla \log P(X_t) + \gamma \nabla \log P(y|X_t). \quad (15)$$

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

$$\nabla \log P(y|X_t) = \nabla \log P(X_t|y) - \nabla \log P(X_t). \quad (16)$$

99 Substituting equation (16) into equation (15) then we get:

$$\nabla \log P(X_t|y) = \nabla \log P(X_t) + \gamma(\nabla \log P(X_t|y) - \nabla \log P(X_t)). \quad (17)$$

$$= \underbrace{(1 - \gamma)\nabla \log P(X_t)}_{\text{unconditional score}} + \underbrace{\gamma\nabla \log P(X_t|y)}_{\text{conditional score}} \quad (18)$$

100 From Tweedie’s formula and equation 5, we can get,

$$\nabla \log p(x_t) = -\frac{1}{\sqrt{1-\bar{\alpha}_t}} \epsilon \quad (19)$$

101 The equation 19 means that estimating ϵ is equivalent to estimating a scaled
 102 version of the score function. So, in this paper, we model the conditional distribution of
 103 precipitation frames in the future given the past precipitation frames $\mathbf{P} = [p_1, p_2, \dots, p_M]$, we
 104 learn two sets of neural networks, $\epsilon_\theta(X_t, t)$ and $\epsilon_\theta(X_t, t, P)$, to approximate the
 105 unconditional and conditional score functions $\nabla \log P(X_t)$ and $\nabla \log P(X_t|y)$, our
 106 conditional diffusion loss function is:

$$L_{\text{condition}} = \mathbb{E}_{t \sim [1, T], X_0 \sim q(X_0), \epsilon \sim N(0, I)} [||\epsilon - \epsilon_\theta(X_t, t, P)||^2]. \quad (19)$$

107

108 2 Details of baseline model

109 2.1 Generative models of radar

110 DGMR holds the current state of the art in precipitation nowcasting, the generator
 111 is built with convolutional and convolutional GRU layers and it was trained with two
 112 adversarial loss and a regularization loss. The first loss is defined by a spatial
 113 discriminator, which ensures spatial consistency. The second loss is defined by a
 114 temporal discriminator, which is a 3D convolutional neural network that aims to impose
 115 temporal consistency. The regularization term encourages the prediction’s mean
 116 precipitation fields to match the mean of past precipitation amount.

117 We utilized Google-Colab to load the saved DGMR model and pconducted
 118 inference on our test dataset, see [https://github.com/deepmind/deepmind-
 119 research/tree/master/nowcasting](https://github.com/deepmind/deepmind-research/tree/master/nowcasting). DGMR exhibits the capability to generate forecasts
 120 up to 90 minutes. However, for the purpose of comparison, we only evaluated its
 121 performance using the first 30 minutes of forecasted results, calculating relevant
 122 metrics.

123 2.2 U-Net

124 We use a U-Net encoder–decoder model as baseline similarly to how it was used
 125 in related studies (Ayzel et al., 2020). This type of model first employs an encoder that
 126 reduces the spatial resolution using pooling and convolutional layers, while the decoder
 127 then increases the resolution by applying up-sampling and convolutional layers to the

128 learned patterns. To prevent gradient vanishing and share the low-level patterns of the
 129 precipitation fields, skip connections are used from the encoder to the decoder
 130 (Srivastava et al., 2015). In this paper, U-Net serves as the baseline for deterministic
 131 forecasting using deep learning.

132 2.3 PySTEPS

133 PySTEPS is an open-source Python library designed for radar precipitation
 134 forecasting and analysis, it is available at <https://github.com/pySTEPS/pysteps>. It offers
 135 a comprehensive range of algorithms, among which STEPS is a widely used
 136 precipitation nowcasting system based on ensembles, considered to be state-of-the-art
 137 of non-ML-based method. In this study, we adopt PySTEPS as a non-machine learning
 138 baseline.

139

140 3 Details of metrics

141 we use the M to denote number of the ensemble members, and f_m to denote the
 142 ensemble member, so the ensemble mean can be written as,

$$\bar{f} = \frac{1}{M} \sum_{m=1}^M f_m \quad (20)$$

143 3.1 MAE

144 The (spatial) mean-absolute-error (MAE) at forecast time step t between ensemble
 145 means \bar{f} and observation f_{obs} is defined as,

$$MAE_t(\bar{f}, f_{obs}) = \frac{1}{p} \sum_{p=1}^p |\bar{f} - f_{obs}| \quad (21)$$

146 where p indexes all the geospatial locations. And we can consider extreme value
 147 prediction accuracy under different precipitation intensities, we use an intensity mask
 148 $[f_{obs} > 4]$ and $[f_{obs} > 8]$ to get the masked prediction and observation \bar{f}_m, f_{m_obs}

$$MAE_{t,mid}(\bar{f}_m, f_{m_obs}) = \frac{1}{p} \sum_{p=1}^p |\bar{f}_m - f_{m_obs}| \quad (22)$$

149

150 3.2 Correlation

151 The spital correlation between ensemble mean and observation is defined as.

$$Corr_t(\bar{f}, f_{obs}) = \frac{\sum_p(\bar{f}_p - \bar{\bar{f}}_p)(f_{obs,p} - \bar{\bar{f}}_{obs,p})}{\sqrt{\sum_p(\bar{f}_p - \bar{\bar{f}}_p)^2} \sqrt{\sum_p(f_{obs,p} - \bar{\bar{f}}_{obs,p})^2}} \quad (23)$$

152 where $\bar{\bar{f}}_p$ means to average in space. In deployment, we flatten the prediction and
 153 observation then use the *corrcoef* function from the *NumPy* library.

154

155 3.3 Critical Success Index

156 The Critical Success Index (CSI) is a statistical measure that quantifies the
 157 accuracy of spatial prediction by evaluating the correct identification of specific events
 158 or outcomes.

159 The CSI is defined as the ratio of true positives (TP) to the sum of true positives,
 160 false positives (FP), and false negatives (FN). Mathematically, it is expressed as,

$$CSI = \frac{TP}{TP+FP+FN} \quad (24)$$

- 161 ● TP represents the number of true positive outcomes, which signifies the accurate
 162 prediction of events or occurrences.
 163 ● FP corresponds to false positives, indicating instances where the event was
 164 predicted, but did not materialize.
 165 ● FN denotes false negatives, signifying cases where the event occurred but was not
 166 correctly predicted.

167 The CSI values range between 0 and 1, where a CSI of 1 indicates perfect spatial
 168 accuracy in prediction, implying that all positive outcomes were correctly forecasted
 169 without any false alarms. Conversely, a CSI of 0 suggests that none of the events were
 170 accurately predicted.

171

172 **3.4 Continuous Ranked Probability Score**

173 CRPS is used to evaluate the calibration and sharpness. It quantifies the
 174 discrepancy between the forecasted cumulative distribution function (CDF) and the
 175 observed CDF, defined as,

$$CRPS = \int_{-\infty}^{+\infty} [F(f_m) - 1(t \leq z)]^2 dz \quad (25)$$

176 where F denotes the CDF of the prediction distribution and $1(t \leq z)$ is an indicator
 177 function that is 1 if $t \leq z$ and 0 otherwise. In the case of a deterministic forecast (like
 178 Unet) the CRPS reduces to the mean absolute error (MAE).

179

180 **3.5 Spread-skill ratio**

181 The SSR evaluates the reliability of the ensemble. It is a ratio that quantifies the
 182 balance between calibration and sharpness, providing insights into the trade-off
 183 between these two critical aspects of predictive modeling.

$$SSR = \frac{Spread}{RMSE} \quad (26)$$

184 where the spread is defined as,

$$Spread = \sqrt{\frac{1}{P} \sum_{p=1}^P Var(f_{m,p})} \quad (27)$$

185 and the RMSE is defined as,

$$RMSE = \sqrt{\frac{1}{P} \sum_{p=1}^P (\bar{f} - f_{obser})^2} \quad (28)$$

186

187 **4 Additional results**

188 **4.1 Skill evaluation**

189 Figure S1 includes PySTEPS metrics calculated over the entire test dataset. Due
 190 to UNet's blurred predictions, it falls short of PySTEPS in terms of CSI8.

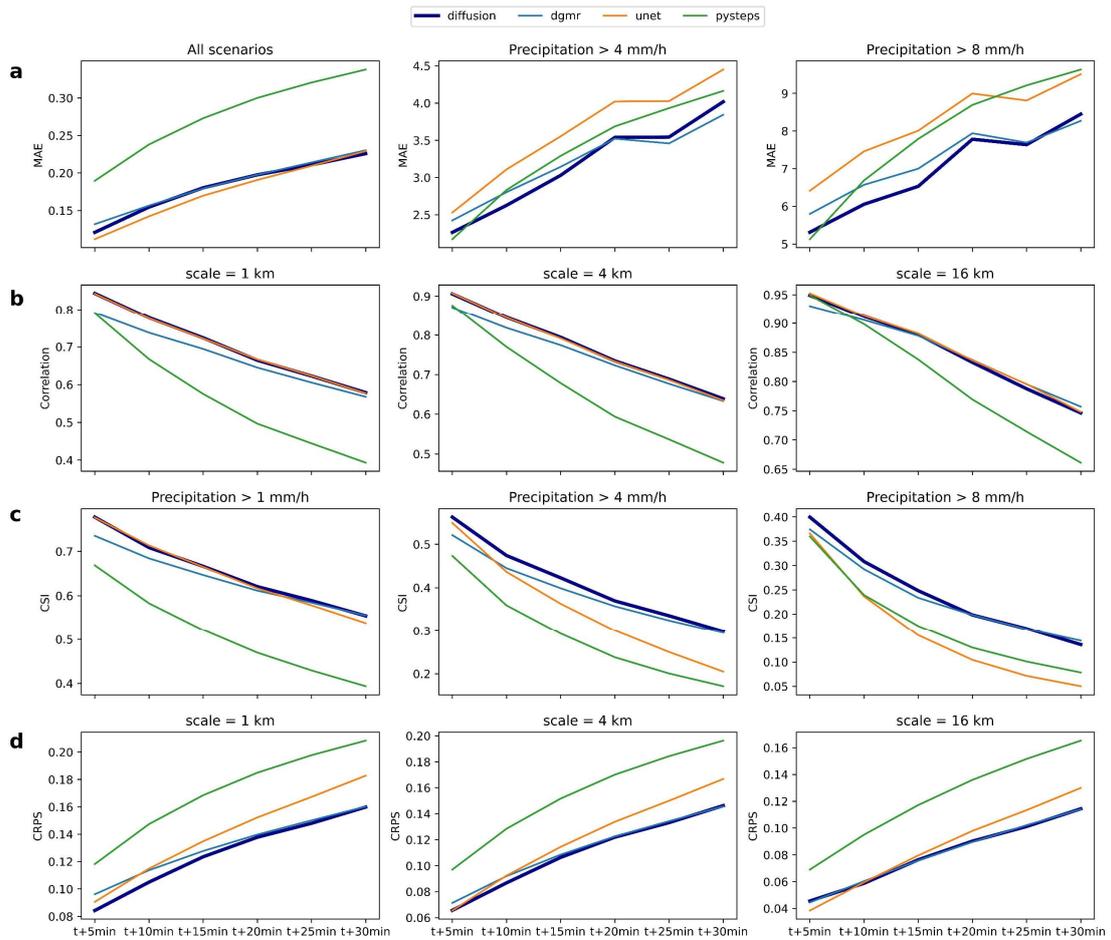
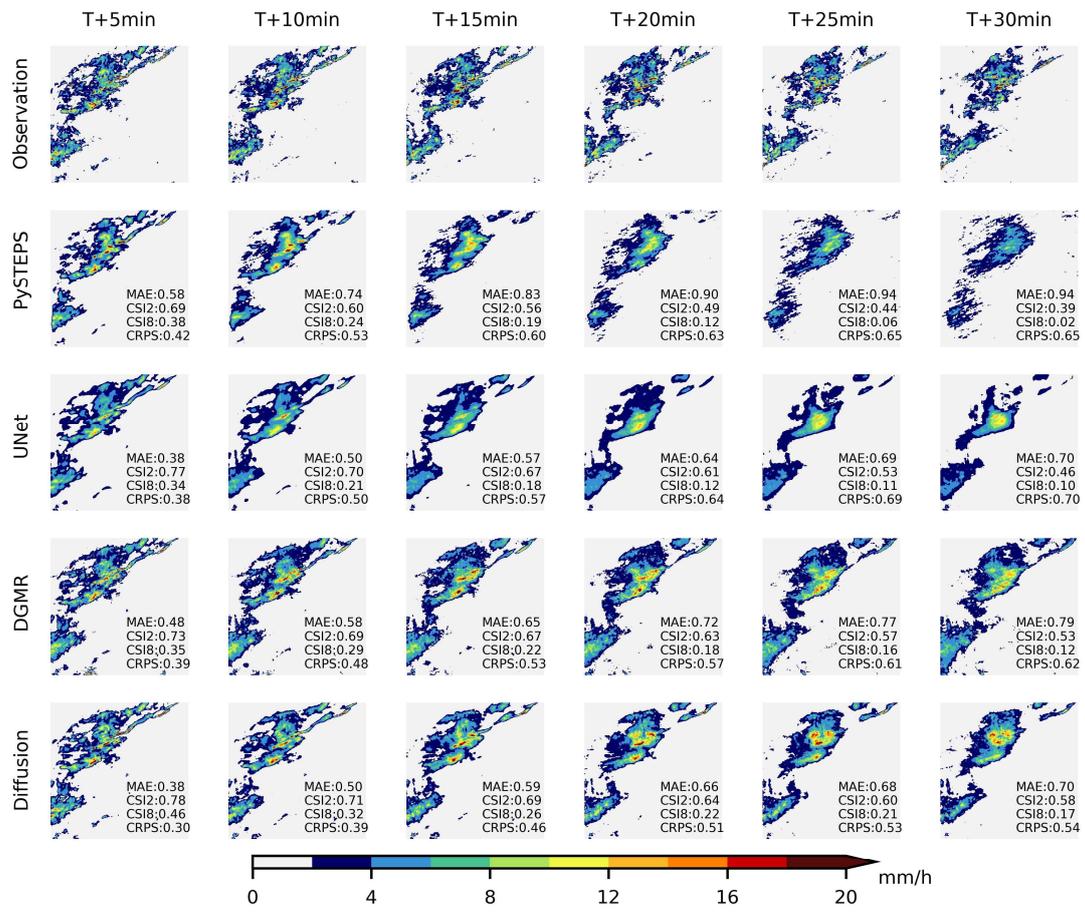


Figure S1. The full version of Figure 2, incorporating metrics for PySTEPS.

191
 192
 193
 194
 195
 196
 197
 198
 199
 200
 201
 202
 203
 204
 205
 206
 207
 208
 209
 210
 211
 212

213

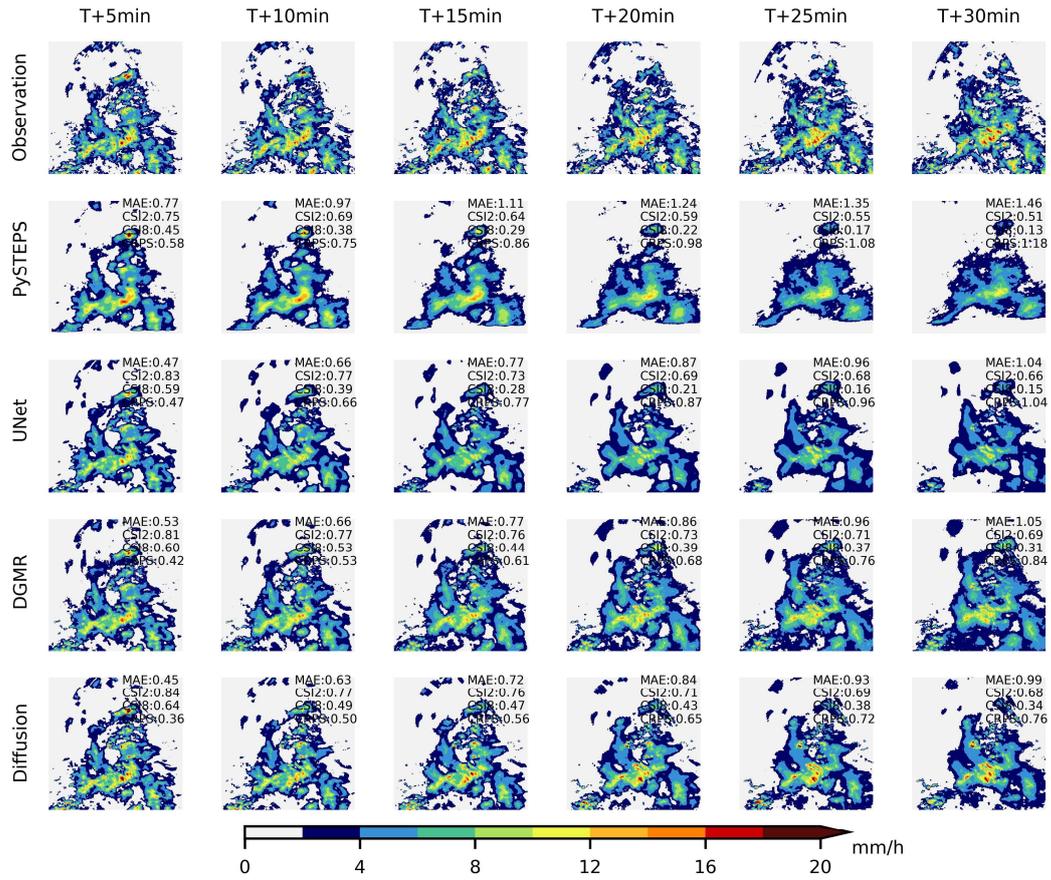
214 **4.2 Additional case**



215

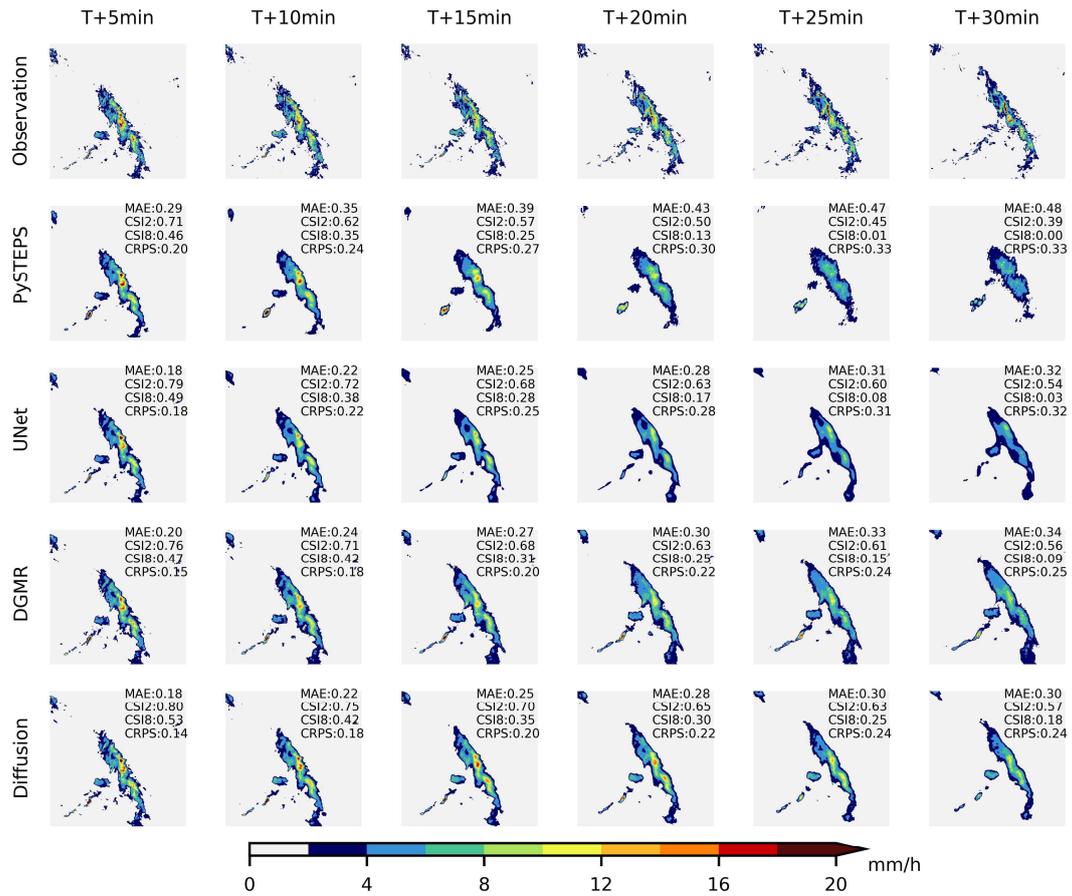
216 **Figure S2.** An additional case in Section 5.1

217



218
219
220

Figure S3. An additional case in Section 5.1

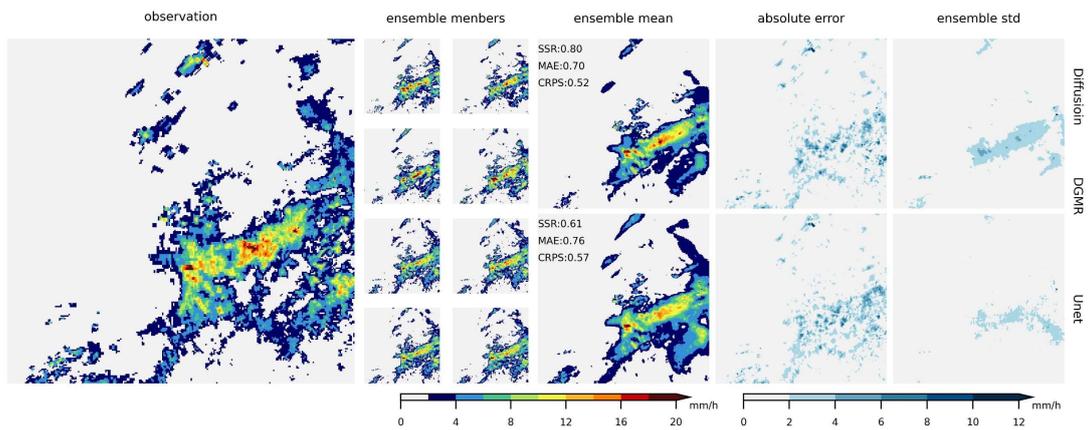


221

222 **Figure S4.** An additional case in Section 5.1

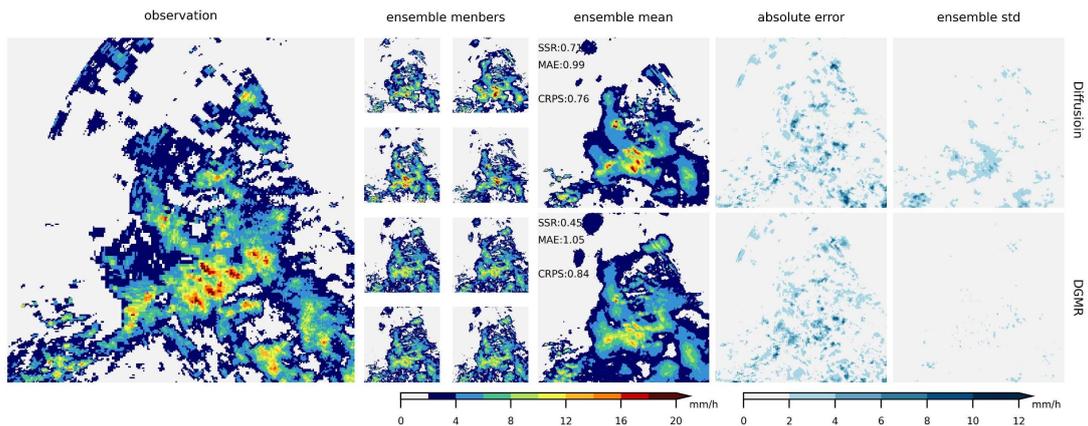
223

224 **4.3 Reliability cases**



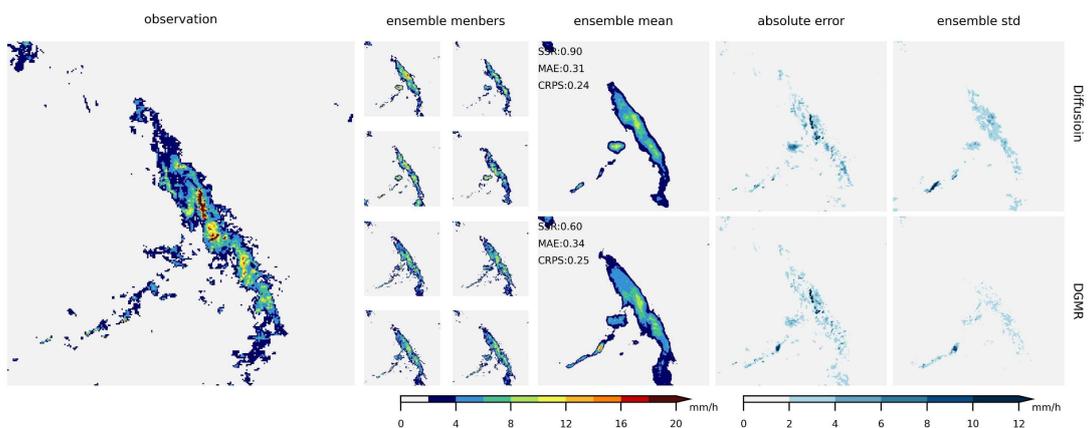
225

226 **Figure S5.** An additional case in Section 5.3



227
228

Figure S6. An additional case in Section 5.3



229

230

Figure S7. An additional case in Section 5.3

231

232 **Reference**

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

236 "Understanding diffusion models: A unified perspective." arXiv preprint arXiv:2208.11970
237 (2022). <https://doi.org/10.48550/arXiv.2208.11970>

238 Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., ... & Yang, M. H. (2022). Diffusion
239 models: A comprehensive survey of methods and applications. ACM Computing
240 Surveys. <https://doi.org/10.48550/arXiv.2209.00796>