



A non-stationary stochastic ensemble generator for radar rainfall fields based on the short-space Fourier transform

Daniele Nerini^{1,2}, Nikola Besic^{1,3}, Ioannis Sideris¹, Urs Germann¹, and Loris Foresti¹

¹Radar, Satellite and Nowcasting Division, MeteoSwiss, Locarno-Monti, Switzerland

²Institute for Atmospheric and Climate Science, ETH, Zurich, Switzerland

³Environmental Remote Sensing Laboratory, EPFL, Lausanne, Switzerland

Correspondence to: Daniele Nerini (daniele.nerini@meteoswiss.ch)

Received: 19 December 2016 – Discussion started: 1 February 2017

Revised: 27 April 2017 – Accepted: 29 April 2017 – Published: 9 June 2017

Abstract. In this paper we present a non-stationary stochastic generator for radar rainfall fields based on the short-space Fourier transform (SSFT). The statistical properties of rainfall fields often exhibit significant spatial heterogeneity due to variability in the involved physical processes and influence of orographic forcing. The traditional approach to simulate stochastic rainfall fields based on the Fourier filtering of white noise is only able to reproduce the global power spectrum and spatial autocorrelation of the precipitation fields. Conceptually similar to wavelet analysis, the SSFT is a simple and effective extension of the Fourier transform developed for space–frequency localisation, which allows for using windows to better capture the local statistical structure of rainfall. The SSFT is used to generate stochastic noise and precipitation fields that replicate the local spatial correlation structure, i.e. anisotropy and correlation range, of the observed radar rainfall fields. The potential of the stochastic generator is demonstrated using four precipitation cases observed by the fourth generation of Swiss weather radars that display significant non-stationarity due to the coexistence of stratiform and convective precipitation, differential rotation of the weather system and locally varying anisotropy. The generator is verified in its ability to reproduce both the global and the local Fourier power spectra of the precipitation field. The SSFT-based stochastic generator can be applied and extended to improve the probabilistic nowcasting of precipitation, design storm simulation, stochastic numerical weather prediction (NWP) downscaling, and also for other geophysical applications involving the simulation of complex non-stationary fields.

1 Introduction

Precipitation exhibits a large spatial and temporal variability over a wide range of scales (Lovejoy and Schertzer, 2006). Such complexity represents a continuous challenge in the practice of observing, understanding and simulating precipitation processes. Despite great efforts, the measurement and prediction of rainfall at fine space–time resolutions is still associated with large errors. Consequently, it is of crucial importance to accurately characterise the uncertainty related to its measurement and prediction, in particular for hydrometeorological applications that are sensitive to input precipitation uncertainty (Zappa et al., 2010).

For medium-range weather forecasting at synoptic and mesoscales, the forecast uncertainty is usually considered by perturbing the initial conditions of numerical weather prediction (NWP) models to generate an ensemble of possible future realisations. However, the NWP approach for uncertainty quantification is often not adapted for several hydrological applications, when high spatial and temporal resolution is required. A typical application for very short-term forecasting (nowcasting) of radar rainfall at a 5 min update rate would need ideally the generation of 50–100 realisations of precipitation fields for the next 6 h in less than 5 min. Assuming 1 s of computational time per realisation, it already requires 60 min on a single processor (50 realisations \times 72 lead times \times 1 s = 60 min). These requirements cannot be met by current operational NWP models and a different approach needs to be taken.

It has been long observed that rainfall fields exhibit spatial and temporal organisation, which is consistent with the

scaling and multifractal framework (Schertzer and Lovejoy, 1987; Menabde et al., 1997). This observation is convenient since it allows for developing parsimonious stochastic simulation methods, which are able to generate realistic precipitation fields that depend only on a few parameters. Stochastic rainfall generators are effective in characterising the uncertainty related to the measurement and forecasting of precipitation and they constitute the core of most statistical design storm techniques, probabilistic nowcasting and NWP downscaling methods.

An ideal stochastic rainfall generator for hydrometeorological applications should be able to

1. simulate large precipitation fields ($> 500 \times 500$ grid points) with minimum computational time ($< 1\text{--}10$ s);
2. reproduce the observed spatial and temporal autocorrelation of the observed precipitation fields;
3. reproduce the observed power-law scaling behaviour of the rainfall fields (ideally the multifractal scaling of different moments of order q);
4. reproduce the observed marginal distribution of rainfall rates (often skewed);
5. take into account the rainfall intermittency and the effects due to the rain/no-rain transition;
6. integrate in a flexible way possible sources of predictability (NWP forecasts, analogues, orography, etc.);
7. handle the spatial and temporal non-stationarity of the statistical properties of precipitation.

This paper focuses particularly on point 7 and presents a non-stationary spatial stochastic rainfall generator, which is able to reproduce the local variability and anisotropy of the observed precipitation fields.

1.1 Brief review of spatial stochastic rainfall generators

There is abundant literature on random field generators using geostatistical (Goovaerts, 1997; Lantuéjoul, 2002) and multifractal approaches (Menabde et al., 1997; Lovejoy and Schertzer, 2006). A widely used covariance-based geostatistical random field generator in hydrology is based on the turning bands method (TBM; see a review in Lantuéjoul, 2002), which reduces the problem of simulating a multi-dimensional Gaussian random field to the simulation of a set of independent uni-dimensional autocorrelated Gaussian processes along random orientations (for practical implementations in hydrology refer to Mantoglou and Wilson, 1982; Tompson et al., 1989). In the context of precipitation, Leblois and Creutin (2013) adapted the TBM to simulate unconditional stochastic fields, which reproduce the correct intermittency and advection of precipitation fields, which was further extended for ensemble precipitation nowcasting by

Casari et al. (2016). Schleiss et al. (2014) also followed the geostatistical approach by using sequential Gaussian simulations (SGS) to generate conditional and unconditional radar rainfall fields, which realistically decay towards zero when moving out from the wet regions (concept of “dry drift”). Despite the speed-up strategies implemented, both the SGS and TBM approaches are still too computationally demanding to generate large precipitation ensembles over extended domains for real-time applications.

Another class of algorithms relies on the fast Fourier transform (FFT), which has efficient and easy to use libraries in most programming languages. The Wiener–Khinchine theorem (Wiener, 1930; Khinchine, 1934) states that the power spectral density (generally referred to as “power spectrum”) of a wide-sense stationary random process can be obtained as the Fourier transform of the autocorrelation (autocovariance) function of that process. Thus, it becomes clear that stochastic simulation methods based on covariance or variogram can be easily reformulated by exploiting the FFT to achieve increased computational efficiency (Marcotte, 1996). In addition, the representation of data in Fourier space can enhance our understanding of the underlying generating mechanism and scaling behaviour of rainfall.

One of the most straightforward and widely used stochastic noise generation techniques is the filtering of a white noise field using a Fourier spectrum characterised by a power-law behaviour (linear relationship in log–log plot of power vs. frequency). The inverse Fourier transform of the filtered noise field directly provides a spatially correlated field, which reproduces the chosen power-law spectrum (for example corresponding to an exponential covariance function). This stochastic generator will be referred to as global FFT noise throughout the paper. The FFT provides an efficient solution to simulate fractional Brownian motion fields (Mandelbrot and Van Ness, 1968), multi-dimensional and also multivariate autocorrelated random fields (Robin et al., 1993). Fractional noise integration is particularly useful to model geophysical processes, which exhibit Fourier spectra following power laws of fractional order (Schertzer and Lovejoy, 1987), in particular precipitation fields (Menabde et al., 1997; Lovejoy and Schertzer, 2006).

There is a long list of probabilistic precipitation estimation, nowcasting, design storm and NWP downscaling techniques that exploit the FFT noise to generate ensembles of stochastic rainfall fields. One of the most advanced probabilistic nowcasting techniques is the Short-Term Ensemble Prediction System (STEPS; Bowler et al., 2006), which represents the uncertainty due to rainfall growth and decay processes by adding stochastic perturbations to a deterministic radar extrapolation forecast. The radar rainfall field is decomposed in Fourier space into an eight-level multiplicative cascade to allow for representing the scale dependence of the predictability of precipitation. The rainfall cascade is blended scale by scale with a cascade of spatially correlated noise fields to consider the different rate of development of pre-

precipitation at different spatial scales. STEPS is operationally used at the Australian Bureau of Meteorology, the UK Met Office and the Royal Meteorological Institute of Belgium, and is in continuous development (see Seed et al., 2013; Foresti et al., 2016). The STEPS framework was also used for design storm generation to produce realistic synthetic rainfall fields (Seed et al., 1999; Niemi et al., 2016). SBM-cast is another well-known probabilistic nowcasting system, which employs the FFT noise technique (Berenguer et al., 2011). It is based on the “String of Beads” model of Pegram and Clothier (2001a, b), which simulates the joint evolution of global radar rainfall field statistics that define the general storm behaviour (wet area ratio and image mean flux, i.e. rainfall fraction and mean rainfall) and the local rainfall. Metta et al. (2009) developed a probabilistic nowcasting system that exploits a Langevin-type model for the evolution of Fourier phases to generate isotropic stochastic rainfall fields. A similar approach was implemented within the RainFARM model of Rebola et al. (2006b) for stochastic NWP precipitation downscaling in space and time. STREAP (Paschalis et al., 2013) also uses the Fourier transform for design storm generation to produce Gaussian random fields with an exponentially decaying isotropic autocorrelation function. Atencia and Zawadzki (2014) used the power spectrum of the last observed radar rainfall field to generate spatially correlated and anisotropic stochastic perturbations for ensemble precipitation nowcasting. Niemi et al. (2014) parametrised the two-dimensional (2-D) power spectrum of precipitation fields to obtain a Fourier filter able to simulate anisotropic stochastic rainfall fields for design storm generation (Niemi et al., 2016).

Stochastic noise generators are also increasingly used to characterise the residual radar measurement uncertainty (Ciach et al., 2007) for ensemble quantitative precipitation estimation (ensemble QPE). An important group of techniques generates the stochastic perturbations by LU or Cholesky decomposition of the covariance matrix, which represents the multiplicative radar and rain gauge errors (see, e.g., Germann et al., 2009; Villarini et al., 2009). Given the direct link between the covariance and the Fourier spectrum via the Wiener–Khinchine theorem, it is not a surprise to observe that FFT-based noise generators are also used for ensemble radar-based QPE (Jordan et al., 2003; Pegram et al., 2011). An interesting question is how to condition the unconditional rainfall ensembles to rain gauges, e.g. as done in the conditional merging of Sinclair and Pegram (2005) or the random mixing approach of Bardossy and Hörning (2016). It is worth mentioning that Velasco-Forero et al. (2009) and Schiemann et al. (2011) also took advantage of the Wiener–Khinchine theorem to obtain 2-D positive definite correlograms directly from radar rainfall images for improved blending with rain gauge measurements.

1.2 Limitation of current stochastic rainfall generators

Apart from few exceptions, the stochastic generators presented above assume spatial stationarity, i.e. uniformity of the generator across space. As a consequence the generator is only able to produce spatially homogeneous noise fields (see point 7 in the requirements of the generator at the beginning of the Introduction). More precisely, they cannot represent and simulate non-stationary fields that display a heterogeneous distribution of autocorrelation range and/or anisotropy.

The non-stationarity in the local statistical properties becomes easily visible when stratiform and convective precipitation coexist in different regions of the radar domain or when there are clear local anisotropies in different directions (see for example Fig. 1). The need for a non-stationary stochastic noise generator is even more relevant in regions with complex orography, such as Switzerland, where the statistical properties and predictability of precipitation are strongly controlled by orographic forcing (Panziera and Germann, 2010; Mandapaka et al., 2012; Foresti and Seed, 2014, 2015). A few approaches have been proposed to adapt the stochastic generators to better capture the local variability of precipitation. Pathirana and Herath (2002) and Badas et al. (2006) proposed to filter or remove the heterogeneities of the precipitation field before applying a homogeneous multifractal downscaling technique and re-adding the trend after running the stochastic simulation. A similar strategy was employed in a recent STEPS implementation, which computes the local mean and variance of the rainfall field on a regular grid to remove the local non-stationarity before adding the homogeneous stochastic perturbations (Seed et al., 2013). However, these strategies are not able to generate noise fields that reflect the local 2-D power spectrum and anisotropy of precipitation fields. When applied for nowcasting of weather situations presenting strong local anisotropies, the stochastic perturbations gradually destroy the original spatial structure and important information is rapidly lost. The need for a more flexible approach to directly consider the non-stationarity of real-world geophysical fields such as rainfall was already pointed out by Lovejoy and Schertzer (2013). The same authors argued that the estimation of the generator over small sections of the field in a moving window fashion suffers from the small sample size. Therefore, they proposed to model the spatial variation of the variables defining the generator using a non-linear extension of generalised scale invariance (GSI), whose estimation, however, is also not trivial (Lovejoy and Schertzer, 2013).

1.3 Contribution of this study

The contribution of this paper is the design of a non-stationary stochastic rainfall generator based on the short-space Fourier transform (SSFT; Hinman et al., 1984). The idea is to extend the short-time Fourier transform (STFT; Kovačević et al., 2013) for 2-D spatial fields both for lo-

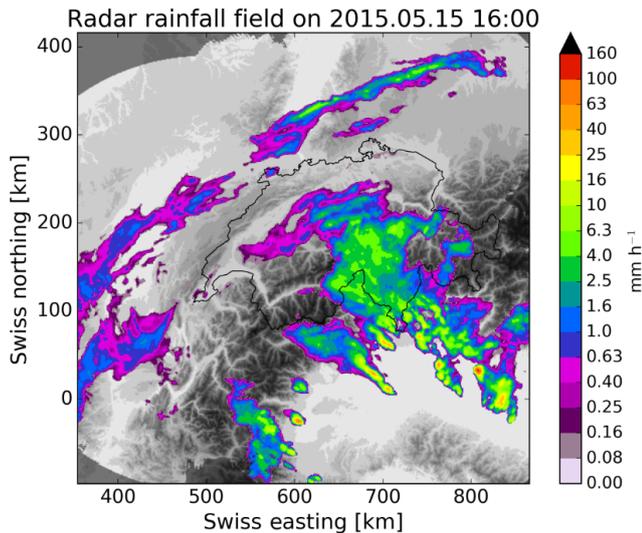


Figure 1. Example of non-stationarities in a radar rainfall field of $512 \text{ km} \times 512 \text{ km}$ centred over Switzerland. Two distinct anisotropies coexist in the same image, whereas convection develops only to the south of the Alps.

cal Fourier analysis and stochastic simulation. Conceptually similar to the wavelet transform (Kumar and Foufoula-Georgiou, 1997; Kovačević et al., 2013), the STFT and SSFT allow for time–frequency and space–frequency localisation respectively; i.e. they enable performing a local moving window Fourier analysis to account for non-stationarity in the statistical properties of the signal (in our case a 2-D precipitation field). The traditional global FFT noise technique can be easily extended by performing local Fourier filtering of white noise in a moving window fashion using the local Fourier transform of the precipitation field as filter. The obtained locally correlated noise field (denoted as SSFT noise) automatically reproduces the local correlation structure of the target precipitation field. In this study we want to investigate to which extent it is reasonable to estimate a local stochastic generator from the data considering the larger uncertainty due to the reduced sample size.

In order to demonstrate the advantage of using a non-stationary noise generator, we selected four non-stationary rainfall fields from the fourth generation of Swiss dual-polarisation weather radar composite. We verified the ability of the generator in replicating the global and local statistical properties of the precipitation fields (Fourier spectrum, anisotropy). The SSFT noise generator will open the doors for enhanced probabilistic precipitation estimation and nowcasting, design storm generation and stochastic NWP downscaling. Geostatistical interpolation methods considering the local anisotropy (e.g. Boisvert et al., 2009; Gyasi-Agyei, 2016) could also be extended to perform local stochastic simulation.

The paper is organised as follows. Section 2 describes the global Fourier-based method for generating spatially corre-

lated stochastic noise fields (FFT noise). Section 3 explains the generation of locally correlated stochastic noise based on the SSFT noise and Sect. 4 makes use of synthetic data produced with parametric methods in order to illustrate and test the functioning of the novel approach. Section 5 presents four selected radar rainfall fields characterised by strong non-stationarities. The four precipitation cases are thus used for the verification of the local SSFT noise generator in comparison with the global FFT noise generator, in particular its ability to reproduce both the global and local structure of precipitation fields. At the end of the Sect. 5, we provide two examples demonstrating the potential applications of the local noise generator. First, we present stochastic simulations of synthetic radar rainfall fields by appropriate local adjustment of the noise fields. Then, we introduce the temporal structure of the precipitation field by using the local noise as a shock term in an auto-regressive process in the context of a simplified stochastic precipitation nowcast. Section 6 concludes the paper and enumerates possible applications and extensions of the generator for stochastic precipitation simulation.

2 Stationary stochastic generator

This section formulates the model to generate 2-D spatially correlated noise and precipitation fields. Such fields should reproduce the spatial scaling, and thus spatial autocorrelation, of the target radar rainfall field. The numerical simulation of such random fields is based on the idea of filtering a Gaussian white noise field using the power-law Fourier spectrum of the precipitation field, also known as fractional integration (Schertzer and Lovejoy, 1987). The following sections explain the basic principles of the Fourier transform and how it can be used for simulation of stationary spatially correlated random fields.

2.1 Fourier transform

The corner stone of signal processing theory (Stankovic et al., 2013) is the idea that a signal $x(t)$ can be represented as an infinite weighted sum of harmonic sine waves of infinitesimally close frequencies f :

$$x(t) = \int_{-\infty}^{+\infty} X(f) e^{j2\pi ft} df, \quad (1)$$

where $X(f)$ represents the complex Fourier frequency representation. It contains both the amplitude ($|X(f)|$) and the phase (φ) information ($X(f) = |X(f)| e^{j\varphi}$). Obtained by means of the Fourier transform (FT), it is actually the same signal, only represented in the frequency domain:

$$X(f) = F\{x(t)\} = \int_{-\infty}^{+\infty} x(t) e^{-j2\pi ft} dt, \quad (2)$$

where $F\{x(t)\}$ denotes the Fourier transform of the signal $x(t)$. Therefore, the Fourier transform could be essentially interpreted as the transformation of a signal from a given space to the space of frequencies, where the energy/average power of a signal is conserved (Parseval's theorem). Signal $x(t)$ and its Fourier counterpart $X(f)$ are often called the FT pair.

If the signal $x(t)$ is deterministic (or transient stochastic), its energy spectral density can be defined as $S(f) = |X(f)|^2$, while in the case of a non-transient stochastic signal we ought to define its power spectral density (commonly referred to as "power spectrum"), being $S(f) = \lim_{T \rightarrow \infty} \frac{|X(f)|^2}{2T}$. Both the energy spectral density and the power spectral density form the Fourier transformation pair with the autocorrelation function of a corresponding signal:

$$R(\tau) = \int_{-\infty}^{+\infty} x(t)x(t+\tau)dt = F^{-1}\{S(f)\}. \quad (3)$$

In the latter case, of non-transient stochastic signals, this relation is called Wiener–Khinchine theorem, and it is very relevant in the analysis of stochastic fields. In other words, the autocorrelation function of a signal can be simply obtained as the inverse FT of its power spectral density.

That leads us to the next step, which is the expansion of the Fourier transform to the analysis of stochastic spatial signals – stochastic fields. The expanded, 2-D Fourier transform (2-D FT) would be, as a matter of fact, the corner stone of image processing theory (Stankovic et al., 2012). Namely, using 2-D FT a field $f(x, y)$ can be transformed from a 2-D space to the 2-D frequency domain:

$$F(f_1, f_2) = F\{f(x, y)\} = \iint_{-\infty}^{+\infty} f(x, y)e^{-j2\pi(f_1x+f_2y)}dxdy, \quad (4)$$

where $f_1 = \frac{1}{x}$ and $f_2 = \frac{1}{y}$ are the spatial frequencies in x and y directions. Analogously to the 1-D signal, which can be represented as a weighted sum of 1-D sinusoids, a field is represented by a weighted sum of sinusoids that vary in the $x - y$ plane:

$$f(x, y) = F^{-1}\{F(f_1, f_2)\} = \iint_{-\infty}^{+\infty} F(f_1, f_2)e^{j2\pi(f_1x+f_2y)}df_1df_2, \quad (5)$$

with the complex Fourier representation $F(f_1, f_2) = |F(f_1, f_2)|e^{j\varphi}$. The power spectral density of a stochastic field is estimated on $\Delta x \times \Delta y$ portion of a plane as $S(f_1, f_2) = \lim_{\Delta x \Delta y \rightarrow \infty} \frac{|F(f_1, f_2)|^2}{\Delta x \Delta y}$. Again, analogously to the 1-D case, it forms the Fourier transformation pair with the 2-D autocorrelation function (Wiener–Khinchine theo-

rem):

$$R(\xi, \zeta) = \iint_{-\infty}^{+\infty} f(x, y)f(x+\xi, y+\zeta)dxdy = F^{-1}\{S(f_1, f_2)\}. \quad (6)$$

In order to obtain the autocorrelation function, the signal needs to be standardised by removal of the mean and division by the standard deviation; otherwise only the non-centred autocovariance is obtained. Using the Wiener–Khinchine theorem and exploiting the speed of FFT one can compute the autocorrelation function of regularly sampled signals very efficiently (Marcotte, 1996; Velasco-Forero et al., 2009; Schiemann et al., 2011). On the other side, the classical approach to derive the spatial autocorrelation function would require computing all the Euclidean distances between pairs of points located at different ranges and along different directions.

Aside from the complete 2-D information about field power spectral density ($S(f_1, f_2)$), it is common to use the 1-D, simplified representation ($S_{1-D}(f)$), obtained by radially averaging $S(f_1, f_2)$ around the $S(0, 0)$ point. This allows us to analyse the scaling behaviour of rainfall, by characterising the power spectral curve with the spectral exponent β , obtained by its power-law approximation in the log–log plot $S_{1-D}^a(f) = \frac{1}{f^\beta}$.

One of the most noteworthy properties of the Fourier transform is a significantly facilitated convolution in the spatial (or temporal) domain. Namely, the convolution of two fields $f_1(x, y)$ and $f_2(x, y)$ forms the FT pair with the product of their respective Fourier transforms $F_1(x, y)$ and $F_2(x, y)$:

$$F\{f_1(x, y) * f_2(x, y)\} = F_1(f_1, f_2)F_2(f_1, f_2). \quad (7)$$

2.2 Global filtering of white noise with FFT

In analysing real, observed fields, we actually never deal with the continuous functions (e.g. $r(x, y)$) but rather with the discrete ones. That would mean that a 2-D field can be represented by a matrix of finite size ($\mathbf{r}_{M \times N}$), with the same being true for its Fourier transform $\mathbf{R}_{M \times N}$. In this case, the transformation from Eq. (4) takes the form of a discrete Fourier transform (DFT):

$$\mathbf{R}(m, n) = \frac{1}{\sqrt{MN}} \sum_{k=0}^{M-1} \sum_{l=0}^{N-1} \mathbf{r}(k, l) e^{-j2\pi\left(\frac{km}{M} + \frac{ln}{N}\right)}, \quad (8)$$

which, particularly if M and N are powers of 2 or have only small prime factors, can be efficiently implemented using FFT algorithms (Cooley and Tukey, 1965). The transformation preserves all the properties presented in the preceding section, out of which the facilitated convolution especially proves to be useful in generating the stationary correlated noise.

Namely, the straightforward way of generating a stationary correlated noise field, i.e. a stochastic field with the overall spatial correlation of a natural rainfall field, is the non-parametric filtering of a white noise field (\mathbf{n}). This basically assumes convolving a generated white Gaussian noise field, characterised with zero mean and unit variance, with the selected rainfall field. If we employ the FFT algorithm to calculate the DFT of a rainfall field (\mathbf{R}), and the DFT of a noise field (\mathbf{N}), a stationary correlated noise field (\mathbf{n}_{scr}) can be derived as

$$\mathbf{n}_{\text{scr}} = \text{FFT}^{-1} \{ \mathbf{R} \circ \mathbf{N} \}, \quad (9)$$

where $[\cdot] \circ [\cdot]$ is the pointwise (Hadamard) product. Given the quasi-constant power spectral density of white noise, this way we mostly alter the spectral phase of a rainfall field, without significantly modifying the amplitude. It makes the resultant spectrum of one realisation of correlated noise field (\mathbf{N}_{scr}) almost identical to the spectrum of the considered rainfall field (\mathbf{R}), whose inverse FFT corresponds to the spatial autocorrelation function. The latter could also be achieved by generating a noise field, which can be used to directly perturb the phase of \mathbf{R} (Metta et al., 2009). In this case we can avoid calculating the Fourier transform of the noise, $\text{FFT} \{ \mathbf{n} \}$, which can slightly contribute to the computational efficiency. Alternatively, the rainfall field spectrum (\mathbf{R}) could be as well approximated by a suitable power spectral law, resulting in the parametric filtering of white noise field.

The most important drawback of the obtained stationary correlated noise field is its spatial stationarity. Namely, as already pointed out, the obtained spectrum (\mathbf{N}_{scr}) corresponds to the one of the selected rainfall field (\mathbf{R}), which implies that the two fields have the same global field anisotropy properties and scaling behaviour. This, however, does not suggest that a correlated noise field exhibits the same local behaviour as its rainfall field counterpart; i.e. two fields do not have the same local field anisotropy properties.

2.3 Non-parametric and parametric filters

In Sect. 2.2 the white noise field was filtered by the actual DFT of the rainfall field without any parametrisation of the distribution of power as a function of frequency and direction. This is referred to as non-parametric filtering of white noise and allows for automatic simulation of the observed anisotropies and variability of the spectral slope as a function of scale (Seed et al., 2013). This approach is particularly effective for cases with sufficient precipitation coverage over the radar domain and in the absence of measurement noise.

On the other hand, for cases with a small fraction of precipitation, the Fourier transform suffers from the small sample size. As a consequence, the power spectrum becomes noisy and it may be hard to observe the anisotropy or the power-law rainfall scaling behaviour. In order to increase the sample size and reduce the noise, a more stable Fourier power spectrum can be computed by averaging the spec-

tra of a sequence of precipitation fields as done by Niemi et al. (2014). This strategy could be helpful in particular when computing local Fourier transforms for the SSFT (see Sect. 3.1).

Instead, it is common to parametrise the 2-D or 1-D power spectra using simple power-law scaling models for rainfall. The parametrisation of 2-D power spectra can be achieved by constructing an anisotropic filter that considers isolines of constant power as a function of frequency based on the principle of GSI (Niemi et al., 2014). A simpler isotropic scaling model can be constructed from the radially averaged 1-D power spectrum by neglecting the field anisotropy. A power-law fit of the rainfall field power vs. spatial frequency in log–log plot is generally sufficient for the task. More refined models acknowledge the presence of a scaling break in the power-law around 10–20 km by fitting two spectral exponents (β_1 and β_2), which represent the different scaling regimes of the large-scale precipitation features and the small-scale convective cells respectively (Gires et al., 2011; Seed et al., 2013).

3 Non-stationary stochastic generator

The following section introduces the concept of SSFT, and explains how the former is used in overcoming the evoked issue of spatial stationarity in the simulation of spatially correlated noise fields.

3.1 Short-space Fourier transform

The concept of SSFT is introduced through its more intuitive 1-D temporal equivalent and then extended to the 2-D spatial case. Namely, the phase of the Fourier representation of a time series contains information about the temporal distribution of spectral components, but its interpretation proves to be difficult. Therefore, due to the increasing need to face the non-stationarity of the signal, numerous advanced methods have been proposed (Stankovic et al., 2013), all of which can be gathered under the umbrella of time–frequency signal analysis (TFSA).

The most intuitive and representative TFSA method is the STFT, which provides Fourier spectra of localised parts of a signal:

$$X^{\text{ST}}(t, f) = \int_{-\infty}^{+\infty} x(t + \tau) w(\tau) e^{-j2\pi f \tau} d\tau, \quad (10)$$

where $w(\tau)$ is the window localising the particular part of a signal around the time instant t ($x(t + \tau)w(\tau)$). The introduced properties of FT (power spectral density, convolution, autocorrelation) are valid for the obtained $X^{\text{ST}}(t, f)$, in the same way they are for the representation of the entire signal.

The original signal can be obtained by integrating over all individual inverse FT:

$$\begin{aligned}
 x(\tau) &= \int_{-\infty}^{+\infty} x(t + \tau)w(\tau)dt \\
 &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} X^{ST}(t, f) e^{j2\pi f\tau} df dt, \tag{11}
 \end{aligned}$$

which, due to the overlapping of segments, provides a better reconstruction, particularly important in the case of highly non-stationary signals. If $\int_{-\infty}^{+\infty} w(\tau) = 1$, despite the overlapping, there is no need to normalise the obtained result. Theoretically, it would be enough to perform the inverse FT at the time step $\Delta t = 2T$, which corresponds to the width of the window.

Analogously to the Fourier transform, the STFT can be expanded to the spatial domain, becoming the SSFT:

$$F^{SS}(x, y, f_1, f_2) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x + \epsilon, y + \mu)w(\epsilon, \mu) e^{-j2\pi(f_1\epsilon + f_2\mu)} d\epsilon d\mu, \tag{12}$$

with the window $w(\epsilon, \mu)$ now localising the particular segment of a field around the position x, y ($f(x + \epsilon, y + \mu)w(\epsilon, \mu)$). The original field is obtained as

$$\begin{aligned}
 f(\epsilon, \mu) &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x + \epsilon, y + \mu) w(\epsilon, \mu) dx dy \\
 &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} F^{SS}(x, y, f_1, f_2) e^{j2\pi(f_1\epsilon + f_2\mu)} df_1 df_2 dx dy. \tag{13}
 \end{aligned}$$

The critical part concerns choosing the size of the window, as well as its shape. Namely, due to the uncertainty principle, a smaller window in the space domain, i.e. a better space resolution, results in a worse frequency resolution, and vice versa. Choosing an appropriate function $w(\epsilon, \mu)$ can decrease the otherwise inevitable uncertainty.

3.2 Window functions

Equivalently to the convolution in space domain, which corresponds to the product in its frequency counterpart (Eq. 7), the product in space domain forms a Fourier pair with the convolution of the spectra. Therefore, the localisation of the field in space, with the window function $w(\epsilon, \mu)$, implies the localisation in frequency through the convolution with the FT of the window function $W(f_1, f_2)$.

The most intuitive choice of a window would be a simple rectangular function of width $2T$ (Fig. 2), which in case of a field analysis, through the product of two separable functions, becomes a square:

$$w_R(\epsilon, \mu) = w_R(\epsilon) w_R(\mu) = \begin{cases} 1 & \text{for } |\epsilon| < T \text{ and } |\mu| < T \\ 0 & \text{elsewhere} \end{cases} \tag{14}$$

and whose FT is as well a product of the FTs of two separable functions:

$$\begin{aligned}
 W_R(f_1, f_2) &= F\{w_R(\epsilon)\} F\{w_R(\mu)\} \\
 &= \frac{\sin(2\pi f_1 T) \sin(2\pi f_2 T)}{\pi^2 f_1 f_2}. \tag{15}
 \end{aligned}$$

However, as it can be seen in Fig. 2, the side lobes of $W_R(f_1, f_2)$ are somewhat strong, causing the poor localisation in the frequency domain. Hence, the need to come up with a more suitable form of a window, which can reduce the overall uncertainty, through the good localisation in both space and frequency domain. Although the best choice in terms of uncertainty would be the Gaussian window, due to the performance in the context of our application, we have adopted the Hanning window:

$$w_H(\epsilon, \mu) = \begin{cases} \frac{1}{4} \left(1 + \cos\left(\frac{\epsilon\pi}{T}\right)\right) \left(1 + \cos\left(\frac{\mu\pi}{T}\right)\right) & \text{for } |\epsilon| < T \text{ and } |\mu| < T \\ 0 & \text{elsewhere} \end{cases}. \tag{16}$$

The corresponding FT decreases proportionally to the f^3 :

$$W_H(f_1, f_2) = \frac{\pi^2 \sin(2\pi f_1 T) \sin(2\pi f_2 T)}{4\pi^2 f_1 f_2 (\pi^2 - 4\pi^2 f_1^2 T^2) (\pi^2 - 4\pi^2 f_2^2 T^2)}, \tag{17}$$

leading to a better localisation in the frequency domain with respect to the rectangular window (see Fig. 2).

3.3 Local filtering of white noise with SSFT

As already stated previously, in the context of numerous diverse applications, the most obvious drawback of a conventionally derived correlated noise field is its spatial stationarity. Tackling that issue, which is the main motivation and idea of this work, is conceptually quite simple: it comes down to the replacement of the Fourier transform, in both non-parametric and parametric filters, with the SSFT.

Namely, instead of deriving either a parametrised or non-parametrised (Eq. 8) Fourier spectrum of the entire rainfall field in order to filter white noise, here we derive a number of local spectra by focusing at different parts of the field. This is done by sliding and centring the window at positions (p_1, p_2) :

$$\begin{aligned}
 \mathbf{R}(p_1, p_2, m, n) & \tag{18} \\
 &= \frac{1}{\sqrt{MN}} \sum_{k=0}^{M-1} \sum_{l=0}^{N-1} \mathbf{r}(p_1 + k, p_2 + l) \mathbf{w}(k, l) e^{-j2\pi\left(\frac{km}{M} + \frac{ln}{N}\right)},
 \end{aligned}$$

where p_1 varies in the range $(1, \frac{M}{\Delta})$ (equivalent for p_2), with $2T \times 2T$ being the size of the non-zero portion of the

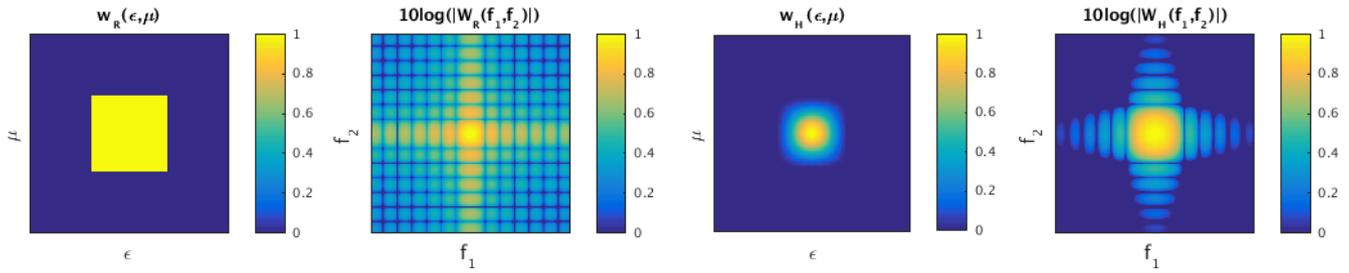


Figure 2. FT pairs for rectangular and Hanning windows (values normalised to unity).

window function. The maximum Δ value corresponds to the size of the window ($2T$) to avoid gaps, yet the overlapping ($\Delta < 2T$) can ensure a smoother representation of the non-stationarity. Although the integration could be performed only in the non-zero part of the matrix \mathbf{w} , the window is zero padded to the size of the field $\mathbf{r}_{M \times N}$, resulting in a better frequency sampling resolution.

In the following step, the local spectra are convolved with a white noise field \mathbf{n} , producing the local spatially correlated noise:

$$\mathbf{n}_{\text{scr}}(p_1, p_2) = \text{FFT}^{-1} \{ |\mathbf{R}(p_1, p_2)| \circ \mathbf{N} \}. \quad (19)$$

By using only the amplitude spectrum of the rainfall field (its absolute value), rather than perturbing the phase of the rainfall field, we are actually imposing the random phase of the white noise field, and slightly perturbing the amplitude. This way we preserve the local anisotropy, while simultaneously benefiting from the phase coherence in the final recomposition; i.e. it decreases the need for the overlapping. The recomposition of the final non-stationary field is obtained by summing over all the local spatially correlated noise fields:

$$\mathbf{n}_{\text{scr}} = \sum_{p_1=1}^{\frac{M}{\Delta}} \sum_{p_2=1}^{\frac{N}{\Delta}} \mathbf{n}_{\text{scr}}(p_1, p_2), \quad (20)$$

where $\frac{M}{\Delta}$ is the number of windows in the k direction and $\frac{N}{\Delta}$ is the number of windows in the l dimension. In other words, if the rainfall field $\mathbf{r}_{M \times N}$ is represented by a set of 10×10 local FFTs (without overlapping), one must sum over the 100 correlated noise fields, which are weighted by the window for spatial localisation of the noise.

4 Synthetic example

The novel concept can be intuitively demonstrated, or rather illustrated, using a synthetic data set. Namely, instead of using Eq. (9) (i.e. the non-parametric approach), here we rely on parametric filtering in order to prescribe a power spectrum to a field of Gaussian white noise. The most straightforward approach is to use an isotropic power-law filter:

$$H(f) = f^{-\beta/2}, \quad (21)$$

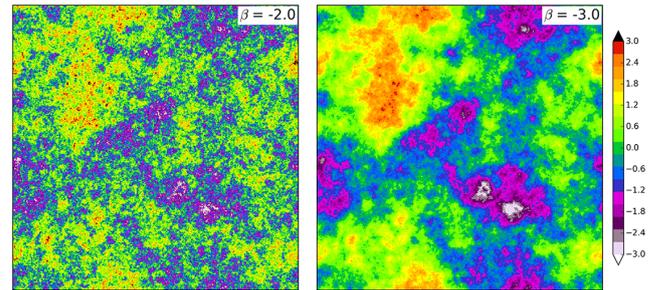


Figure 3. Examples of $512 \text{ px} \times 512 \text{ px}$ synthetic fields of correlated noise produced with an isotropic power-law filtering of Gaussian white noise using two different values of β . Both noise fields have been drawn from the standard normal distribution and share the same random seed.

where $H(f)$ is the parametric filter as a function of f , the spatial frequency or wave number. It follows a stationary, isotropic field of correlated noise, whose spectral slope is nearly equal to β . It can be seen in Fig. 3 that β controls the smoothness of field, or autocorrelation range, as a more negative β results in a more pronounced spectral slope and consequently in a higher power at larger wavelengths.

The same filtering can be applied locally, so that a set of β is used to control the autocorrelation range within successive windows over the same field of Gaussian white noise. Figure 4a presents a set of non-stationary autocorrelation functions in terms of autocorrelation range overlaid on the resulting correlated noise. In other words, these parametric 2-D autocorrelation functions are used to filter white noise, locally, using a set of windows without overlapping. This image represents our target field. It is easy to notice that the autocorrelation range of the background noise field decreases when going from the top left to the bottom right corner of the image. In the following experiment we compute the global and local Fourier transforms of the target noise field in order to filter the same white noise field but with non-parametric filters. In other words, we want to test whether the local SSFT is able to learn the prescribed parametric spectra directly from the target noise field. Given such heterogeneous target field, the application of the global non-parametric generator will fail to capture the trend in the spec-

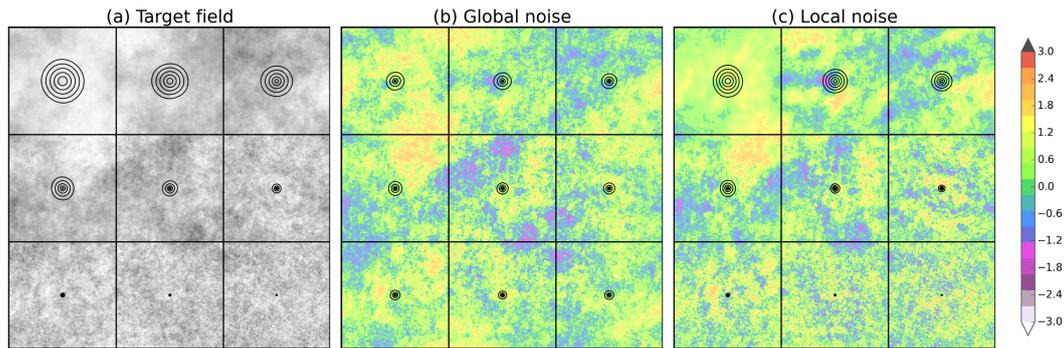


Figure 4. (a) An example of $512 \text{ px} \times 512 \text{ px}$ non-stationary correlated Gaussian noise generated using parametric filtering. The spectral slope β increases from -3.0 to -1.5 from the top towards the bottom of the image as represented by black contours at 0.5, 0.6, 0.7, 0.8 and 0.9 correlation coefficients. Non-parametric stochastic realisations of the target field using a global generator (b) and a local SSFT generator (c) with Hanning windows centred over the grid boxes in the image. All simulations have been drawn from the standard normal distribution and share the same random seed.

tral slope. This can be seen in Fig. 4b, where the noise field has homogeneous characteristics throughout the image and it is confirmed by the black contours representing the resulting local non-parametric 2-D autocorrelation functions. Conversely, the local non-parametric generator applied in Fig. 4c closely reproduces the autocorrelation structure that was prescribed in the target image. The spatial coherence between different windows, i.e. the continuity in absolute values at the edges of consecutive windows, is the result of convolving the white noise field with the amplitude spectrum of the rainfall field only. A reasonable amount of overlapping between windows (e.g. 50 %) would be sufficient to obtain a smoother representation of the non-stationary anisotropy or autocorrelation range in presence of stronger gradients.

As a next step, the method was tested in its ability to reproduce the locally varying anisotropy of a synthetic target image. This target image was produced by means of the GSI model as presented by Niemi et al. (2014). The method was firstly introduced by Lovejoy and Schertzer (1985) to describe the scaling behaviour of anisotropy within multifractal fields. The GSI method has the advantage of being able to model the scale dependence of anisotropy using only few parameters. It should be noted that the parametrisation of anisotropy can also be done with 2-D anisotropic autocorrelation functions as can be found in the field of geostatistics. In such cases, the corresponding filter in the Fourier space is obtained with the Wiener–Khinchine theorem (Eq. 6).

In a similar way as in Fig. 4, the set of arbitrary parameters of the GSI model was spatially varied in order to produce a target image with changing anisotropy. The result of the simulation is presented in Fig. 5a. We tested again the ability of the global and local non-parametric generators to reproduce the spatial heterogeneities introduced in the target field and the results are presented in Fig. 5b and c. While generally the global generator could only learn and reproduce the same average anisotropy, the local approach managed to localise the

distinct patterns of the field. In particular, the method is effective in producing the correct angle of anisotropy.

In the non-parametric local realisations presented in Figs. 4c and 5c, one can notice a certain contraction of the contours of the autocorrelation functions compared to the references in the target field. This means that the non-parametric realisations exhibit slightly lower correlation ranges. We believe that this is caused by a combination of minor source of errors associated with the use of the Hanning window and imperfections in the white noise. Hardly preventable, these minor deviations are the integral part of the method. It should also be noted that some degree of localisation can still be found in the result of the global non-parametric simulations. This is believed to be the effect of the residual phase information contained in the global Fourier spectrum used for the white noise filtering (Sect. 2.2). However, this will not be the case with a parametric filter, as no phase information would be available.

4.1 Effect of window size

The side length of the window function used to estimate the local Fourier spectrum is expected to affect the accuracy of the estimation. In fact, if too large a window is used, then the localisation gets less and less informative and the assumption of stationarity within the window becomes weaker. The global approach represents a particular case, when the window size equals the image size. On the other hand, if a too small a window is used, we lose information at wavelengths larger than the window itself and the filter may become ill-defined due to a limited sample size.

In Fig. 6, non-zero synthetic realisations of $512 \text{ px} \times 512 \text{ px}$ were produced using Eq. (21) and $\beta = -3.0, -2.5, -2.0$ and -1.5 . We then applied the SSFT method with decreasing window sizes of 512, 256, 128 or 64 pixel width to produce 20 stochastic realisations for each of the windows. Finally, the average 1-D power

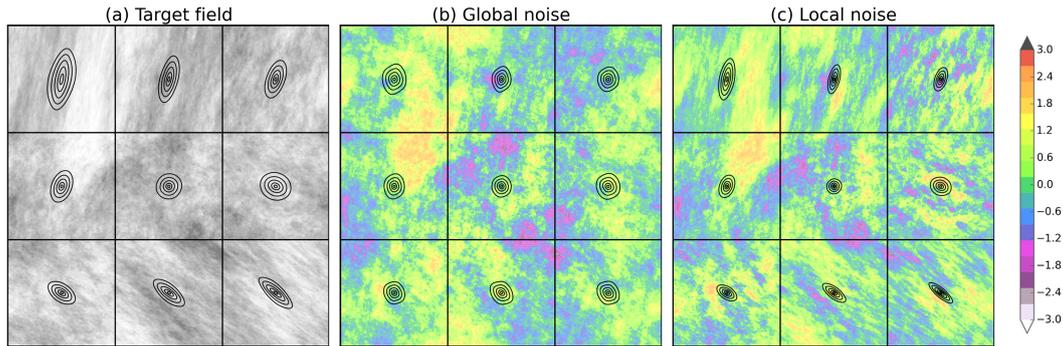


Figure 5. (a) An example of $512 \text{ px} \times 512 \text{ px}$ non-stationary correlated Gaussian noise generated using parametric filtering. The GSI parameters (see Niemi et al., 2014, for details) linearly evolve from $G(-0.12, -0.12, -0.12)$ and $I_s = 5$ at the upper-left corner to $G(0.12, 0.12, 0.12)$ and $I_s = 5$ at the lower-right corner of the image. The black contours show the corresponding autocorrelation functions at 0.5, 0.6, 0.7, 0.8 and 0.9 correlation coefficients. Non-parametric stochastic realisations of the target field using a global generator (b) and a local SSFT generator (c) with Hanning windows centred over the grid boxes in the image. All simulations have been drawn from the standard normal distribution and share the same random seed.

spectra of the resulting fields of correlated noise were plotted in Fig. 6a, alongside the benchmark representing the 1-D power spectra of the original image. In order to facilitate the comparison, all spectra were centred by removing the mean value. The idea was to investigate the impact of decreasing window sizes on the capacity to reproduce the global power spectra of the original image. The effect of reducing the window size is visible in Fig. 6 as we observe a loss of power for wavelengths larger than the window itself. The magnitude of the deviation is more and more important for smaller window sizes and can therefore represent a limitation in the reproduction of scales larger than the chosen window size. However, Fig. 6 also shows that the effect gets less important as the spectral slope of the original field decreases. This outcome suggests that the choice of the window size, or rather its lower limit, should consider the spatial characteristics of the target field, namely the presence of large-scale features. The same is observed with real precipitation fields as discussed later (Sect. 5.3).

5 Weather radar rainfall fields

Precipitation fields from the MeteoSwiss operational radar product for QPE (Germann et al., 2006) were employed for testing the method with real data. In 2015, this weather radar product corresponded to the Cartesian composite of four C-band Doppler radars equipped with dual polarisation (Germann et al., 2015). It should be noted that the radar on the Weissfluhgipfel became operational only in 2016 and it is thus not included in this analysis (Fig. 7).

The raw volumetric radar data undergo a sophisticated signal processing in order to guarantee the best possible precipitation estimates at ground level in a predominately alpine environment. The processing chain includes ground clutter removal by means of a clutter-elimination decision tree includ-

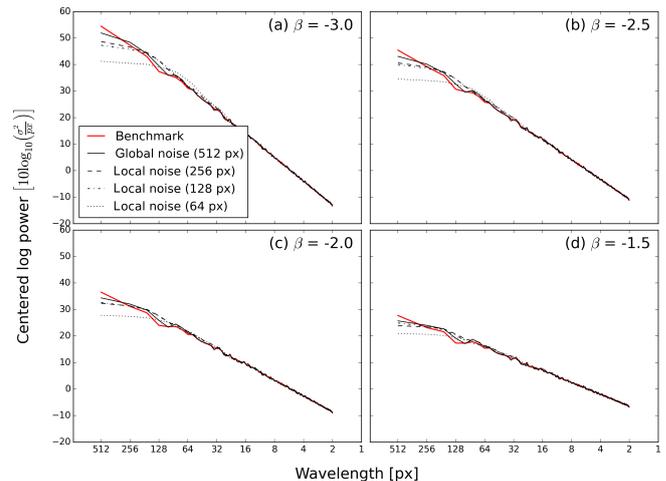


Figure 6. The radially averaged 1-D centred power spectra of synthetic realisations with spectral slope (a) $\beta = -3.0$, (b) $\beta = -2.5$, (c) $\beta = -2.0$ and (d) $\beta = -1.5$ (in red). In black are the corresponding average spectra of 20 non-parametric realisations for decreasing window sizes. The same experiment was realised with real precipitation fields (see Fig. 11).

ing dual-polarisation moments, visibility correction, correction for the vertical profile of reflectivity (VPR), global and local bias correction. For more details, the reader is referred to Germann et al. (2006).

Spatial resolution of the final Cartesian radar grid is 1 km and the temporal sampling is 5 min. In this study, the original images are cropped to a $512 \text{ km} \times 512 \text{ km}$ domain as illustrated by the red box in Fig. 7. Rain rates are converted back to reflectivity (dBZ) using the operational MeteoSwiss Z-R relationship:

$$\text{dBZ} = 10 \log_{10} \left(316 R^{1.5} \right), \quad (22)$$

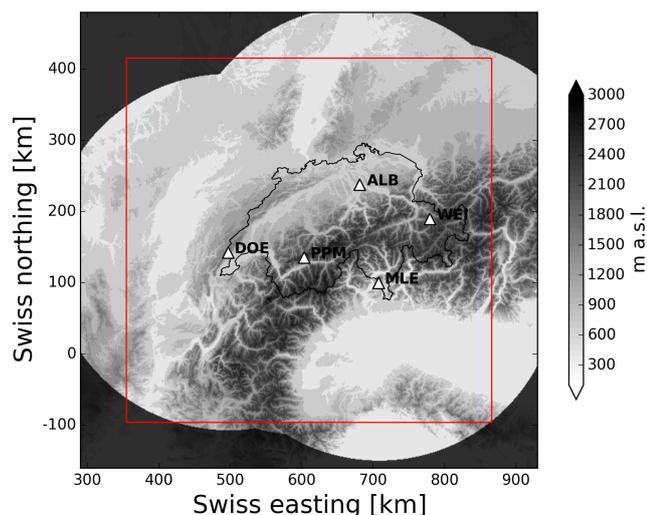


Figure 7. The 512 km \times 512 km study domain (red box) plotted over the Alpine orography, the Swiss national borders and the radar composite mask. Locations of the 4th generation operational MeteoSwiss weather radars are reported as white triangles. ALB: Albis (925 m a.s.l.); DOE: La Dôle (1675 m a.s.l.); MLE: Monte Lema (1625 m a.s.l.); PPM: Pointe de la Plaine Morte (2942 m a.s.l.); WEI: Weissfluhgipfel (2850 m a.s.l., operational since 2016).

where R is the rain rate in mm h^{-1} . Given the skewed distribution of rain rates, the logarithmic transformation leads to approximately normally distributed data, which provides better conditions for the computation of the Fourier transform. The rain/no-rain threshold is set at 0.08 mm h^{-1} (corresponding to 8.54 dBZ). In order to reduce the discontinuity produced by the thresholding, we opted for subtracting 8.54 dBZ in rainy pixels, while keeping all the non-rainy pixels to zero.

The choice of the rainfall threshold and of the value associated with zero rain in dBZ is quite critical and affects the Fourier spectrum, in particular at high spatial frequencies. In fact, a larger step at the rain/no-rain transition has the effect of increasing the power at high frequencies, thus decreasing the absolute value of the spectral slope β . As a consequence, the generated noise fields will have more power at high spatial frequencies (more spatial detail and shorter correlation structure).

5.1 Non-stationary rainfall cases

The first case study is taken from an event on 30 March 2015 and is presented in Fig. 8a. The main north-westerly flow is associated with strong orographic blocking on the northern slopes of the Alps, producing persistent stratiform precipitation on a large portion of the country. At the same time, a squall line develops to the north and rapidly moves to the south-east following the displacement of the cold front. These two features display contrasting spatial characteristics: the pre-frontal, orographic precipitation has low intensity and a long range of spatial autocorrelation, while the cold front

produces convective activity organised on a straight line with a different orientation angle. The squall line is clearly visible in the 2-D power spectrum as a horizontal line passing through the centre of the spectrum (Fig. 8a, centre). In the 2-D autocorrelation function in the bottom panel of Fig. 8a, these two distinct patterns translate at larger scales into a main anisotropy of approximately 45° (anticlockwise angle of rotation from due east) and into a different one at small scales with approximately 0° angle.

The second case study is from 15 May 2015 at 16:00 UTC and is illustrated in Fig. 8b. The south-easterly flow pushes convective cells to the southern Alps where they are blocked by the mountain ridge. At the same time, narrow bands of precipitation are distributed on the northern and western portions of the image while moving south-west. Overall, the apparent motion field has a anticlockwise rotation about the centre of the image and is responsible for the appearance of the local anisotropies. These two distinct precipitation structures have pronounced anisotropies at 90° from each other meaning that they cancel out once averaged together over the field, finally resulting in a fairly isotropic 2-D autocorrelation function, as shown in the bottom panel of Fig. 8b.

The third case study belongs to an event on 8 June 2015 at 02:30 UTC (Fig. 9a). The image is characterised by a distinct active line of thunderstorms along the southern Prealps. This is originated from weak easterly currents bringing humid air towards the Alps that is then destabilised by the transition of an occluded front. Precipitation in the northern side of the Alps is instead associated with a low pressure at ground level. These two main features have strong differences in both orientation and correlation ranges that are caused by the presence of the Alps.

The fourth and final case study is from 15 June 2015 at 11:45 UTC (Fig. 9b). This case is not characterised by any distinct large-scale anisotropy as the precipitation patterns look almost isotropic. Instead, it is easy to notice a north-south gradient in the autocorrelation ranges of the rainfall patterns. In fact, the northern side of the Alps is principally characterised by patches of stratiform precipitation of low intensity, while in the southern part of the image there are small-scale precipitation structures only in correspondence of topographical features.

5.2 Reproduction of local non-stationarities

Both global and local noise generators were tested to detect and model the local non-stationarities in the weather radar images presented above. The local SSFT generator was applied with a window side length of 256 km (i.e. half of the original field size), a Hanning window function and no overlapping. Such settings allow for a clear representation and interpretation of the results. The global FFT generator can be seen as a particular case of the local SSFT generator with the window side length equal to the field size, 512 km in this case.

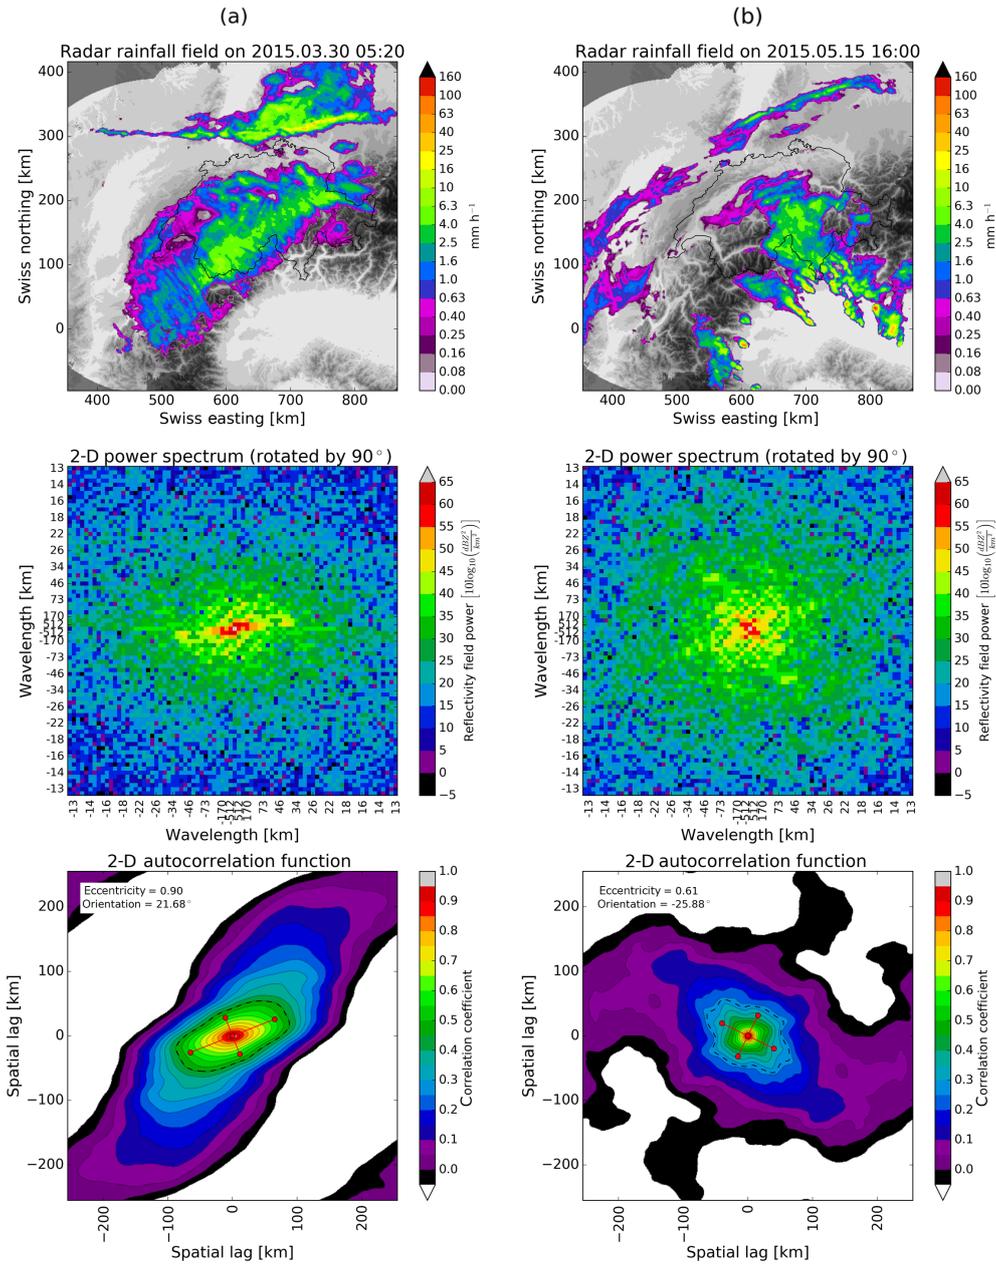


Figure 8. Radar rainfall fields (top row), 2-D Fourier spectra zoomed on wavelengths > 13 km and rotated by 90° (centre row) and corresponding 2-D autocorrelation functions (bottom row) for case studies on (a) 30 March 2015 and (b) 15 May 2015. The autocorrelation function is obtained as inverse FFT of the 2-D power spectrum. The 90° rotation is performed in order to align the anisotropies of the 2-D spectra and spatial autocorrelation functions. The anisotropy is estimated by eigenvalue decomposition of the autocorrelation function using values only inside the dashed region and shifted to start from 0. The eccentricity is obtained as $\sqrt{\frac{1-a_{\text{minor}}}{a_{\text{major}}}}$ and is constrained between 0 (perfect circle) and 1 ($a_{\text{minor}} \ll a_{\text{major}}$).

Figure 10 presents for all case studies the 2-D autocorrelation functions within the four windows of 256 km side length as shown by the grid. The latter is computed and averaged over 20 realisations of global and local noise, while only the first realisation is shown in the background as an example. It should be noted that the notion of local is defined

by the choice of the window size, as there is the assumption of homogeneity of statistical properties within a single window. The visual comparison shows that for all four cases, only the local approach can effectively reproduce the non-stationarities as defined in the reference images on the left. It is easy to notice how the global generator systematically pro-

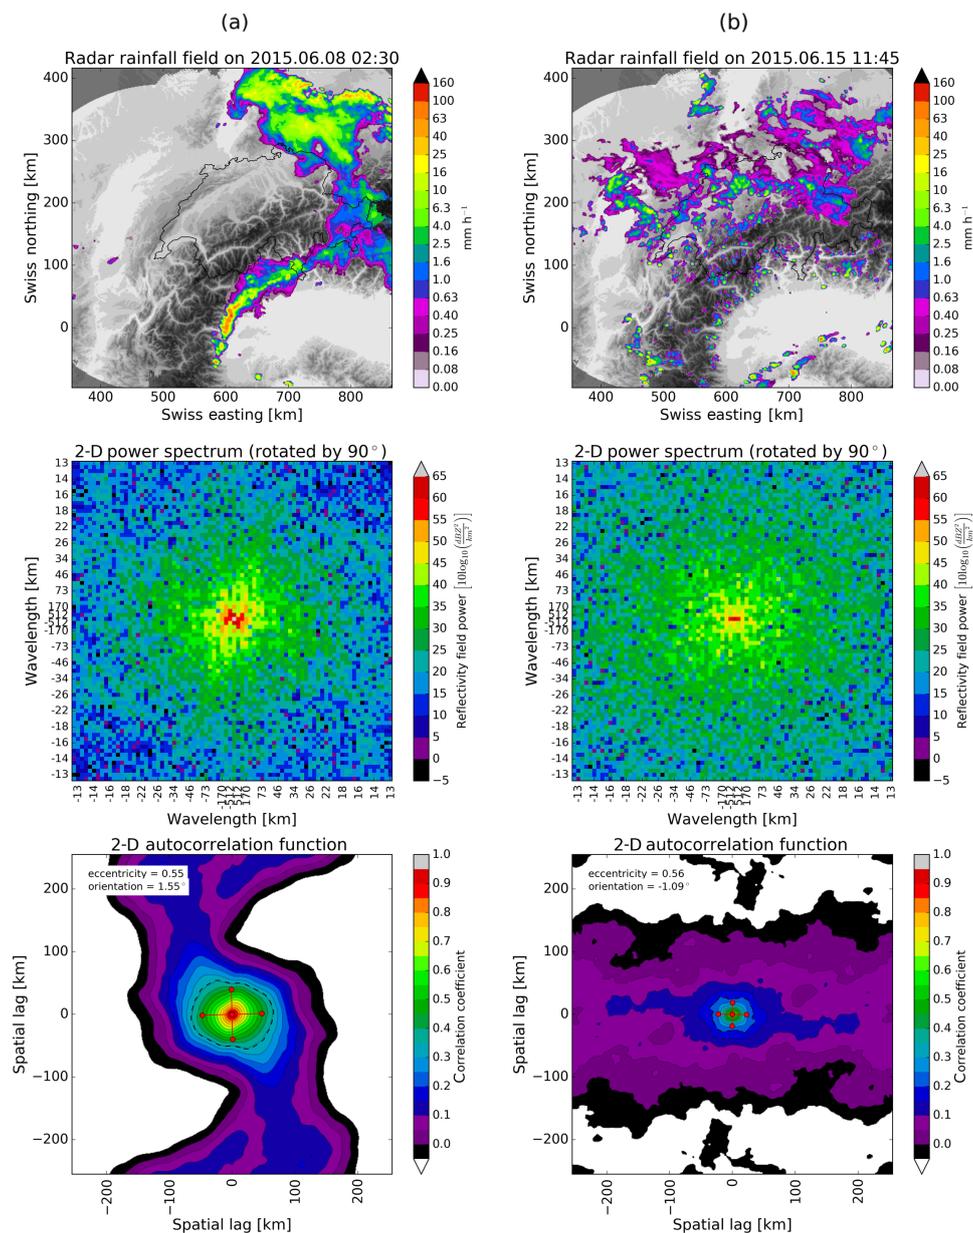


Figure 9. Radar rainfall fields (top row), 2-D Fourier spectra zoomed on wavelengths > 13 km and rotated by 90° (centre row) and corresponding 2-D autocorrelation functions (bottom row) for case studies on (a) 8 June 2015 and (b) 15 June 2015. More details are provided in the caption of Fig. 8.

duces average autocorrelation functions that are driven by the main anisotropies observed in the field. As already discussed in the synthetic examples in Sect. 4, small differences between the reference 2-D autocorrelation functions and those computed from the obtained SSFT noise fields are believed to be linked to sources of errors caused by the windowed approach, such as the use of the Hanning window.

Arguably, even for a 256 km window the assumption of stationarity appears as weak and this is noticeable by the

presence of abrupt changes between some of the adjacent windows. The use of smaller windows and the introduction of some degree of overlapping between successive windows will guarantee a smoother transition over the field.

5.3 Reproduction of the global power spectrum

The goal of the global FFT noise generator is to impose the same power spectrum observed in the precipitation radar im-

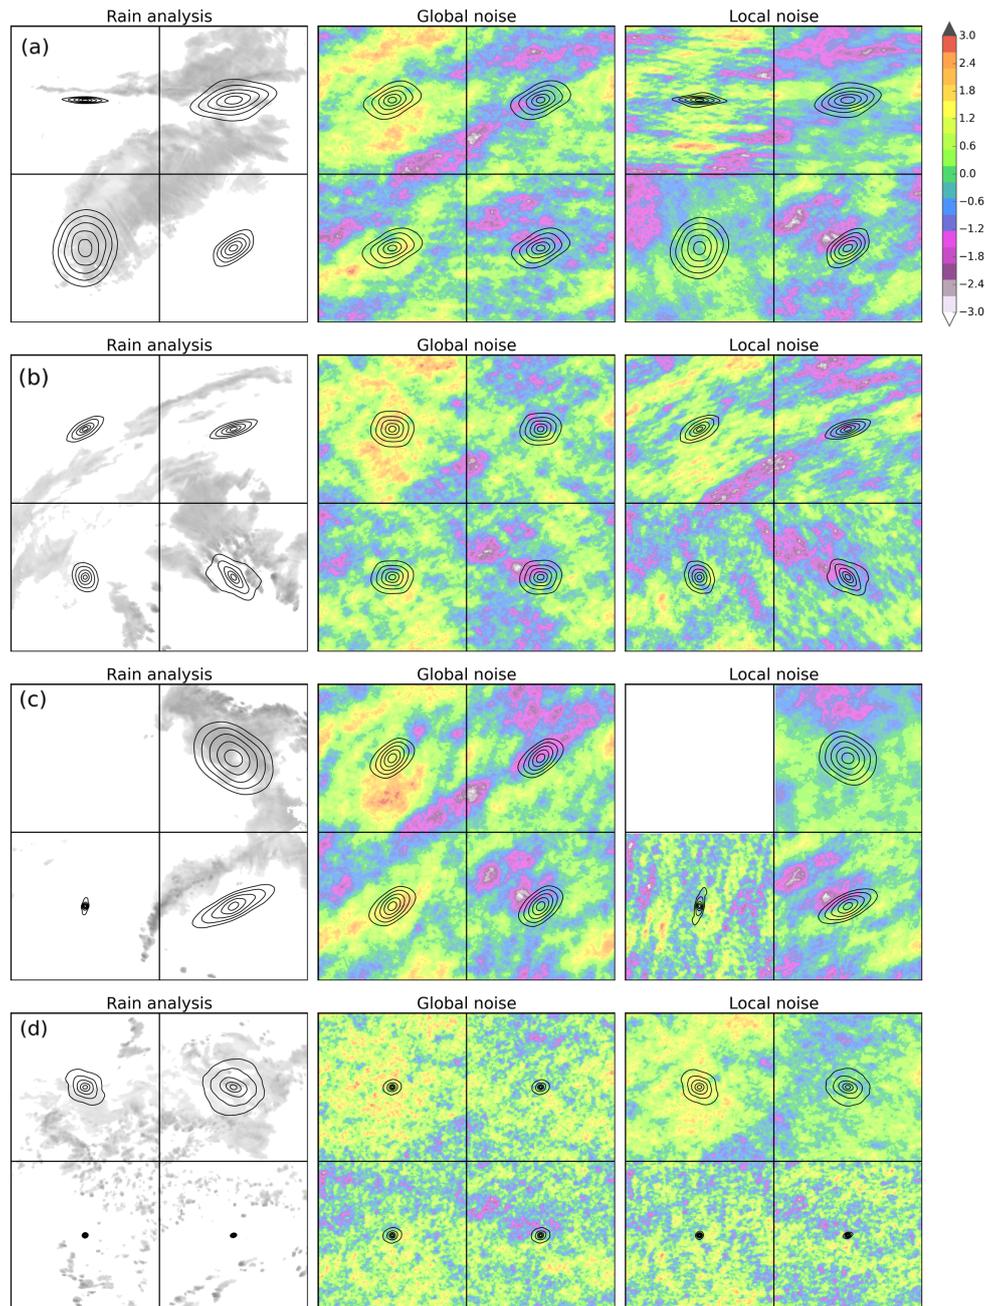


Figure 10. Localised 2-D autocorrelation functions (black contours for 0.5, 0.6, 0.7, 0.8 and 0.9 correlation coefficients) for all four case studies (rain rates in black and white) and the corresponding non-parametric stochastic simulations using both the global and local generators. The image domain is $512 \text{ km} \times 512 \text{ km}$. All noise fields have been drawn from the standard normal distribution and share the same random seed.

age to a field of Gaussian white noise. The local SSFT noise generator has the advantage of being able to localise this operation in order to account for non-stationarities. However, a legitimate question is whether the local SSFT noise generator is still able to reproduce the global power spectrum, despite representing the precipitation as a convolution of localised Fourier transforms. The analysis over synthetic data

(see Fig. 6) showed that there is a loss of power in the spectrum for wavelengths larger than the window size. The same analysis can be done with the precipitation fields presented above, thus introducing all the complexities associated with real data. For this verification we computed the radially averaged 1-D power spectra of the target rainfall field, the global FFT noise and the local SSFT noise using three different win-

dow sizes (264, 128 and 64 km). For the sake of comparison, all the power spectra were centred to have a zero mean.

The 1-D power spectra in Fig. 11 show that overall both global and local FFT approaches manage to impose to uncorrelated white noise a similar scaling behaviour as observed in the reference rain analysis. However, it is easy to notice the loss of power at large wavelengths associated with the reduction of the window size. Moreover, the magnitude of the deviation varies between the four case studies. Interestingly, the last case study in Fig. 11d displays the least significant deviations at large wavelengths, while it is also associated with the lowest spectral slope (i.e. the least negative β), as the precipitation field is mainly dominated by small-scale features (see Fig. 9b). This case relates well to the synthetic experiment of Sect. 4.1, where we observed lower deviations for increasing (i.e. less negative) β . Again, these results support the idea that the choice of the optimal window size is case dependent.

5.4 Stochastic rainfall fields using local noise adjustment

A first application of the SSFT generator is related to design storms. In order to dimension the size of hydraulic infrastructure (sewer systems, hydroelectric dams, artificial banks and bridges, etc.), it is important to model the maximum expected flood and discharge under different weather situations. The probable maximum precipitation (PMP) can be computed by running NWP models using extreme input conditions or by stochastic simulation of a large number of synthetic rainfall fields, e.g. using statistics from extreme events that occurred in the past.

The FFT noise and SSFT noise fields already reproduce the power spectrum and spatial autocorrelation function of the rainfall fields, but they do not replicate the same marginal statistics (empirical statistical distribution, mean and variance, wet area ratio, etc.). There are different techniques that are used in order to transform the noise fields so that they reproduce the statistics of rainfall fields. A well-established approach is to threshold the noise field to have the same wet area ratio (WAR) of the target rainfall field and to rescale the non-zero values to have the same marginal mean precipitation, usually known as “shift and scale” (Pegram and Clothier, 2001a, b; Berenguer et al., 2011). A similar method attempts to match the observed mean and standard deviation of the precipitation field (Seed et al., 1999; Niemi et al., 2014). More sophisticated methods apply a quantile–quantile transformation to match exactly the same empirical distribution of observed precipitation values, which is also known as probability matching or anamorphosis (Metta et al., 2009; Schleiss et al., 2014).

All these transformation techniques share the same dilemma; i.e. changing the marginal statistics of the noise fields perturbs their power spectra (Metta et al., 2009). More precisely, the noise fields accurately reproduce the power spectrum of the precipitation fields but not their marginal statistics. On the other hand, the simulated precipitation fields accurately reproduce the marginal statistics but the power spectra are slightly perturbed. Following the approach of Pegram and Clothier (2001a, b), a local adjustment of the noise is applied.

1. Observed wet area ratio (WAR_0), marginal mean (MM_0) and marginal standard deviation (MSD_0) are computed based on the observed radar rainfall field, Z_0 , in dBZ units.
2. From the noise field, Z_N , a threshold T is computed as the $(1 - \text{WAR}_0)$ th percentile, such that the number of pixels above or equal to T is the same as the number of wet pixels in the observed field.
3. The marginal mean and marginal standard deviation with respect to T (thus MM_T and MSD_T) are computed from Z_N .
4. The noise field Z_N and the threshold T are renormalised to match the marginal distribution of Z_0 as follows:

$$Z_N^* = \frac{(Z_N - \text{MM}_T)}{\text{MSD}_T} \text{MSD}_0 + \text{MM}_0, \quad (23)$$

$$T^* = \frac{(T - \text{MM}_T)}{\text{MSD}_T} \text{MSD}_0 + \text{MM}_0. \quad (24)$$
5. Finally, all values in Z_N^* below T^* are set to zero and the field can be back transformed into rain rates by inverting the Z - R relationship in Eq. (22).

Steps 1 to 5 can be applied at once to the whole image or more locally within smaller windows. In this case, values of MM_0 , MSD_0 , MM_T , MSD_T and T^* need to be assigned to every pixel by interpolation to avoid the appearance of discontinuities in the rescaled fields. As final adjustment, a global quantile–quantile transform (probability matching) is applied to reproduce the cumulative distribution of the observed rain rates.

We applied the above procedure at a resolution of 64 km to a number of realisations of correlated Gaussian noise generated using the rainfall field of case study 1 (30 March 2015). The original field and the three stochastic realisations are presented in Fig. 12 for both globally and locally generated noise. The use of a locally modulated noise allows one to better reproduce some of the specific features of the original rainfall field, such as the squall line in the upper portion of the image. Not all realisations produced this local feature with the same realism and still some residual noise affects the quality of the output. Nevertheless, we consider the examples in Fig. 12 as a significant improvement in the context of stochastic simulation of precipitation fields.

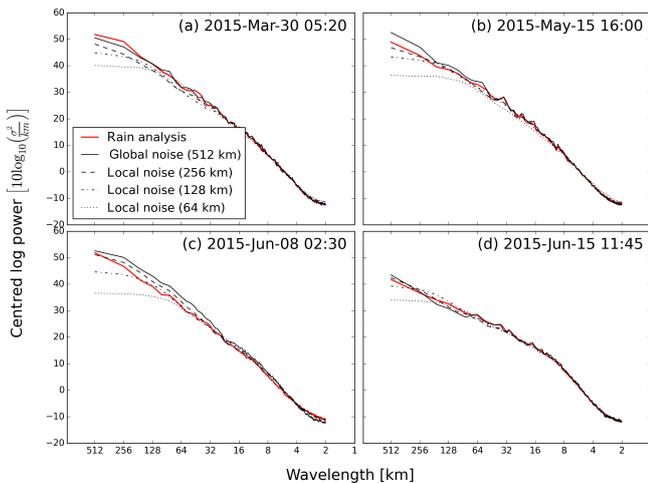


Figure 11. The 1-D centred power spectra of precipitation fields for the four case studies are represented in red. In black are the corresponding average spectra of 20 non-parametric noise field realisations with decreasing window sizes.

5.5 Imposing the temporal structure of precipitation: the nowcasting example

Probabilistic precipitation nowcasting represents a common application of fields of stochastic noise and it can be helpful to illustrate how to simulate the temporal structure of precipitation. The unknown temporal evolution of precipitation due to growth and decay processes (often considered in Lagrangian coordinates) is usually modelled by auto-regressive models (AR) of the order of 1 or 2 (see Pegram and Clothier, 2001a, b; Seed, 2003; Bowler et al., 2006; Berenguer et al., 2011; Seed et al., 2013; Atencia and Zawadzki, 2014). A time series following a first-order auto-regressive process (AR(1)) can be constructed as follows:

$$X_{t+1} = \rho X_t + \sqrt{1 - \rho^2} G(0, 1), \quad (25)$$

where t is the current time step, X is the standardised variable to be predicted, ρ is the temporal autocorrelation of the time series and $G(0, 1)$ is the shock term, which in the case of a single 1-D time series is taken as simple white Gaussian noise with zero mean and unit variance. In the context of nowcasting, ρ is an estimation of the Lagrangian autocorrelation of radar rainfall fields, which is a measure of predictability for the rate of development of rainfall due to growth and decay processes in moving coordinates (Germann et al., 2002). A cascade of auto-regressive processes can also be considered to account for the scale dependence of the predictability of precipitation as done in STEPS (Seed, 2003; Bowler et al., 2006).

The problem becomes more complicated when trying to model the temporal evolution of 2-D precipitation fields, which requires preserving the joint spatial correlation between different points in the domain. In other words, using

a 2-D field of white noise as shock term would end up destroying the spatial structure of the precipitation field after a few time steps. The shock term should serve as a model for the expected forecast errors, which are known to be correlated in both space and time. The classical approach is to filter a field of stationary FFT noise using an AR(1) model that evolves with the same Lagrangian autocorrelation of the precipitation field. This generates a time series of spatially and temporally correlated noise fields that can be used as shock terms in Eq. (25) (Bowler et al., 2006). The auto-regressive filtering of correlated noise fields allows one to condition the future rainfall realisations to the last observed rainfall field and to effectively model the growth of forecast errors by imposing appropriate spatial and temporal correlations.

However, in the presence of non-stationary precipitation fields, it is not adequate to simply use stationary FFT noise, which does not represent the local characteristics of the field. In fact, as the time progresses the global noise will gradually destroy the local anisotropies and inject too much variability within areas of smoother stratiform rain, or conversely not enough variability in convective regions. Therefore, we expect the local SSFT noise to provide a more realistic and locally adaptive shock term.

In order to test this hypothesis, we generated two time series of stochastic precipitation fields, one using the global and one using the local shock term. Figure 13 illustrates an example of stochastic precipitation nowcasting using a simple AR(1) model to simulate the temporal variability of rainfall due to unknown growth and decay processes. The correlation of the AR(1) process has been set to the Eulerian temporal autocorrelation with the previous radar rainfall field. The visual comparison for increasing lead times (+5, +30 and +60 min) highlights the advantages of using a locally generated shock term. In fact, the evolution of the image looks more realistic when local noise is employed, as the non-stationarity is preserved both in terms of anisotropy and correlation length. When a globally generated noise is used in the presence of strong non-stationarities, spatial patterns are quickly destroyed, as the image spatial structure converges towards a global average that looks artificial. In other words, thanks to the local modulation of the shock term, growth and decay processes take place in a spatially consistent way so that new cells appear with the same anisotropy and spatial autocorrelation as existing nearby cells. On the other hand, a marked improvement in forecast skills is yet to be demonstrated, as performance scores are expected to be driven by the ability of the forecast system to predict the rainfall fraction and average intensity over a given area. For instance, Gyasi-Agyei (2016) did not find improvements in the kriging interpolation of daily rain gauges when including radar-based locally varying anisotropy.

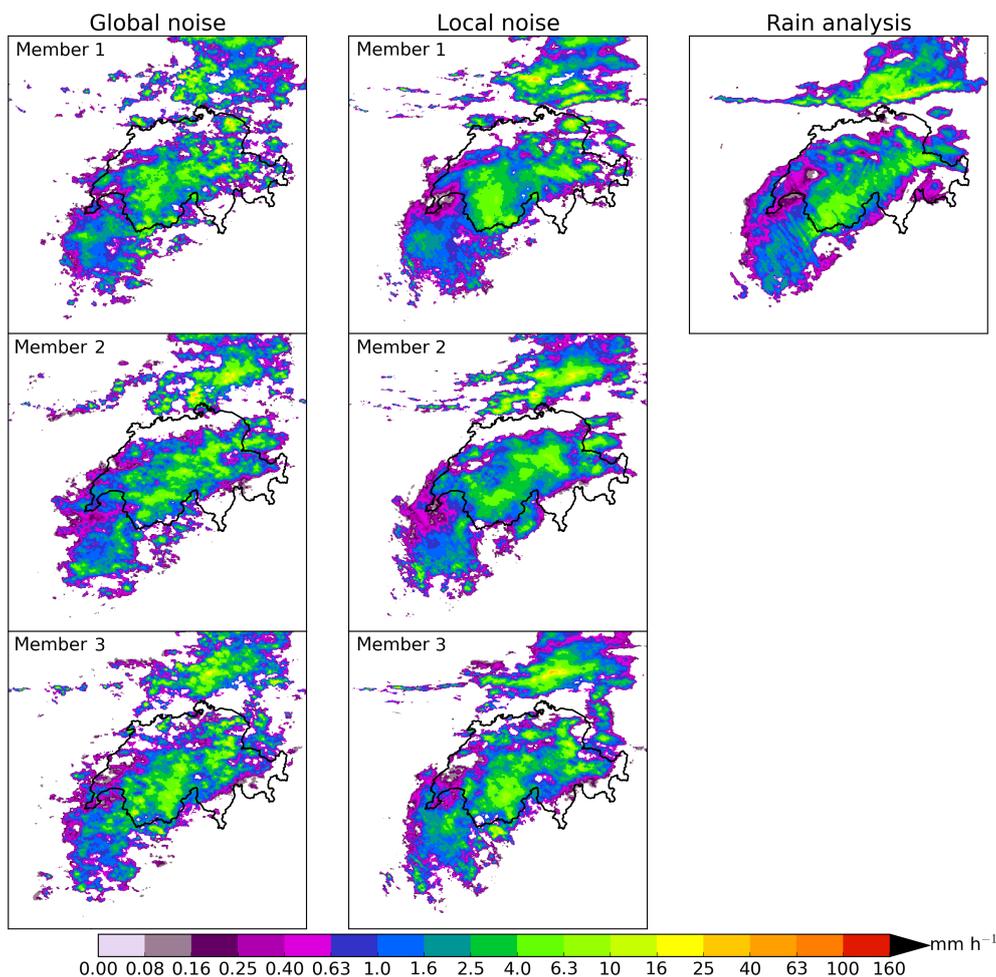


Figure 12. Three stochastic realisations conditioned to case study 1 (30 March 2015) after local adjustment at 64 km resolution of correlated global FFT (left) and local SSFT (centre) Gaussian noise. The top right panel shows the actual observed field. Parameters of the local SSFT generator: window size = 128 km, overlapping = 50 %. The national borders of Switzerland are included as spatial reference (black contour).

6 Summary and conclusions

In this paper we proposed a novel approach for the generation of non-stationary stochastic rainfall fields. The main idea behind this work consists in the simulation of a stochastic rainfall field in order to preserve not only global spatial correlation properties of the target precipitation field, but also the spatial distribution of its local correlation ranges and anisotropies. This is achieved by replacing the 2-D Fourier transform in both non-parametric and parametric conventional filtering of white noise, with the short-space Fourier transform (SSFT). Namely, instead of convolving the white noise field with the entire precipitation field, we perform the convolution locally, by employing a carefully parametrised moving window. This way, we obtain a spatially correlated noise field, which locally reproduces the spectral properties of the observed rainfall field, i.e. which conserves its local correlation ranges and anisotropies.

We illustrated the potential of the non-stationary generator with synthetic and real precipitation data and demonstrated the advantages of the local approach in the presence of non-stationarities. We showed how the non-stationary correlated noise can be locally adjusted to generate realistic stochastic simulations of radar precipitations fields. As a final experiment, we introduced the temporal structure of rainfall in the context of precipitation nowcasting using an auto-regressive process. We demonstrated that the local shock term can help to preserve the spatial heterogeneities of the original field during the simulation of the temporal evolution of precipitation.

The negative effect of localising the Fourier transform is a loss of power at scales larger than the window size, whose magnitude depends on the global spectral slope. In other words, precipitation fields displaying large-scale patterns suffer more from such loss of power. Ways to account for such an effect are currently under study. Other sources of errors associated with the method itself, such as the use of a

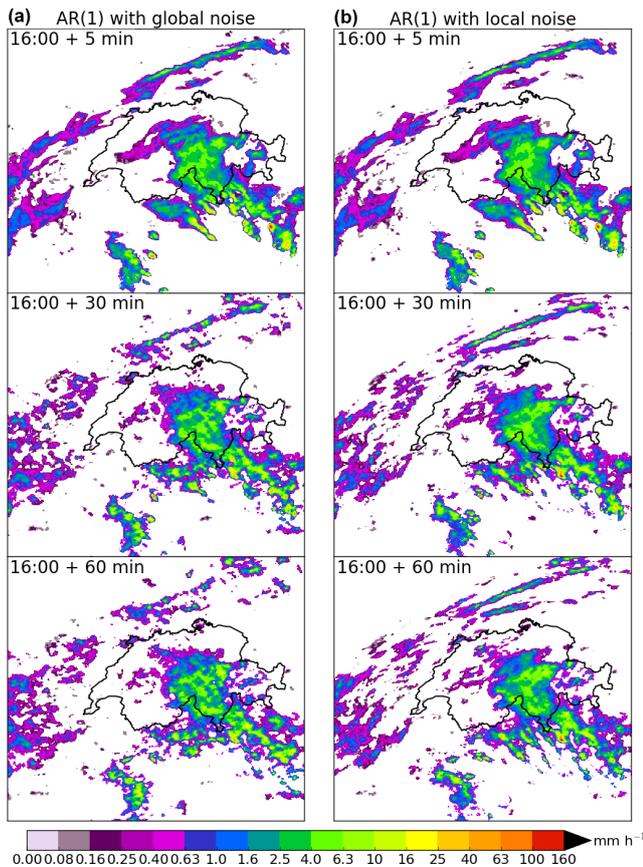


Figure 13. Example of simplified stochastic rainfall nowcasting for case study 2 (15 May 2015). Here the advection component is neglected and the movement of the rainfall field is assumed to be zero. Growth and decay processes are simulated stochastically using an AR(1) process with a globally (a) or locally (b) generated shock term. The lag-1 autocorrelation coefficient was set equal to 0.94. A local adjustment as in Sect. 5.4 was included at 64 km resolution. The national borders of Switzerland are included as spatial reference (black contour). Both simulations share the same random seed.

Hanning window, were found to slightly affect the reproduction of local autocorrelation structure. However, we believe that these limitations are compensated for by the benefits introduced by the new approach.

7 Future perspectives

The non-stationary stochastic generator was designed in the context of ensemble quantitative precipitation estimation and nowcasting, blending and downscaling of NWP models, and design storms modelling. As already mentioned in the Introduction (Sect. 1.1), several stochastic rainfall models based on the global Fourier filtering of white noise can be reformulated using the local non-stationary SSFT generator.

The residual radar measurement uncertainty can be effectively modelled by the non-stationary generator to obtain QPE ensembles, which reproduce the local statistical characteristics and anisotropy of the observed rainfall fields (Jordan et al., 2003; Ciach et al., 2007; Villarini et al., 2009; Germann et al., 2009; Cecinati et al., 2017). In addition, we believe that current radar rain gauge merging and adjustment techniques (e.g. Sideris et al., 2014) can be further extended to generate QPE ensembles conditioned onto the rain gauges, for example by following the conditional merging approach of Sinclair and Pegram (2005).

A demonstrative application of the local SSFT generator for stochastic precipitation nowcasting was presented in Sect. 5.5. A complete nowcasting model would also need an optical flow technique for the Lagrangian extrapolation of the stochastic radar rainfall fields. In the example presented in this paper, the stochastic perturbations are generated by using the same set of local Fourier spectra estimated at analysis time for all forecast lead times. An interesting scientific question would be to analyse the persistence and predictability of the local Fourier spectra, which in turn will control the future evolution of the local properties of the stochastic rainfall fields. This question is particularly important to design ensemble precipitation nowcasting systems that better represent the forecast uncertainty.

For longer forecast ranges (e.g. 24–48 h), the statistical properties of the radar rainfall field at analysis time become even more obsolete and should not be used to generate the stochastic perturbations. Current FFT-based downscaling techniques simplify the problem by using a fixed climatological isotropic parametric filter (Rebora et al., 2006a; Pierce et al., 2010). As a consequence all the downscaled NWP fields share the same small-scale precipitation features, independently on the NWP field that needs to be downscaled. A more flexible approach would involve the generation of the stochastic perturbations that adapt to the local properties and anisotropy of the NWP forecast.

Finally, the extension of the non-stationary generator for design storms would need additional components to model the storm arrival process (duration of dry and wet spells) and the temporal evolution of the mean areal statistics (global IMF and WAR); see e.g. Paschalis et al. (2013).

The SSFT generator makes a certain number of assumptions and can be further developed. In its current implementation the rainfall field is convolved by rectangular or Hanning windows of fixed size, which assumes the stationarity of rainfall properties within the window. However, the dimension or length of stationary regions within a rainfall field could also depend on spatial location. For example, there may be a small region of 50 km × 50 km characterised by cellular convection in one part of the domain and another region of 200 km × 100 km with frontal stratiform rain. Adaptive windows that dilate and contract according to the length of the stationary region could be an interesting extension of the method. Given the link between the short-time Fourier

transform (STFT) and the wavelet decomposition (Kovačević et al., 2013), it would be interesting to develop a non-stationary wavelet-based stochastic noise generator for comparison with the local SSFT approach.

Another interesting problem that deserves more research relates to the choice of the optimal window type. In particular, the use of different windows, for instance a top flat Hanning, might help to improve the reproduction of the true local correlation structure of rainfall.

A practical problem of the local Fourier analysis is the inevitable loss of power at wavelengths larger than the window size, which becomes more accentuated for decreasing window sizes. For smaller window sizes, the local Fourier spectrum could be parametrised and the scaling behaviour extrapolated to scales larger than the window size to compensate for the loss of power. Another option that we are currently exploring is the implementation of a nested approach where the rainfall field is subsequently divided into smaller and smaller boxes for Fourier analysis. In this way, the information coming from boxes at different sizes can be merged into one optimal local Fourier spectrum.

Finally, the Fourier filtering of white Gaussian noise is a particular case of a broader class of continuous-in-scale multifractal simulation methods (Lovejoy and Schertzer, 2010). Instead of filtering white Gaussian noise, the universal multifractal simulation techniques are based on the filtering of Lévy noise, which provides more flexibility for modelling of multiscaling precipitation fields. An interesting extension of the method would be to exploit the SSFT approach for non-stationary continuous-in-scale multifractal simulation.

Data availability. We have provided on GitHub a basic implementation of the code in Python together with the four radar rainfall images used in this study. This supplementary material is available at the following address: <https://doi.org/10.5281/zenodo.803285>.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This study was supported by the Swiss National Science Foundation Ambizione project “Precipitation attractor from radar and satellite data archives and implications for seamless very short-term forecasting” (PZ00P2_161316). We thank Isztar Zawadzki, Geoff Pegram and the anonymous reviewer for the comments and fruitful discussion. We are also grateful to Tero Niemi for providing the basic code and explanations on the generalised-scale invariance and to Carlos Velasco-Forero for his comments.

Edited by: R. Uijlenhoet

Reviewed by: G. Pegram and one anonymous referee

References

- Atencia, A. and Zawadzki, I.: A comparison of two techniques for generating nowcasting ensembles. Part I: Lagrangian ensemble technique, *Mon. Weather Rev.*, 142, 4036–4052, 2014.
- Badas, M. G., Deidda, R., and Piga, E.: Modulation of homogeneous space-time rainfall cascades to account for orographic influences, *Nat. Hazards Earth Syst. Sci.*, 6, 427–437, <https://doi.org/10.5194/nhess-6-427-2006>, 2006.
- Bárdossy, A. and Hörning, S.: Gaussian and non-Gaussian inverse modeling of groundwater flow using copulas and random mixing, *Water Resour. Res.*, 52, 4504–4526, 2016.
- Berenguer, M., Sempere-Torres, D., and Pegram, G. G. S.: SBMcast – an ensemble nowcasting technique to assess the uncertainty in rainfall forecasts by Lagrangian extrapolation, *J. Hydrol.*, 404, 226–240, 2011.
- Boisvert, J. B., Manchuk, J. G., and Deutsch, C. V.: Kriging in the presence of locally varying anisotropy using non-Euclidean distances, *Math. Geosci.*, 41, 585–601, <https://doi.org/10.1007/s11004-009-9229-1>, 2009.
- Bowler, N. E. H., Pierce, C. E., and Seed, A. W.: STEPS: A probabilistic precipitation forecasting scheme which merges an extrapolation nowcast with downscaled NWP, *Q. J. Roy. Meteor. Soc.*, 132, 2127–2155, 2006.
- Casari, A., Ramos, M. H., Javelle, P., and Leblois, E.: A space-time geostatistical approach for ensemble rainfall nowcasting, *E3S Web Conf.*, 7, 18001, <https://doi.org/10.1051/e3sconf/20160718001>, 2016.
- Cecinati, F., Rico-Ramirez, M. A., Heuvelink, G. B. M., and Han, D.: Representing radar rainfall uncertainty with ensembles based on a time-variant geostatistical error modelling approach, *J. Hydrol.*, 548, 391–405, 2017.
- Ciach, G. J., Krajewski, W. F., and Villarini, G.: Product-error-driven uncertainty model for probabilistic quantitative precipitation estimation with NEXRAD data, *J. Hydrometeorol.*, 8, 1325–1347, 2007.
- Cooley, J. and Tukey, J.: An Algorithm for the Machine Calculation of Complex Fourier Series, *Math. Comput.*, 19, 297–301, 1965.
- Foresti, L. and Seed, A.: The effect of flow and orography on the spatial distribution of the very short-term predictability of rainfall from composite radar images, *Hydrol. Earth Syst. Sci.*, 18, 4671–4686, <https://doi.org/10.5194/hess-18-4671-2014>, 2014.
- Foresti, L. and Seed, A.: On the spatial distribution of rainfall nowcasting errors due to orographic forcing, *Meteorol. Appl.*, 22, 60–74, 2015.
- Foresti, L., Reyniers, M., Seed, A., and Delobbe, L.: Development and verification of a real-time stochastic precipitation nowcasting system for urban hydrology in Belgium, *Hydrol. Earth Syst. Sci.*, 20, 505–527, <https://doi.org/10.5194/hess-20-505-2016>, 2016.
- Germann, U. and Zawadzki, I.: Scale-Dependence of the Predictability of Precipitation from Continental Radar Images. Part I: Description of the Methodology, *Mon. Weather Rev.*, 130, 2859–2873, 2002.
- Germann, U., Galli, G., Boscacci, M., and Bolliger, M.: Radar precipitation measurement in a mountainous region, *Q. J. Roy. Meteor. Soc.*, 132, 1669–1692, 2006.
- Germann, U., Berenguer, M., Sempere-Torres, D., and Zappa, M.: REAL–Ensemble radar precipitation estimation for hydrology in a mountainous region, *Q. J. Roy. Meteor. Soc.*, 135, 445–456, 2009.

- Germann, U., Boschetti, M., Gabella, M., and Sartori, M.: Radar design for prediction in the Swiss Alps, *Meteorological Technology International*, 4, 42–45, 2015.
- Gires, A., Tchiguirinskaia, I., Schertzer, D., and Lovejoy, S.: Multifractal and spatio-temporal analysis of the rainfall output of the Meso-NH model and radar data, *Hydrolog. Sci. J.*, 55, 380–396, 2011.
- Goovaerts, P.: *Geostatistics for natural resources evaluation*, Oxford University Press, Oxford, UK, 1997.
- Gyasi-Agyei, Y.: Assessment of radar-based locally varying anisotropy on daily rainfall interpolation, *Hydrolog. Sci. J.*, 61, 1890–1902, 2016.
- Hinman, G., Bernstein, J., and Staelin, D.: Short-space Fourier transform image processing, Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '84, 166–169, <https://doi.org/10.1109/ICASSP.1984.1172374>, 1984.
- Jordan, P. W., Seed, A. W., and Weinmann, P. E.: A stochastic model of radar measurement errors in rainfall accumulations at catchment scale, *J. Hydrometeorol.*, 4, 841–855, 2003.
- Khintchine, A.: Korrelationstheorie der stationären stochastischen Prozesse, *Math. Ann.*, 109, 604–615, 1934.
- Kovačević, J., Goyal, V. K., and Vetterli, M.: *Fourier and Wavelet Signal Processing*, Cambridge University Press, Cambridge, 2013.
- Kumar, P. and Foufoula-Georgiou, E.: Wavelet analysis for geophysical applications, *Rev. Geophys.*, 35, 385–412, 1997.
- Lantuéjoul, C.: *Geostatistical simulation: models and algorithms*, Springer-Verlag, Berlin, Heidelberg, 2002.
- Leblois, E. and Creutin, J.-D.: Space-time simulation of intermittent rainfall with prescribed advection field: Adaptation of the turning band method, *Water Resour. Res.*, 49, 3375–3387, <https://doi.org/10.1002/wrcr.20190>, 2013.
- Lovejoy, S. and Schertzer, D.: Generalized scale invariance in the atmosphere and fractal models of rain, *Water Resour. Res.*, 21, 1233–1250, <https://doi.org/10.1029/WR021i008p01233>, 1985.
- Lovejoy, S. and Schertzer, D.: Multifractals, cloud radiances and rain, *J. Hydrol.*, 322, 59–88, 2006.
- Lovejoy, S. and Schertzer, D.: On the simulation of continuous in scale universal multifractals, part I: spatially continuous processes, *Comput. Geosci.*, 36, 1393–1403, 2010.
- Lovejoy, S. and Schertzer, D.: *The Weather and Climate: Emergent Laws and Multifractal Cascades*, Cambridge University Press, Cambridge, UK, 2013.
- Mandelbrot, B. B. and Van Ness, J. W.: Fractional Brownian motions, fractional noises and applications, *SIAM Rev.*, 10, 422–437, 1968.
- Mandapaka, P. V., Germann, U., Panziera, L., and Hering, A.: Can Lagrangian Extrapolation of Radar Fields Be Used for Precipitation Nowcasting over Complex Alpine Orography?, *Weather Forecast.*, 27, 28–49, 2012.
- Mantoglou, A. and Wilson, J. L.: The turning bands method for simulation of random fields using line generation by a spectral method, *Water Resour. Res.*, 18, 1379–1394, 1982.
- Marcotte, D.: Fast variogram computation with FFT, *Comput. Geosci.*, 22, 1175–1186, [https://doi.org/10.1016/S0098-3004\(96\)00026-X](https://doi.org/10.1016/S0098-3004(96)00026-X), 1996.
- Menabde, M., Harris, D., Seed, A., Austin, G., and Stow, D.: Multi-scaling properties of rainfall and bounded random cascades, *Water Resour. Res.*, 33, 2823–2830, 1997.
- Metta, S., Rebor, N., Ferraris, L., von Hardenberg, J., and Provenzale, A.: PHAST: a phase-diffusion model for stochastic nowcasting, *J. Hydrometeorol.*, 10, 1285–1297, 2009.
- Niemi, T. J., Kokkonen, T., and Seed, A. W.: A simple and effective method for quantifying spatial anisotropy of time series of precipitation fields, *Water Resour. Res.*, 50, 5906–5925, <https://doi.org/10.1002/2013WR015190>, 2014.
- Niemi, T. J., Guillaume, J. H. A., Kokkonen, T., Hoang, T. M. T., and Seed, A. W.: Role of spatial anisotropy in design storm generation: Experiment and interpretation, *Water Resour. Res.*, 52, 69–89, <https://doi.org/10.1002/2015WR017521>, 2016.
- Panziera, L. and Germann, U.: The relation between airflow and orographic precipitation on the southern side of the Alps as revealed by weather radar, *Q. J. Roy. Meteor. Soc.*, 136, 222–238, 2010.
- Paschalis, A., Molnar, P., Faticchi, S., and Burlando, P.: A stochastic model for high-resolution space-time precipitation simulation, *Water Resour. Res.*, 49, 8400–8417, 2013.
- Pathirana, A. and Herath, S.: Multifractal modelling and simulation of rain fields exhibiting spatial heterogeneity, *Hydrol. Earth Syst. Sci.*, 6, 695–708, <https://doi.org/10.5194/hess-6-695-2002>, 2002.
- Pegram, G. G. S. and Clothier, A. N.: High resolution space-time modelling of rainfall: the “String of Beads” model, *J. Hydrol.*, 241, 26–41, 2001a.
- Pegram, G. G. S. and Clothier, A. N.: Downscaling rainfields in space and time, using the String of Beads model in time series mode, *Hydrol. Earth Syst. Sci.*, 5, 175–186, <https://doi.org/10.5194/hess-5-175-2001>, 2001b.
- Pegram, G., Lloret, X., and Sempere-Torres, D.: Radar rainfall: separating signal and noise fields to generate meaningful ensembles, *Atmos. Res.*, 100, 226–236, 2011.
- Pierce, C., Hirsch, T., and Bennett, A. C.: Formulation and evaluation of a post-processing algorithm for generating seamless, high resolution ensemble precipitation forecasts, *Forecasting R&D Technical Report 550*, MetOffice, Exeter, UK, 2010.
- Rebor, N., Ferraris, L., von Hardenberg, J., and Provenzale, A.: Rainfall downscaling and flood forecasting: a case study in the Mediterranean area, *Nat. Hazards Earth Syst. Sci.*, 6, 611–619, <https://doi.org/10.5194/nhess-6-611-2006>, 2006a.
- Rebor, N., Ferraris, L., von Hardenberg, J., and Provenzale, A.: RainFARM: Rainfall downscaling by a filtered autoregressive model, *J. Hydrometeorol.*, 7, 724–738, 2006b.
- Robin, M. J. L., Gutjahr, A. L., Sudicky, E. A., and Wilson, J. L.: Cross-correlated random field generation with the direct Fourier Transform Method, *Water Resour. Res.*, 29, 2385–2397, <https://doi.org/10.1029/93WR00386>, 1993.
- Schertzer, D. and Lovejoy, S.: Physical modelling and analysis of rain and clouds by anisotropic scaling multiplicative processes, *J. Geophys. Res.*, 92, 9696–9714, 1987.
- Schiemann, R., Erdin, R., Willi, M., Frei, C., Berenguer, M., and Sempere-Torres, D.: Geostatistical radar-rain gauge combination with nonparametric correlograms: methodological considerations and application in Switzerland, *Hydrol. Earth Syst. Sci.*, 15, 1515–1536, <https://doi.org/10.5194/hess-15-1515-2011>, 2011.
- Schleiss, M., Chamoun, S., and Berne, A.: Stochastic simulation of intermittent rainfall using the concept of “dry drift”, *Water Resour. Res.*, 50, 2329–2349, 2014.

- Seed, A. W., Srikanthan, R., and Menabde, M.: A space and time model for design storm rainfall, *J. Geophys. Res.*, 104, 31623–31630, <https://doi.org/10.1029/1999JD900767>, 1999.
- Seed, A. W.: A dynamic and spatial scaling approach to advection forecasting, *J. Appl. Meteorol.*, 42, 381–388, 2003.
- Seed, A. W., Pierce, C. E., and Norman, K.: Formulation and evaluation of a scale decomposition-based stochastic precipitation nowcast scheme, *Water Resour. Res.*, 49, 6624–6641, 2013.
- Sideris, I. V., Gabella, M., Erdin, R., and Germann, U.: Real-time radar–raingauge merging using spatio-temporal co-kriging with external drift in the alpine terrain of Switzerland, *Q. J. Roy. Meteor. Soc.*, 140, 1097–1111, 2014.
- Sinclair, S. and Pegram, G. G. S.: Combining radar and rain gauge rainfall estimates using conditional merging, *Atmos. Sci. Lett.*, 6, 19–22, 2005.
- Stankovic, L., Dakovic, M., and Thayananthan, T.: Time-frequency signal analysis with applications, 1st Edn., Artech House, Boston, 2013.
- Stankovic, S., Orovic, I., and Sejdic, E.: *Multimedia Signals and Systems*, Springer-Verlag, New York, 2012.
- Tompson, A. F. B., Ababou, R., and Gelhar, L. W.: Implementation of the three-dimensional turning bands random field generator, *Water Resour. Res.*, 25, 2227–2243, <https://doi.org/10.1029/WR025i010p02227>, 1989.
- Velasco-Forero, C. A., Sempere-Torres, D., Cassiraga, E. F., and Gómez-Hernández, J. J.: A non-parametric automatic blending methodology to estimate rainfall fields from rain gauge and radar data, *Adv. Water Resour.*, 32, 986–1002, 2009.
- Villarini, G., Krajewski, W. F., Ciach, G. J., and Zimmerman, D. L.: Product-error-driven generator of probable rainfall conditioned on WSR-88D precipitation estimates, *Water Resour. Res.*, 45, W01404, <https://doi.org/10.1029/2008WR006946>, 2009.
- Wiener, N.: Generalized harmonic analysis, *Acta Math.*, 55, 117–258, 1930.
- Zappa, M., Beven, K. J., Bruen, M., Cofiño, A. S., Kok, K., Martin, E., Nurmi, P., Orfila, O., Roulin, E., Schröter, K., Seed, A., Szturc, J., Vehviläinen, B., Germann, U., and Rossa, A.: Propagation of uncertainty from observing systems and NWP into hydrological models: COST-731 Working Group 2, *Atmos. Sci. Lett.*, 11, 83–91, 2010.