

Precipitation Nowcasting by a Spectral-Based Nonlinear Stochastic Model

Original

Precipitation Nowcasting by a Spectral-Based Nonlinear Stochastic Model / Metta, S., von Hardenberg, J., Ferraris, L., Rebora, N., Provenzale, A.. - In: JOURNAL OF HYDROMETEOROLOGY. - ISSN 1525-755X. - 10:5(2009), pp. 1285-1297. [10.1175/2009JHM1120.1]

Availability:

This version is available at: 11583/2815026 since: 2020-04-22T14:21:56Z

Publisher:

AMER METEOROLOGICAL SOC

Published

DOI:10.1175/2009JHM1120.1

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Precipitation Nowcasting by a Spectral-Based Nonlinear Stochastic Model

SABINO METTA

ISAC-CNR, Turin, Italy

JOST VON HARDENBERG*

ISAC-CNR, Lecce, Italy

LUCA FERRARIS AND NICOLA REBORA

CIMA Research Foundation, Savona, Italy

ANTONELLO PROVENZALE

ISAC-CNR, Turin, Italy

(Manuscript received 24 October 2008, in final form 17 April 2009)

ABSTRACT

A novel rainfall nowcasting method based on the combination of an empirical nonlinear transformation of measured precipitation fields and the stochastic evolution in spectral space of the transformed fields is introduced. The power spectrum and the amplitude distribution of precipitation are kept constant during the forecast, and a Langevin-type model is used to evolve the Fourier phases. The application of the method to a study case is illustrated, and it is shown that, with this procedure, a forecast skill can be obtained that is superior to those provided by Eulerian or Lagrangian persistence for a lead time of up to two hours.

1. Introduction

Short-range forecasting (nowcasting) of intense precipitation events has many applications in hydrometeorological risk management, including flash-flood warning and surface or air traffic control. Extensive weather radar networks are now operational in several countries, providing detailed snapshots of local precipitation conditions at high spatial and temporal resolution. The availability of these data allows for the development of nowcasting techniques that use radar fields as initial conditions and forecast local precipitation up to a lead time of about one hour (Wilson et al. 1998, 2004).

Some of the nowcasting procedures are deterministic in nature, providing one single “best,” forecast (Dixon and

Wiener 1993). However, probabilistic forecasts are essential to quantify prediction uncertainties (Krzysztofowicz 2001). Physically based atmospheric models can be used for this purpose, but they usually require large computational times and detailed boundary conditions that may not be readily available in real-time applications. For this reason, most operational nowcasting techniques rely upon some form of stochastic modeling approach. Some of the existing methods are based on Bayesian approaches (Xu and Chandrasekar 2005; Fox and Wikle 2005). Other methods stochastically evolve individual rain cells or the entire precipitation field (Andersson and Ivarsson 1991; Mellor et al. 2000; Grecu and Krajewski 2000; Seed 2003; Xu and Chandrasekar 2005).

An important point in building a nowcasting procedure is that predictability of precipitation at short temporal scales depends on the spatial scale of the structures considered (Wilson et al. 1998; Germann and Zawadzki 2002). This has motivated the development of stochastic nowcasting methods operating in the Fourier domain, such as Spectral Prognosis (S-PROG; Seed 2003) or the

* Current affiliation: ISAC-CNR, Turin, Italy.

Corresponding author address: Antonello Provenzale, Institute of Atmospheric Sciences and Climate (ISAC-CNR), Corso Fiume 4, I-10133 Turin, Italy.
E-mail: a.provenzale@isac.cnr.it

method proposed by Xu and Chandrasekar (2005). In these techniques, the spectral representation allows for naturally taking into account the scale dependence of the statistical properties.

In addition to scale-dependent predictability, there are other important properties of spatiotemporal precipitation fields that need to be taken into account: (i) both the amplitude distribution and the correlation structure of the fields—that is, their power spectra—are often only weakly varying on the nowcasting time scale (Zawadzki 1973) and (ii) precipitation structures are often persistent in a Lagrangian framework, leading to a good nowcasting skill of simple Lagrangian persistence methods applied either to the entire field or to individual precipitation structures (Germann and Zawadzki 2004).

These two requirements lead us to the development of a new spectral-based nowcasting procedure, based on the empirical nonlinear transformation of precipitation fields provided by radar measurements and the stochastic evolution of the transformed fields in spectral space. In this approach, the initial one-point distribution and power spectrum of the precipitation field are kept constant, and a stochastic Ornstein–Uhlenbeck process is used for the time evolution of the Fourier phases of the Gaussianized precipitation field. As we will show, this procedure is able to provide an ensemble, probabilistic nowcasting of precipitation fields up to a lead time of two hours. The method introduced here automatically includes large-scale advection of precipitation structures, and it reproduces the nonlinear and intermittent nature of rain fields. In addition, the use of spectral space instead of physical space assures that the spatial correlations of precipitation fields are preserved.

The rest of this paper proceeds as follows: we introduce the “phase stochastic” (PhaSt) nowcasting technique and its implementation details in section 2 and illustrate its application to an intense precipitation event observed by a meteorological radar in Italy in section 3. We conclude in section 4 with a summary and some outlook.

2. Nowcasting in Fourier space

a. Model definition

The PhaSt model adopted here requires two initial precipitation fields, to be used as initial conditions, and it includes three main steps.

- We take an empirical nonlinear transformation of the two precipitation fields used as initial conditions, $p(x, y, t = 0)$ and $p(x, y, t = -\Delta t)$, generating two Gaussian fields, $g(x, y, 0)$ and $g(x, y, -\Delta t)$. The

transformation G is defined by rank ordering the values of p and substituting them with rank-ordered values obtained by sampling a Gaussian distribution (Schreiber and Schmitz 1996; Ferraris et al. 2003a). Here, x and y are spatial coordinates, t is time, $t = 0$ is the time when the nowcasting starts, Δt is the radar sampling time, p is precipitation, and g is the Gaussian field obtained by the nonlinear transformation $g = G(p)$.

- We take the Fourier transform of the Gaussianized fields and obtain their Fourier spectra, $\hat{g}(k_x, k_y, 0)$ and $\hat{g}(k_x, k_y, -\Delta t)$, signified by the caret. From these, we obtain for each wavenumber (k_x, k_y) the Fourier phase, ϕ , and an estimate of the Fourier angular frequency, herein defined as $\omega = d\phi/dt$. Fourier phases are then evolved in time by a stochastic process while Fourier amplitudes are kept fixed. The spectrum with the evolved Fourier phases is inverted to generate a nowcasted Gaussian field at the time t of interest, $g(x, y, t)$. This evolved field has the same power spectrum as the initial Gaussianized field, $g(x, y, 0)$. Different realizations of the stochastic process allow for generating different evolutions of the precipitation field and for creating an ensemble of precipitation nowcasts.
- We apply an inverse nonlinear transformation to pass from the evolved Gaussian field $g(x, y, t)$ to the nowcasted precipitation field, $\tilde{p}(x, y, t) = \tilde{G}^{-1}[g(x, y, t)]$. The nonlinear transformation is again defined empirically by rank ordering the values of $g(x, y, t)$ and then substituting them with the rank-ordered values of the initial field $p(x, y, 0)$. In this way, the nowcasted precipitation field $\tilde{p}(x, y, t)$ has exactly the same one-point amplitude distribution as the initial field. Note that the inverse transformation \tilde{G}^{-1} is not exactly the inverse of G , although in practice it is rather close. Because the power spectrum is not invariant for a nonlinear transformation (Balmforth et al. 1999), the power spectrum of \tilde{p} is very similar to that of $p(x, y, 0)$ but not exactly the same.

There are several stochastic models that can be used to evolve the Fourier phases. The simplest option is to use a random walk—that is, to assume the angular frequencies $\omega = d\phi/dt$ to be Gaussian and uncorrelated random variables. To allow for the presence of time correlations in the angular frequencies, we can resort to a Langevin-type model. This type of stochastic model has been widely used in the turbulence literature as a simple closure model for isotropic turbulence (Herring and Kraichnan 1972) to model dispersion (van Dop et al. 1985) and to create synthetic turbulent velocity fields (Kraichnan 1970; Fung and Vassilicos 1998). The random walk behavior is recovered in the limit of vanishing correlation time.

In the Langevin-type model adopted here, the temporal evolution of the Fourier phase $\phi(k_x, k_y)$ at a given wavenumber (k_x, k_y) is written in terms of a linear Ornstein–Uhlenbeck stochastic process for the angular frequency:

$$d\phi(k_x, k_y, t) = \omega(k_x, k_y, t) dt \quad \text{and} \quad (1)$$

$$d\omega(k_x, k_y, t) = \frac{\omega'(k_x, k_y) - \omega(k_x, k_y, t)}{T_c(k_x, k_y)} dt + \sqrt{\frac{2\sigma^2(k_x, k_y)}{T_c(k_x, k_y)}} dW(k_x, k_y, t). \quad (2)$$

Here, $\omega(k_x, k_y, t)$ is the angular frequency; $\omega'(k_x, k_y)$ is a relaxation frequency to be discussed below; $dW(k_x, k_y, t)$ is a random increment drawn from a normal distribution with zero mean and second-order moment $\langle dW(k_x, k_y, t) dW(k'_x, k'_y, t') \rangle = \delta(k_x - k'_x) \delta(k_y - k'_y) \delta(t - t')$, where W is a Wiener process; $T_c(k_x, k_y)$ is the correlation time; and $\sigma^2(k_x, k_y)$ is the variance of the distribution of angular frequencies at wavenumber (k_x, k_y) . The first term on the right-hand side of Eq. (2) represents a deterministic relaxation on the time scale T_c to the reference angular frequency ω' , and the second term represents the stochastic driver. The Ornstein–Uhlenbeck process [Eqs. (1) and (2)] generates angular frequencies that have a Gaussian distribution with zero mean and variance σ^2 and an exponentially decaying temporal autocorrelation, $C(k_x, k_y, \tau) = \langle \omega(k_x, k_y, t) \omega(k_x, k_y, t + \tau) \rangle_t / \sigma^2(k_x, k_y) = \exp[-\tau/T_c(k_x, k_y)]$, where $\langle \cdot \rangle_t$ indicates a time average. In this simple model, all Fourier phases are assumed to evolve independently.

The choice of the reference angular frequency, $\omega'(k_x, k_y)$, can either be $\omega' = 0$ or, as we discuss below, $\omega'(k_x, k_y) = \omega(k_x, k_y, t = 0)$. The latter choice allows the model to reproduce the temporal evolution encoded in the initial fields. In fact, by choosing $\omega' = \omega(t = 0)$ and $\sigma^2 = 0$ from Eq. (2), we obtain $d\omega = 0$ and $\omega(t) = \omega' = \omega(t = 0)$ at all times. Large-scale advection with rigid translation of the initial rain field corresponds to angular frequencies of the form $\omega_{\text{adv}}(k_x, k_y) = c_{0,x}k_x + c_{0,y}k_y$, where $c_{0,x}$ and $c_{0,y}$ are the constant phase speeds along the x and y directions. If the initial fields provide a frequency distribution with this type of dependence on wavenumber, then large-scale advection is automatically encoded in the PhaSt method: By determining $\omega'(k_x, k_y)$ from the two initial frames, we obtain the properties of the large-scale advection field, similar to what could be recovered by maximum correlation techniques (Zawadzki et al. 1994). Alternatively, if one had reasons to force the presence of

large-scale advection in the temporal evolution of the rain fields, then one could choose the relaxation frequencies accordingly—that is, $\omega'(k_x, k_y) = \omega_{\text{adv}}(k_x, k_y)$. Clearly, if advection is too strong, precipitation structures are advected too quickly and then the model cannot reproduce the event. In particular, as discussed in quantitative terms in the section on parameter estimate, Fourier phases should vary much less than 2π in one time step—that is, $c_{0,x}$ and $c_{0,y}$ should not be too large for a given radar sampling time.

On the one hand, rain fields have been shown to be characterized by a nonlinear and intermittent nature, often associated with the presence of multifractal behavior and interpreted in terms of multiplicative cascades (Lovejoy and Mandelbrot 1985; Schertzer and Lovejoy 1987; Gupta and Waymire 1993; Kumar and Foufoula-Georgiou 1993). On the other hand, Ferraris et al. (2003a) have shown that the multifractality of measured rain fields cannot be distinguished from that generated by the static nonlinear transformation of a linear Gaussian process. In a further study, Ferraris et al. (2003b) have shown that three main classes of rainfall models—namely, (i) multifractal cascades, (ii) the superposition of rain cells, and (iii) the nonlinear transformation of a linear Gaussian field—all display similar performance in reproducing the observed precipitation data. Rebola et al. (2006) then developed a downscaling model based on the static nonlinear transformation of a Gaussian field with a power-law spectrum and showed that it correctly reproduces the statistics of high-resolution precipitation fields measured by meteorological radars. The model introduced here is an extension of that type of approach and, consistent with the analysis results mentioned above, is based on the assumption that the main nonlinearity of spatial–temporal precipitation dynamics can be approximated by the static nonlinear transformations G and \tilde{G}^{-1} , and that the residual correlations between the different Fourier phases of the evolving Gaussianized fields remain absent or very weak.

The use of a stochastic process for the evolution of Fourier phases allows for generating many realizations, to be used as members of an ensemble of precipitation nowcasts. All ensemble members are characterized by the same amplitude distribution and very similar power spectra. However, the phase evolution (i.e., the positioning of rainfall structures) evolves differently in the different realizations, providing an estimate of the probability of occurrence of precipitation at a given point in space and a given moment in time. At each wavenumber, the spread in the phase evolution is controlled by the value of σ , and the relaxation rate to the reference phase ω' is controlled by the value of T_c .

b. Model implementation

In the model implementation, we first Gaussianize, by rank ordering, the two precipitation fields used as initial conditions at $t = 0$ and $t = -\Delta t$. A delicate point is how to treat the zeros (pixels with no precipitation) possibly existing in the initial fields. Here we opt for assigning all points with a precipitation value below a given minimum threshold to the same value of the transformed variable. This leads to a truncated Gaussian distribution and does not significantly modify the resulting power spectrum.

The spectral amplitudes and the initial spectral phases, $\phi(k_x, k_y, 0)$, are determined from the Gaussianized precipitation field at the initial time $t = 0$. First-order finite differences with the phases at the previous time step, $t = -\Delta t$ (this is why we need at least two initial precipitation fields), provide the initial conditions for the Fourier angular frequencies, $\omega(k_x, k_y, 0) = [\phi(k_x, k_y, 0) - \phi(k_x, k_y, -\Delta t)]/\Delta t$.

The Ornstein–Uhlenbeck model used to evolve the Fourier phases can be discretized in time as

$$\omega_{t+\Delta t} = \omega_t \left(1 - \frac{\Delta t}{T}\right) + \frac{\omega'}{T} \Delta t + \sqrt{\frac{2\sigma^2 \Delta t}{T}} \left(1 - \frac{\Delta t}{2T}\right) W_t$$

and

$$\phi_{t+\Delta t} = \phi_t + \omega_{t+\Delta t} \Delta t, \quad (4)$$

where W_t is drawn from a Gaussian distribution with zero mean and unit variance. The variance of the angular frequencies has been normalized to match the variance of the continuous case, Eq. (2). The correlation time in the continuous case T_c is related to that of the discretized model T by

$$T = \frac{\Delta t}{1 - \exp(-\Delta t/T_c)}. \quad (5)$$

When $T_c \gg \Delta t$, one has $T \approx T_c$. For $T = \Delta t$, this model generates temporally uncorrelated angular frequencies and a Brownian random walk of Fourier phases. Note that when $\omega' = \omega_0$ and $\sigma^2 = 0$, one obtains $\omega_t = \omega_0$ at all times; when $\sigma^2 \neq 0$, the ensemble average of the set of stochastic realizations corresponds to the evolution obtained by keeping constant the initial frequencies.

Because all Fourier phases of the Gaussianized fields evolve independently of each other, they quickly lose any initial cross correlation (if it existed). This ensures that the fields generated by the inverse Fourier transform of the evolved spectra remain Gaussian, in keeping with the assumption that the main nonlinearity is associated with the static nonlinear transformation G used to

Gaussianize the field. The empirical inverse transformation \tilde{G}^{-1} is finally implemented by rank ordering again: we sort both the nowcasted Gaussian field values and the initial precipitation intensities, and we substitute each nowcasted value with the value in the initial field having the same rank. This operation only moderately modifies the correlation structure of the fields (Balmforth et al. 1999), and all nowcasted fields have power spectra that remain very close to that of the initial precipitation field.

Two-dimensional Fourier transforms assume a double periodic domain, so that the nowcasted fields can display precipitation structures falling across the borders. To avoid this unphysical effect, one should always apply the PhaSt technique to a larger domain, padding the initial observation fields with zeros, and then extracting the area corresponding to the original domain from the nowcasting outputs.

c. Parameter estimate

Before applying the PhaSt method to real precipitation fields, we explore whether and how the model parameters can be recovered from measured data. To this end, we generate a sequence of synthetic precipitation fields with an a priori choice of model parameters to try to recover their value from the analysis of the data.

The synthetic precipitation fields have spatial resolution $\Delta x = 1$ km and temporal resolution $\Delta t = 10$ min. We generate a total number of 18 fields, corresponding to an event with a total duration of three hours. The fields are generated on a large domain and then cut on an area of 128×128 km² to avoid spurious spatial periodicity. We consider a case with a constant value of the correlation time, $T = 1.5\Delta t$, and a value of σ that is linearly growing with wavenumber, $\sigma = c_o k$, where $k = (k_x^2 + k_y^2)^{1/2}$ is the radial wavenumber and $c_o = 0.15$ km min⁻¹. The initial radial power spectrum is fixed as $P(k) \propto k^{-2.7}$, isotropic in wavenumber space and consistent with the analysis of precipitation data (Ferraris et al. 2003a,b). Here, $P(k)$ indicates the average of the squared Fourier amplitudes on the radial shell $(k, k + \Delta k)$, and the total variance of the field is thus $\sum k P(k) \Delta k$. The fields generated in this way are spatially isotropic (in a statistical sense). The Gaussian fields generated by inverting the Fourier spectra at different times are then passed through an exponential filter to generate synthetic precipitation fields.

We then apply the PhaSt method to the synthetic fields. After padding the fields with zeroes to a larger domain of 192×192 km², the first step is to estimate angular frequencies from finite differences of Fourier phases. Because phases are defined on the periodic domain $[0, 2\pi]$, an “unwrapping” procedure is required




FIG. 1. (a) Standard deviation of the angular frequencies reconstructed from the synthetic fields generated to test the PhaSt procedure, as a function of radial wavenumber k . (b) Correlation time T as a function of radial wavenumber k , estimated from the data by fitting an exponentially decaying function to the autocorrelation of the angular frequencies. The bars represent one standard deviation around the mean, computed from the distribution of angular frequencies on circular shells of width Δk .

before computing finite differences: when an estimated Fourier phase increases by more than π in a time interval Δt , we subtract 2π from it; when the phase jump is smaller than $-\pi$, we add 2π . This procedure works well when angular frequencies are small and Fourier phases evolve slowly, something that usually happens at low wavenumbers. When the angular frequency is large, the corresponding Fourier phase can change by more than π in the time interval Δt . In this case, the whole unwrapping procedure ceases to work and angular frequency estimates become ambiguous.

From the reconstructed angular frequencies, we can obtain estimates of σ and T . Figure 1a shows the estimate of the standard deviation of the angular frequencies recovered from the Gaussianized fields $\sigma(k)$ as a function of the radial wavenumber k , obtained by averaging the variances over circular shells of width $\Delta k = 2\pi/L$ and then taking the square root. Figure 1b shows the estimate of the correlation time of the angular frequencies of the Gaussianized fields, again obtained as an average over circular shells of width Δk . Each value of $T(k_x, k_y)$ has been obtained by fitting a decaying exponential to the temporal autocorrelation of the time series of the corresponding angular frequency and then using Eq. (5) that relates T_c and T . Similar results are obtained by first averaging the autocorrelation functions for each wavenumber in the shell and then fitting the mean autocorrelation to an exponential.

An approximate linear dependence of σ on k is evident up to moderate wavenumbers. At higher wavenumbers, the quantity $\sigma\Delta t$ saturates at the upper limit

$\pi/\sqrt{3}$, owing to the fact that the phases are defined modulus 2π . When the phases change by more than π in a time step Δt , the distribution becomes randomized and the angular frequencies cannot be properly reconstructed. As a result, one gets an upper limiting value for $\sigma(k)$. The value $\pi/\sqrt{3}$ is the value of $\sigma\Delta t$ corresponding to a random uniform distribution of Fourier phases. Experimentation with other parameter values shows that when c_o is large enough, there is always a value of k beyond which the reconstructed value of the variance saturates. In real situations, the wavenumber at which this saturation occurs depends on how rapidly the precipitation field evolves and on the sampling time Δt .

As shown by Fig. 1b, at small k the value of $T(k)$ provides an estimate of the correlation time, even though a constant value is not recovered. The estimated value of T rapidly decays to Δt for growing k . This confirms that, at large k , Fourier phases undergo a Brownian motion and indicates that the value of T should be estimated from the small wavenumbers.

3. Application to a study case

To illustrate the workings of the PhaSt method, we now apply this nowcasting technique to one example of a typical precipitation event observed by the meteorological radar on Mt. Settepani, Liguria, Italy. The synoptic scenario of this event shows a deep low-pressure area extending between northern France and northern Africa and a southern warm moist flow over Italy, leading to intense precipitation events in northwestern Italy.

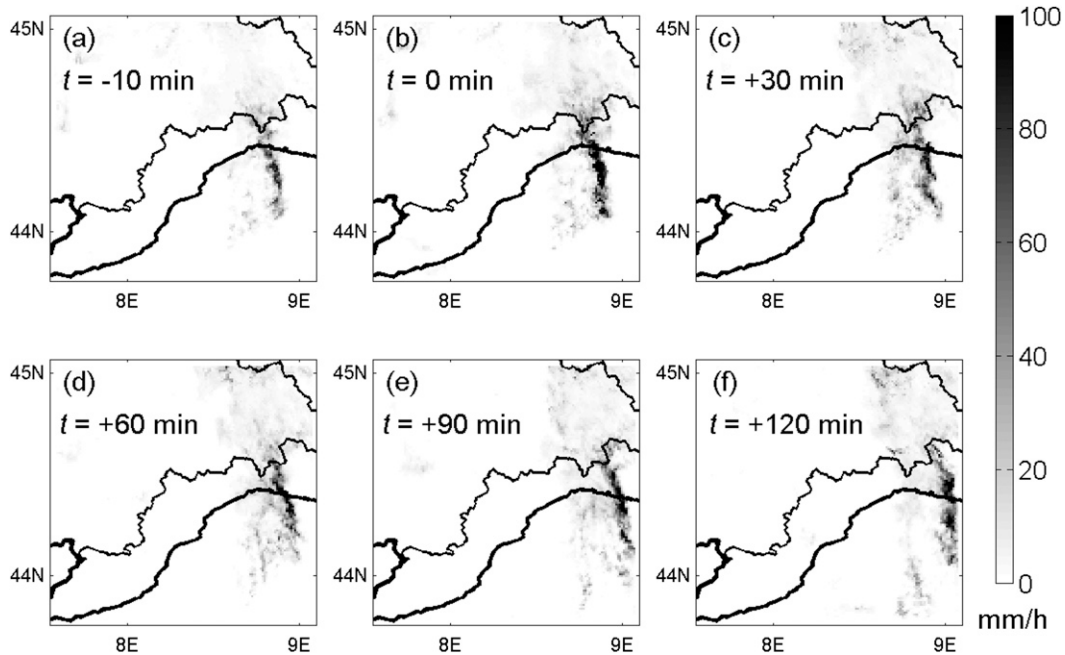


FIG. 2. Sequence of precipitation intensity fields, accumulated over 10 min, observed by the meteorological radar on Mt. Settepani, Italy. The startup time of the nowcasting procedure, $t = 0$, is fixed at 0730 UTC 24 Nov 2003.

The radar frames analyzed here cover a domain of $128 \times 128 \text{ km}^2$, with resolution $\Delta x = 1 \text{ km}$. Different scans are obtained with temporal resolution $\Delta t = 10 \text{ min}$, starting at 0720 UTC 24 November 2003 and continuing for more than two hours, providing a total of 14 consecutive precipitation fields (Fig. 2). The particular sequence analyzed here is an example of an intense, localized precipitation pattern, characterized by significant wind advection.

a. Verification of model assumptions

First, we explore whether the statistical properties of the precipitation fields in the event used as an example are consistent with the assumptions of the PhaSt method. Figures 3a and 3b show the power spectrum and the amplitude distribution of the precipitation field at different times. Both statistics remain approximately




FIG. 3. (a) Spatial power spectra, as a function of radial wavenumber, of the precipitation fields. (b) Distributions of precipitation intensity for the event shown in Fig. 2. The spectra of the Gaussianized fields are very similar. Different curves refer to different times from $t = 0$ to $t = 12\Delta t$.




FIG. 4. (a) Standard deviation of reconstructed angular frequencies as a function of radial wavenumber k , for the event shown in Fig. 2. The inset shows the reconstructed field $\sigma(k_x, k_y)$. (b) Same analysis as in Fig. 1b, for the event shown in Fig. 2.

constant in time, consistent with the assumptions of the nowcasting model introduced here and the observations of Zawadzki (1973).

Then, to proceed with model verification, we first pad the original field with zeros, obtaining a field with resolution $\Delta x = 1$ km and total size 192×192 km². We then proceed with the empirical Gaussianization of the fields and obtain a temporal sequence of angular frequencies, $\omega(k_x, k_y, t)$, by finite differences of the spectral phases of the Gaussianized fields. From these, the variance $\sigma^2(k_x, k_y)$ of the angular frequency distribution at each wavenumber can be estimated. By computing the autocorrelation of $\omega(k_x, k_y, t)$ at different time lags and fitting a decaying exponential function, we estimate the correlation time $T_c(k_x, k_y)$ and then, using Eq. (5), the value of $T(k_x, k_y)$. At this point, one can either keep the information on potential anisotropies in the x and y directions, at the cost of a limited statistics in the estimates of σ and T , or average over isotropic wavenumber shells at radial wavenumber k and width Δk . In this way, the information on anisotropy is lost but a more robust estimate of model parameters is obtained.

Figure 4a reports the standard deviation of the angular frequencies $\sigma(k)$ as a function of the radial wavenumber k , obtained by averaging the variances over circular shells of width Δk and then taking the square root. An approximate linear dependence of σ on k is evident up to moderate wavenumbers. Consistent with the results discussed earlier, the quantity $\sigma\Delta t$ saturates at the upper limit $\pi/\sqrt{3}$ at higher wavenumbers. For this event and $\Delta t = 10$ min, saturation occurs at a wavenumber k corresponding to length scale $L_{\text{sat}} = 2\pi/k_{\text{sat}} \approx 13$ km.

Notice, however, that the complete reconstructed field $\sigma(k_x, k_y)$, shown in the inset of Fig. 4a, displays clear anisotropy. In this case, the use of an isotropic form $\sigma(k)$ would represent only a coarse approximation to the full structure of $\sigma(k_x, k_y)$.

Figure 4b shows the average correlation time $T(k)$, again averaged over radial shells of width Δk . At small values of k , the value of $T(k)$ indicates that the angular frequencies have significant temporal correlation. At larger values of k , the value of $T(k)$ decays to the value $T = \Delta t$, corresponding to uncorrelated frequencies. The length scale L_{sat} obtained from the saturation of $\sigma(k)$ is consistent with the value beyond which temporal correlations in angular frequencies vanish.

Overall, Figs. 3 and 4 show that the precipitation statistics of this test case are consistent with the assumptions of the PhaSt nowcasting model—that is, temporally correlated angular frequencies, a variance growing with wavenumber, and an amplitude distribution and a power spectrum approximately constant in time.

b. Hindcast

We now apply the PhaSt method to the test case discussed earlier, assuming in-sample knowledge of model parameters. The standard deviation $\sigma(k_x, k_y)$ and the correlation time $T(k_x, k_y)$ at each wavenumber (k_x, k_y) are obtained from the whole dataset. We determine the reference angular frequencies $\omega'(k_x, k_y)$ as the time average of the angular frequencies obtained from the analysis of the entire time series $\omega'(k_x, k_y) = \langle \omega(k_x, k_y, t) \rangle_t$. No assumption of isotropy in spectral space is made.




FIG. 5. Sequence of precipitation intensity fields, accumulated over 10 min, nowcasted by the PhaSt method assuming in-sample knowledge of model parameters.

Figure 5 shows a single stochastic realization obtained from the nowcasting model. Comparison with Fig. 2 indicates that the PhaSt approach is capable of reproducing a realistic temporal evolution of the precipitation field and that it captures the eastward advection of the high-intensity precipitation structures. By construction, this sequence of nowcasted fields is characterized by the same amplitude distribution and very similar power spectrum at all times. Figure 6 shows the power spectrum of the nowcasted fields at different times.

To quantitatively compare the nowcasted fields with the observed sequence, we show in Fig. 7 the average cross correlation ρ between nowcasted and observed fields, as a function of forecast time. The graph shows the cross correlations for an ensemble of 100 stochastic realizations of the PhaSt method, corresponding to different choices of the evolution of Fourier phases, and the ensemble mean. The figure also shows the cross correlations obtained from three other nowcasting methods: 1) simple Eulerian persistence, where the field at time $t = 0$ is assumed to persist; 2) Lagrangian per-

sistence, where the average advection velocity is estimated from the first two radar images of the sequence by using a maximum correlation technique (Zawadzki et al. 1994); and 3) random nowcasting fields having the same amplitude distribution and correlation structure of the observed precipitation field at $t = 0$ but characterized by a random, uniform distribution of the spectral phases of the Gaussianized fields (this generates a random positioning of precipitation structures with no temporal correlation). For the latter case, the figure reports the limit below which 95% of the 3000 random realizations fall.

As shown in Fig. 7, both Eulerian and Lagrangian persistence display a rapidly decreasing correlation with the observations and, even though Lagrangian persistence appears to provide a better forecast than Eulerian persistence for this event, after about 90 min they both provide nowcasts that are actually worse than the reference random nowcast. On the contrary, all PhaSt nowcasting fields keep high correlation with observations up to about two hours of forecast time, and they never fall below the skill of the reference random




FIG. 6. Spatial power spectra, as a function of radial wavenumber, of the nowcasted fields generated by the PhaSt method. Different curves refer to different times.

forecast. In addition, the spread around the ensemble mean of the different realizations provides an estimate of forecast uncertainty.

A further confirmation of the good performance of the PhaSt method comes from the comparison between the statistics of the accumulated precipitation data and those of model outputs. Figure 8 shows the amplitude distribution (left) and power spectra (right) for the radar data (dashed curve) and the PhaSt fields (solid curve), accumulated over 0.5 (top), 1 (middle), and 1.5 h (bottom). Although the amplitude distribution of the data and that of the PhaSt fields are the same at every time step, there is no a priori reason why they should remain the same for accumulated precipitation. The fact that the amplitude distributions and the power spectra of the data and of the model outputs remain so close to each other confirms that for this event, the PhaSt method is able to reproduce the main statistical properties of the evolution of the rain field.

c. Simplifications adopted for operational applications

In operational applications, only the past history of the event is available, and a robust estimate of model parameters can become difficult to determine. For this reason, it is necessary to significantly reduce the model complexity. One significant simplification is to assume a constant value for T , independent of wavenumber, and an isotropic form of σ growing linearly with wavenumber, $\sigma = c_o k$. Therefore, the only two free parameters of




FIG. 7. Cross correlation, as a function of forecast time, between observed radar fields and nowcasts for Eulerian persistence (squares), Lagrangian persistence (circles), and 100 nowcasting fields generated using the PhaSt method, assuming in-sample knowledge of the model parameters (thin lines) and their ensemble mean (thick line). The triangles show the limit below which 95% of 3000 random realizations (with the same amplitude distribution and power spectrum of the precipitation data and random, uncorrelated Fourier phases) are found.

the model become T and c_o . Next we discuss the performance of the PhaSt method when this simplification is adopted, using as an example the rainfall event analyzed earlier.

Figure 9a shows the cross correlation between the model forecasts and the observed fields obtained by fixing $T = 15$ min and $c_o = 0.15$ km min⁻¹. Not surprisingly, the ensemble mean now shows lower skill than in the in-sample case. However, even in this case all nowcasting ensemble members have a skill that is larger than the reference random forecast—up to two hours of forecast time. Some ensemble members still achieve very high-correlation values, indicating the good skill of the method even with the simplified approach adopted here.

Figure 9b shows the critical success index (CSI; Wilks 1995), with a threshold of 10 mm h⁻¹, for the same data. The information provided by this measure is consistent with the results of the cross-correlation test, with the PhaSt nowcasting showing better skill than Lagrangian and Eulerian persistence at all times.

The spread of the ensemble gauges uncertainty in the nowcast. Probabilities of exceedence of fixed precipitation thresholds can be derived by counting the number of ensemble members satisfying this condition. Figure 10

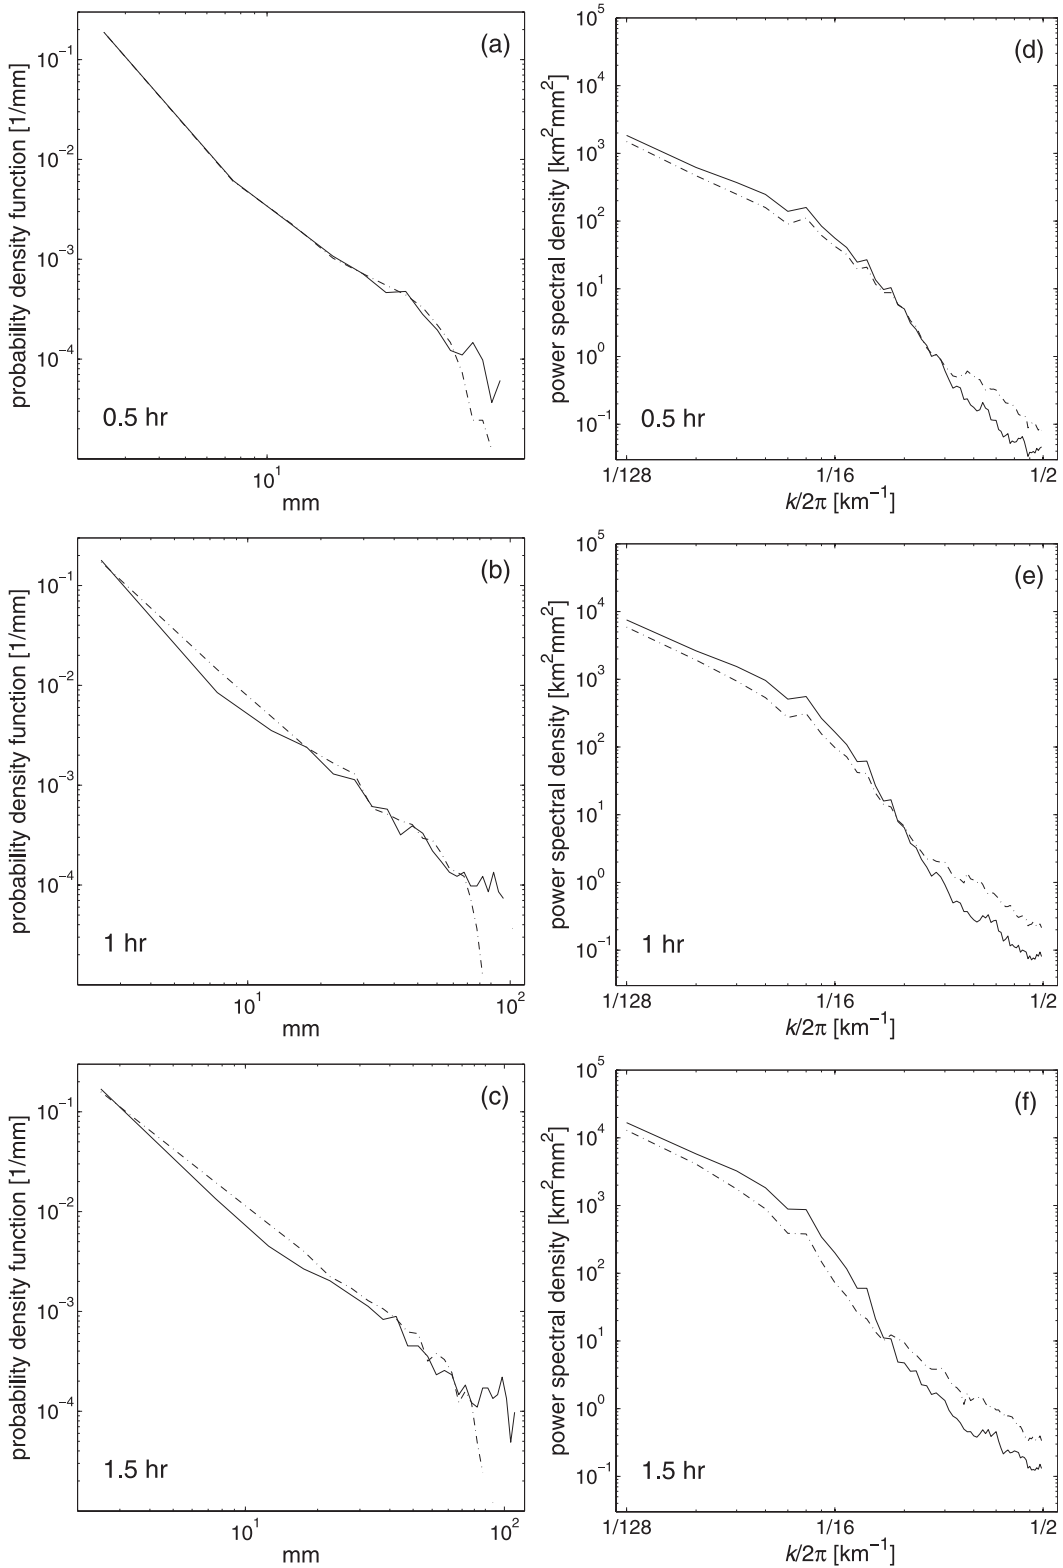


FIG. 8. (left) Amplitude distributions and (right) spatial power spectra, as a function of radial wavenumber, of the accumulated precipitation data (dashed curve) and the accumulated nowcasted fields generated by the PhaSt method (solid curve). Accumulation times of (top) 0.5, (middle) 1, and (bottom) 1.5 h.




FIG. 9. (a) Cross correlations and (b) critical success indices for a threshold of 10 mm h^{-1} , as a function of forecast time, between observed radar fields and nowcasts for Eulerian persistence (squares), Lagrangian persistence (circles), and 100 nowcasting fields generated using the simplified version of the PhaSt method with $\sigma(k) = c_o k$, where $c_o = 0.15 \text{ km min}^{-1}$, and $T = 1.5\Delta t$, where $\Delta t = 10 \text{ min}$. The thin lines show each ensemble member and the thick line indicates the ensemble mean. Refer to Fig. 7 for meaning of triangles.

shows the skill of the PhaSt ensemble in providing a probabilistic forecast. The figure reports the Brier scores (Wilks 1995) for the probability of exceeding a precipitation intensity of 10 mm h^{-1} , as a function of forecast time. For comparison, we report also the hit rates H of the two deterministic Eulerian and Lagrangian nowcasts and the Brier score obtained using a reference random nowcasting having the same amplitude distribution and power spectrum of the initial radar fields and random Fourier phases. The Brier score and the hit rate (or accuracy) are related: we transform observations to a binary verification set $\mathcal{O}(x, y, t)$ —where x and y indicate the position of the pixel of the radar image and t indicates time—by defining a given threshold in precipitation intensity. If the nowcasting is probabilistic, then counting the number of ensemble members exceeding the threshold provides a probability of exceedence, $\mathcal{P}(x, y, t)$. The Brier score B is defined as the average squared difference between \mathcal{P} and the verification set \mathcal{O} . If the forecast is deterministic, then it can be transformed to a binary forecast set $\mathcal{Y}(x, y, t)$. The average squared difference between \mathcal{Y} and the verification \mathcal{O} is equivalent to a Brier score B in case of an ensemble with only one member. In this case, the quantity $H = 1 - B$ is known as “hit rate” or “accuracy” (Wilks 1995).

Also with this probabilistic skill measure, both the Eulerian and Lagrangian persistence nowcasts perform worse than the reference random nowcasting already at relatively short times, whereas the PhaSt nowcasting performs better at all times.

A delicate point in the application of the PhaSt method to operational nowcasting is the choice of the values of T and c_o . An estimate of the correlation time T can be obtained either from the analysis of a sequence of a few past rain fields (assuming persistence of the statistical properties of the event) or fixed from climatology. Preliminary analysis of different events indicates that the value of T is usually smaller than about 20 min in the low wavenumber range, consistent with the typical lifetimes of precipitation structures at kilometer scale. Experimentation with different values of T in the model indicates that nowcasting accuracy is only weakly sensitive to the precise value of T . A reasonable “first guess” is thus to fix $T = 15 \text{ min}$, as in the example discussed earlier.

In principle, the value of c_o can be fixed from the behavior of $\sigma(k)$ in the low wavenumber range, as obtained from a few past rain fields. For the event studied here, Fig. 4 suggests a slope $c_o \approx 0.3 \text{ km min}^{-1}$ at low wavenumbers. At least for this event, however, there is strong anisotropy in spectral space, and the value of c_o obtained from radial average can be misleading. Experimentation with different values of c_o indicates that the average performance of the nowcasting procedure improves as c_o is reduced. At the same time, reducing c_o (and thus σ) leads to a smaller spread of the ensemble, which collapses to zero for $c_o = 0$. A compromise should thus be accepted between the desire for improving the performance of the ensemble mean and the need for obtaining a representative spread of the ensemble.




FIG. 10. Values of $1 - H$ for Eulerian persistence (squares) and Lagrangian persistence (circles). The thick line shows B for an ensemble of 100 nowcasting fields generated using the simplified version of the PhaSt method with $\sigma(k) = c_o k$, where $c_o = 0.15 \text{ km min}^{-1}$ and $T = 1.5\Delta t$, where $\Delta t = 10 \text{ min}$. The triangles indicate the Brier scores for an ensemble of 100 reference random forecasts with the same amplitude distribution and power spectrum of the precipitation data and random, uncorrelated Fourier phases.

4. Summary and conclusions

In this paper we have introduced a novel ensemble nowcasting method based on the stochastic evolution of the Fourier phases of a Gaussianized precipitation field. Because the model is formulated in spectral space, it naturally allows for representing an uncertainty varying with scale. In this model, Lagrangian advection of the precipitation field can be captured by estimating the angular frequencies by finite differences of the Fourier phases of two Gaussianized initial fields. In its simplest version, the model has only two free parameters—that is, the correlation time and the variance entering the stochastic process that regulates the evolution of the angular frequencies. Because the model implements a simple stochastic process, large ensembles of realizations can be created with only modest computational burden, allowing for estimating the probability of occurrence of intense precipitation events in a small area and at a given time.

We illustrated the applicability of this nowcasting method to an example precipitation event, showing that in this case its skill is superior to simple Lagrangian advection and Eulerian persistence—up to about two hours of forecast time. Given the short lifetime of localized precipitation structures (of the order of 20–30 min),

this is probably about the maximum predictability time that can be reached with statistical or deterministic methods in the absence of an assimilation network of the full set of meteorological variables at the scales of individual rainfall structures (i.e., about 1 km). Future explorations should verify the performance of the PhaSt method on a set of precipitation events with different characteristics.

In operational applications, the most difficult part of the method is the correct estimate of the correlation time T and the variance σ^2 . One option is to tune these parameters on a large set of example events, creating a climatology of model parameter values. Another option is to estimate these parameters from a few measured fields—once the storm has started. As an alternative, these parameters can be allowed to vary, generating a larger (“super”) ensemble that includes both individual realizations for given parameter values and different choices for the parameters of the stochastic evolution.

REFERENCES

- Andersson, T., and K.-I. Ivarsson, 1991: A model for probability nowcasts of accumulated precipitation using radar. *J. Appl. Meteor.*, **30**, 135–141.
- Balmforth, N., A. Provenzale, E. Spiegel, M. Martens, C. Tresser, and C. W. Wu, 1999: Red spectra from white and blue noise. *Proc. Roy. Soc. London*, **B266**, 311–314.
- Dixon, M., and G. Wiener, 1993: TITAN: Thunderstorm identification, tracking, analysis, and nowcasting—A radar-based methodology. *J. Atmos. Oceanic Technol.*, **10**, 785–797.
- Ferraris, L., S. Gabellani, U. Parodi, N. Rebora, J. von Hardenberg, and A. Provenzale, 2003a: Revisiting multifractality in rainfall fields. *J. Hydrometeorol.*, **4**, 544–551.
- , —, N. Rebora, and A. Provenzale, 2003b: A comparison of stochastic models for spatial rainfall downscaling. *Water Resour. Res.*, **39**, 1368–1383, doi:10.1029/2003WR002504.
- Fox, N. I., and C. K. Wikle, 2005: A Bayesian quantitative precipitation nowcast scheme. *Wea. Forecasting*, **20**, 264–275.
- Fung, J. C. H., and J. C. Vassilicos, 1998: Two-particle dispersion in turbulentlike flows. *Phys. Rev.*, **57**, 1677–1690, doi:10.1103/PhysRevE.57.1677.
- Germann, U., and I. Zawadzki, 2002: Scale-dependence of the predictability of precipitation from continental radar images. Part I: Description of the methodology. *Mon. Wea. Rev.*, **130**, 2859–2873.
- , and —, 2004: Scale dependence of the predictability of precipitation from continental radar images. Part II: Probability forecasts. *J. Appl. Meteor.*, **43**, 74–89.
- Greco, M., and W. Krajewski, 2000: A large-sample investigation of statistical procedures for radar-based short-term quantitative precipitation forecasting. *J. Hydrol.*, **239**, 69–84.
- Gupta, V., and E. Waymire, 1993: A statistical analysis of mesoscale rainfall as a random cascade. *J. Appl. Meteor.*, **32**, 251–267.
- Herring, J. R., and R. H. Kraichnan, 1972: Comparison of some approximations for isotropic turbulence. *Statistical Models and Turbulence*, M. Rosenblatt and C. M. Van Atta, Eds., Lecture Notes in Physics, Vol. 12, Springer, 148–194.
- Kraichnan, R. H., 1970: Diffusion by a random velocity field. *Phys. Fluids*, **13**, 22–31.

- Krzysztofowicz, R., 2001: The case for probabilistic forecasting in hydrology. *J. Hydrol.*, **249**, 2–9.
- Kumar, P., and E. Foufoula-Georgiou, 1993: A multicomponent decomposition of spatial rainfall fields. 2. Self-similarity in fluctuations. *Water Resour. Res.*, **29**, 2533–2544.
- Lovejoy, S., and B. Mandelbrot, 1985: Fractal properties of rain and a fractal model. *Tellus*, **37A**, 209–232.
- Mellor, D., J. Sheffield, P. E. O’Connell, and A. V. Metcalfe, 2000: A stochastic space-time rainfall forecasting system for real time flow forecasting I: Development of MTB conditional rainfall scenario generator. *Hydrol. Earth Syst. Sci.*, **4**, 603–615.
- Rebora, N., L. Ferraris, J. von Hardenberg, and A. Provenzale, 2006: RainFARM: Rainfall downscaling by a filtered autoregressive model. *J. Hydrometeor.*, **7**, 724–738.
- Schertzer, D., and S. Lovejoy, 1987: Physical analysis and modeling of rain and clouds by anisotropic scaling multiplicative processes. *J. Geophys. Res.*, **92**, 9693–9714.
- Schreiber, T., and A. Schmitz, 1996: Improved surrogate data for nonlinearity tests. *Phys. Rev. Lett.*, **77**, 635–638.
- Seed, A. W., 2003: A dynamical and spatial scaling approach to advection forecasting. *J. Appl. Meteor.*, **42**, 381–388.
- van Dop, H., F. T. M. Nieuwstadt, and J. C. R. Hunt, 1985: Random walk models for particle displacements in inhomogeneous unsteady turbulent flows. *Phys. Fluids*, **28**, 1639–1653.
- Wilks, D. S., 1995: *Statistical Methods in the Atmospheric Sciences: An Introduction*. International Geophysics Series, Vol. 59, Academic Press, 467 pp.
- Wilson, J. W., N. A. Crook, C. K. Mueller, J. Sun, and M. Dixon, 1998: Nowcasting thunderstorms: A status report. *Bull. Amer. Meteor. Soc.*, **79**, 2079–2099.
- , E. E. Ebert, T. R. Saxen, R. D. Roberts, C. K. Mueller, M. Sleigh, C. E. Pierce, and A. Seed, 2004: Sydney 2000 forecast demonstration project: Convective storm nowcasting. *Wea. Forecasting*, **19**, 131–150.
- Xu, G., and V. Chandrasekar, 2005: Operational feasibility of neural-network-based radar rainfall estimation. *IEEE Geosci. Remote Sens. Lett.*, **2**, 13–17.
- Zawadzki, I. I., 1973: Statistical properties of precipitation patterns. *J. Appl. Meteor.*, **12**, 459–472.
- , J. Morneau, and R. Laprise, 1994: Predictability of precipitation patterns: An operational approach. *J. Appl. Meteor.*, **33**, 1562–1571.