APPROXIMATE SIMULATION OF LINEAR CONTINUOUS TIME MODELS DRIVEN BY ASYMMETRIC STABLE LÉVY PROCESSES

Marina Riabiz and Simon J. Godsill

Signal Processing and Communications Laboratory, Cambridge University Engineering Department, Cambridge, CB2 1PZ, UK - {mr622, sjg30} @cam.ac.uk

ABSTRACT

In this paper we extend to the multidimensional case the modified Poisson series representation of linear stochastic processes driven by α-stable innovations. The latter has been recently introduced in the literature and it involves a Gaussian approximation of the residuals of the series, via the exact characterization of their moments. This allows for Bayesian techniques for parameter or state inference that would not be available otherwise, due to the lack of a closed-form likelihood function for the α-stable distribution. Simulation results are presented to validate the introduced extension and the quality of the approximation of the distribution. Finally, we show an example of generation from the process.

Index Terms— α-stable multidimensional Lévy processes, Poisson series representation, conditionally Gaussian distribution, residual approximation.

1. INTRODUCTION

Time series arising in the natural sciences, engineering and finance are frequently characterised by high data-rate and irregular sampling, a situation well represented by continuous-time state-space models. The most straightforward form of these, in terms of analytical tractability, is the linear model driven by Brownian motion, for which a wide literature has been developed, see e.g. [1, 2]. However, its simplicity is limiting in the kind of phenomena that it can account for, in particular heavy tails and jumps/ discontinuities. We refer, for instance, to [3] for a review of application areas, spanning from financial models to climatological data, to audio and image processing. A possible tool to deal with such data is the jump-diffusion model, see [4], that adds a jump process to the Gaussian diffusion. On the contrary, we consider here a single non-Gaussian driving noise that exhibit such behaviours, namely α-stable Lévy-processes driving linear state-space systems [5, 6].

The α-stable distribution [7] extends the Central Limit Theorem to heavy-tailed cases with non-finite variance. Its parameters allow extreme values (as well as skewness), while including the Gaussian distribution as a special case. The lack of closed-form density expressions makes direct inference intractable. An approximate conditionally Gaussian framework targeting skewed α-Stable distributions has been recently formulated by [8, 9], and it generalizes the scale mixture of normals approach given by [10, 11, 12] for the symmetric case. This approach truncates the Poisson series representation of α-stable random variables, see [6], and computes the parameters of Gaussian approximations of the residuals, using exact formulae for their moments. The advantage of this representation is to allow Bayesian inference via standard Bayesian computational tools such as Monte Carlo Expectation Maximization, Markov chain Monte Carlo [13] and particle filters [14]. A general approximate framework for the simulation of Lévy-processes is given in [15, 16, 17]. This is related to the approach in our paper in that it approximates small jumps with a Gaussian term, but considers pure Lévy processes, i.e. not the general linear state space models that we are concerned with here.

The main contribution of this paper is the extension and experimental validation of the existing framework to the case of multidimensional state-space models. In Section 2 we describe the continuous-time models driven by α-stable noise that we are targeting. Section 3 gives the theoretical expression for the multidimensional extension of the Poisson series representation of α-stable stochastic integrals, and Section 4 two Gaussian approximations of the residuals of the series. In Section 5 we provide experimental validation of the approximations and show a simulation from the process.

2. LINEAR CONTINUOUS TIME α-STABLE LÉVY PROCESSES

In the following we consider linear continuous-time processes of order P, $\mathbf{x}_t$, driven by a scalar noise l_t, that can be expressed in the differential form

$$d\mathbf{x}_t = \mathbf{A}\mathbf{x}_t dt + \mathbf{h} dl_t,$$

where $\mathbf{x}_t = [x_{1,t}, \dots, x_{P,t}]'$, $\mathbf{A}$ is a $P \times P$ matrix describing the interaction of the components of the process, and $\mathbf{h}$ is a P-dimensional vector indicating the direct effect of the noise

The first author acknowledges partial funding from EPSRC-DTG-2013/EU.

l_t on each state component $x_{i,t}$. Observe that, depending on the structure of $\mathbf{A}$ and $\mathbf{h}$, $x_{i,t}$ can denote the i-th derivative of $x_{1,t}$, for $i > 1$. The above stochastic differential equation has solution

$$\mathbf{x}_{t+\delta} = e^{\mathbf{A}\delta}\mathbf{x}_t + \int_0^\delta e^{\mathbf{A}(\delta-u)}\mathbf{h}dl_u, \tag{1}$$

see e.g. [18]. In order to represent skewed time series with extreme values, we model l_t as a non-symmetric α-stable Lévy process, meaning that $l_0 = 0$ almost surely and that l_t has independent α-stable increments $l_{t+\delta} - l_t \sim S_\alpha\left(\tilde{\sigma}\delta^{1/\alpha}, \beta, 0\right)$, $\forall\delta > 0$. We write $Y \sim S_\alpha(\sigma, \beta, \mu)$ to denote that the random variable Y has α-stable distribution in the S^1 parametrization, whose characteristic function is referenced in [19]. The parameter $\alpha \in (0, 2]$ represents the tail thickness, $\beta \in [-1, 1]$ the skewness, $\sigma > 0$ the scale and $\mu \in \mathbb{R}$ the location. The expression of $\tilde{\sigma}$ is determined according to the modified Poisson series representation (MPSR) of α-stable random variables given in [8, 9] and presented in the next Section. The latter is motivated by the fact that the α-stable characteristic function cannot be inverted to closed-form density expressions, except for a few special cases. Thus, approximations are needed to allow for inference.

3. POISSON SERIES REPRESENTATION

The Poisson series representation (PSR) of α-stable random variables was originally introduced by [20]. Here we refer to its version in [6], where a result is provided also for the series expansion of stochastic integrals with respect to α-stable random measures. However, while the PSR for random variables allows for a conditionally Gaussian representation of the distribution, the one for stochastic integrals does not. Thus a modified Poisson series representation (MPSR), that meets this requirement, was presented in [8] and proved in [9]. Our contribution is to generalize this formulation, given for stochastic integrals of real-valued functions, to vector-valued functions, namely

$$\boldsymbol{\xi} := \int_0^\delta \mathbf{f}(u)dl_u,$$

Observe that, in (1), $\mathbf{f}(u) = e^{\mathbf{A}(\delta-u)}\mathbf{h}$. Applying the MPSR to $\boldsymbol{\xi}$, we obtain the following equality in distribution $\stackrel{\mathcal{D}}{=}$

$$\boldsymbol{\xi} \stackrel{\mathcal{D}}{=} \delta^{1/\alpha}\sum_{i=1}^{\infty} W_i\Gamma_i^{-1/\alpha}\mathbf{f}(V_i) - b_i^{(\alpha)}\mathbf{k}, \tag{2}$$

where $\{\Gamma_i\}_{i=1}^\infty$ are the arrival times of a Poisson process with unit arrival rate, $\{V_i\}_{i=1}^\infty$ are i.i.d. uniform random variables in $[0, \delta]$ and $\{W_i\}_{i=1}^\infty$ are i.i.d. random variables independent of Γ_i and V_i, with conditions on the moments given in [9]. Denoting with $\mathbb{E}[\cdot]$ the expected value,

$$\mathbf{k} = \mathbb{E}[W_i]\mathbb{E}\left[\mathbf{f}(V_i)\right],$$

$$b_i^{(\alpha)} = \begin{cases} 0 & \text{if } 0 < \alpha < 1 \\ \dfrac{\alpha}{\alpha-1}\left(i^{\frac{\alpha-1}{\alpha}} - (i-1)^{\frac{\alpha-1}{\alpha}}\right) & \text{if } 1 < \alpha < 2. \end{cases}$$

For simplicity of notation, in the paper we do not deal with the case $\alpha = 1$ (a pole for the S^1 parametrization). Then, if $W_i \stackrel{i.i.d}{\sim} \mathcal{N}(\mu_W, \sigma_W^2)$, the MPSR implies that, conditionally on the full sequences of latent variables $\{\Gamma_i, V_i\}_{i=1}^\infty$, $\boldsymbol{\xi}$ has Gaussian distribution

$$\boldsymbol{\xi}\big|\,\{V_i, \Gamma_i\}_{i=1}^\infty \sim \mathcal{N}\left(\mu_W\mathbf{m}, \sigma_W^2\mathbf{S}\right), \tag{3}$$

with moments proportional to the following series

$$\mathbf{m} := \delta^{1/\alpha}\sum_{i=1}^{\infty}\Gamma_i^{-1/\alpha}\mathbf{f}_i - b_i^{(\alpha)}\mathbb{E}\left[\mathbf{f}_i\right], \tag{4}$$

$$\mathbf{S} := \delta^{2/\alpha}\sum_{i=1}^{\infty}\Gamma_i^{-2/\alpha}\mathbf{f}_i\mathbf{f}_i', \tag{5}$$

where $\mathbf{f}_i := \mathbf{f}(V_i)$ from now on. However, for simulation purposes, the full sequences of latent variables cannot be produced, meaning that the series (2)-(4)-(5) need to be truncated. In [9] two Gaussian approximations of the residuals of the series are proposed, in order to preserve the overall conditionally Gaussian structure of $\boldsymbol{\xi}$. We provide the multidimensional version of this procedure in the following section.

4. RESIDUAL APPROXIMATION

The heavy tailed behaviour of the series (2)-(4)-(5) is determined by the first terms, due to the fact that the sequence $\{\Gamma_i\}_{i=1}^\infty$ is monotonically increasing and $\alpha > 0$ (the lower α is, the heavier the tails of the distribution are). With this in mind, we assume that the first M terms of the series are given, where $M \sim$ Poisson(c) is a random number of terms such that $\Gamma_M \leq c$ and $\Gamma_{M+1} > c$. The residuals of the series are not Gaussian, if the latent variables successive to the Mth are unknown. Here we compute the moments of the residuals and discuss the quality of the approximation of their distribution in Section 5.

4.1. Moments of the residual series

Following the procedure proposed in [9], in this first method we compute the moments of the residual of (2) as follows

$$\boldsymbol{\xi}\big|\,\{V_i, \Gamma_i\}_{i=1}^\infty \stackrel{\mathcal{D}}{=} \delta^{1/\alpha}\left(\sum_{i=1}^{M} W_i\Gamma_i^{-1/\alpha}\mathbf{f}_i + \mathbf{r}\right),$$

with

$$\mathbf{r} := \lim_{d\to\infty}\mathbf{r}_d,$$

$$\mathbf{r}_d := \sum_{i:\Gamma_i\in[c,d]} W_i\Gamma_i^{-1/\alpha}\mathbf{f}_i - \sum_{n:\Gamma_n\in[0,d]} b_n^{(\alpha)}\mathbf{k}.$$

Then, denoting with $\mathrm{Var}[\cdot]$ the variance-covariance matrix,

$$\mathbb{E}\left[\mathbf{r}\right] = \lim_{d\to\infty} \mathbb{E}\left[\mathbf{r}_d\right], \quad \mathrm{Var}\left[\mathbf{r}\right] = \lim_{d\to\infty} \mathrm{Var}\left[\mathbf{r}_d\right].$$

The following result can be proved.

Lemma 4.1 *The moments of the residual of the series* (2) *are*

$$\mathbb{E}\left[\mathbf{r}\right] = \mu_W \frac{\alpha}{1-\alpha} c^{\frac{\alpha-1}{\alpha}} \mathbf{q}(\delta),$$
$$\mathrm{Var}\left[\mathbf{r}\right] = (\sigma_W^2 + \mu_W^2) \frac{\alpha}{2-\alpha} c^{\frac{\alpha-2}{\alpha}} \mathbf{Q}(\delta),$$

where $\mathbf{q}(\delta) \coloneqq \mathbb{E}\left[\mathbf{f}_i\right],\ \mathbf{Q}(\delta) \coloneqq \mathbb{E}\left[\mathbf{f}_i\mathbf{f}_i'\right].$

Observe that, to perform the computations, we need to evaluate $\mathbf{q}(\delta)$ and $\mathbf{Q}(\delta)$. In the stochastic process (1), these are integrals of matrix exponentials on a bounded support. Given that the integral of a vector or matrix-valued function is its componentwise integral, this can be done either numerically or according to the procedure proposed in [21]. The latter involves computing the exponential of an auxiliary matrix and combining its blocks. This method is not exact either, because matrix exponentials are in turn defined as series, and numerical methods may be required to evaluate them, see [22].

4.2. Moments of the residual series mean and covariance

Here we characterize the moments of the residuals in the series (4) and (5), modelling their joint distribution as multivariate Gaussian. The main motivation behind this approach is that it allows to account for the correlation existing between (4) and (5), caused by the same random variables being involved. It also preserves the structure of (3), including only the parameter σ_W^2 in the variance-covariance matrix of the approximation of $\boldsymbol{\xi}$. This proves to be useful for Bayesian inference methods, as shown in [9]. In summary

$$\mathbf{m} = \delta^{1/\alpha}\left[\sum_{i=1}^{M} \Gamma_i^{-1/\alpha}\mathbf{f}_i + \mathbf{r}^m\right],$$
$$\mathbf{S} = \delta^{2/\alpha}\left[\sum_{i=1}^{M} \Gamma_i^{-2/\alpha}\mathbf{f}_i\mathbf{f}_i' + \mathbf{R}^S\right],$$

with

$$\mathbf{r}^m \coloneqq \lim_{d\to\infty} \mathbf{r}_d^m, \qquad \mathbf{R}^S \coloneqq \lim_{d\to\infty} \mathbf{R}_d^S,$$

and $\mathbf{r}_d^m$ a P-dimensional vector and $\mathbf{R}_d^S$ a $P\times P$ symmetric positive definite matrix, respectively given by

$$\mathbf{r}_d^m \coloneqq \sum_{i:\Gamma_i\in[c,d]} \Gamma_i^{-1/\alpha}\mathbf{f}_i - \sum_{n:\Gamma_n\in[0,d]} b_n^{(\alpha)}\mathbf{k},$$
$$\mathbf{R}_d^S \coloneqq \sum_{i:\Gamma_i\in[c,d]} \Gamma_i^{-2/\alpha}\mathbf{f}_i\mathbf{f}_i'.$$

In the computation, we vectorize the lower diagonal part of $\mathbf{R}^S$ and $\mathbf{R}_d^S$ to the $(P^2+P)/2$-dimensional vectors $\mathbf{r}^S \coloneqq \mathrm{vech}(\mathbf{R}^S)$ and $\mathbf{r}_d^S \coloneqq \mathrm{vech}(\mathbf{R}_d^S)$, where $\mathrm{vech}(\cdot)$ indicates the the half-vectorization of a symmetric matrix. Defining $\mathbf{r}^{\mathrm{tot}} \coloneqq \left[(\mathbf{r}^m)', (\mathbf{r}^S)'\right]'$, we have that

$$\mathbb{E}\left[\mathbf{r}^{\mathrm{tot}}\right] = \lim_{d\to\infty}\begin{bmatrix}\mathbb{E}[\mathbf{r}_d^m]\\ \mathbb{E}[\mathbf{r}_d^S]\end{bmatrix},$$
$$\mathrm{Var}\left[\mathbf{r}^{\mathrm{tot}}\right] = \lim_{d\to\infty}\begin{bmatrix}\mathrm{Var}[\mathbf{r}_d^m] & \mathrm{Cov}[\mathbf{r}_d^m, \mathbf{r}_d^S]\\ \mathrm{Cov}[\mathbf{r}_d^S, \mathbf{r}_d^m] & \mathrm{Var}[\mathbf{r}_d^S]\end{bmatrix},$$

where $\mathrm{Cov}[\cdot,\cdot]$ denotes the cross-covariance matrix. Consequently, the following result is obtained.

Lemma 4.2 *The moments of the residuals of the series* (4) *and* (5) *are*

$$\mathbb{E}\left[\mathbf{r}^{tot}\right] = \begin{bmatrix}\frac{\alpha}{1-\alpha}c^{\frac{\alpha-1}{\alpha}}\mathbf{q}(\delta)\\ \frac{\alpha}{2-\alpha}c^{\frac{\alpha-2}{\alpha}}\mathit{vech}\left(\mathbf{Q}(\delta)\right)\end{bmatrix},$$
$$\mathrm{Var}\left[\mathbf{r}^{tot}\right] = \begin{bmatrix}\frac{\alpha}{2-\alpha}c^{\frac{\alpha-2}{\alpha}}\mathbf{Q}(\delta) & \frac{\alpha}{3-\alpha}c^{\frac{\alpha-3}{\alpha}}\mathbb{E}\left[\mathbf{f}_i\mathbf{g}_i'\right]\\ \frac{\alpha}{3-\alpha}c^{\frac{\alpha-3}{\alpha}}\mathbb{E}\left[\mathbf{g}_i\mathbf{f}_i'\right] & \frac{\alpha}{4-\alpha}c^{\frac{\alpha-4}{\alpha}}\mathbb{E}\left[\mathbf{g}_i\mathbf{g}_i'\right]\end{bmatrix},$$

where $\mathbf{g}_i \coloneqq \mathit{vech}\left(\mathbf{f}_i\mathbf{f}_i'\right).$

We observe that the moments of $\mathbf{r}_d, \mathbf{r}_d^m, \mathbf{r}_d^S$ have a higher convergence rate to those of $\mathbf{r}, \mathbf{r}^m, \mathbf{r}^S$ the lower is α. Furthermore the mean values coincide if $\alpha > 1$.

Now that we have computed the moments of the residual terms, we can approximate them as normally distributed

$$\mathbf{r} \overset{\text{approx}}{\sim} \mathcal{N}\left(\mathbb{E}[\mathbf{r}], \mathrm{Var}[\mathbf{r}]\right),$$
$$\mathbf{r}^{\mathrm{tot}} \overset{\text{approx}}{\sim} \mathcal{N}\left(\mathbb{E}[\mathbf{r}^{\mathrm{tot}}], \mathrm{Var}[\mathbf{r}^{\mathrm{tot}}]\right),$$

respectively. However, we accentuate that this is an approximation, given that we do not know the latent variables successive to the Mth. For conciseness, in the following we experimentally validate the approximation on $\mathbf{r}$, but analogous considerations hold for $\mathbf{r}^{\mathrm{tot}}$. We show that approximating the distribution of the residual as Gaussian is accurate if a sufficiently large series truncation limit, $c = c(\alpha)$, is chosen.

5. RESULTS

In particular, in this section we consider $\mathbf{h} = [0,\ldots,0,1]'$ and two scenarios for the transition matrix $\mathbf{A}$:

$$\mathbf{A}_1 = \left[\begin{array}{c|c}\mathbf{0}_{P-1} & \mathbf{I}_{P-1,P-1}\\ \hline -a_P & -a_{P-1}\ldots - a_1\end{array}\right], \quad \mathbf{A}_2 = \begin{bmatrix}\theta_1 & 1\\ 0 & \theta_2\end{bmatrix},$$

where $\mathbf{0}_{P-1}$ is a $(P-1)$-dimensional null vector, $\mathbf{I}_{P-1,P-1}$ is the $(P-1)\times(P-1)$ identity matrix, and stationarity is achieved if the parameters $\{a_i\}_{i=1}^P \in \mathbb{R}^+$, and $\{\theta_i\}_{i=1}^2$ have negative real part. The matrix $\mathbf{A}_1$ represents a continuous-time autoregressive model of order P, CAR(P), where $x_{1,t}$ is the state and $x_{i,t}, i > 1$ are its derivatives, see [23]. $\mathbf{A}_2$ corresponds to a model with components reverting to their mean value, as in [24, 25].

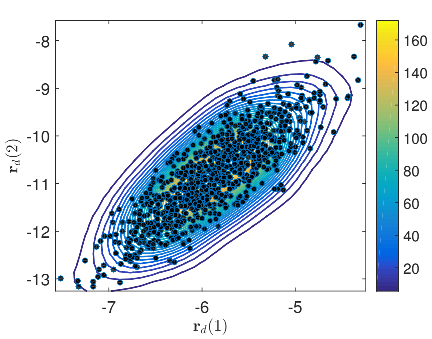

Fig. 1. Contour plot of the sample distribution of $\mathbf{r}_d$.

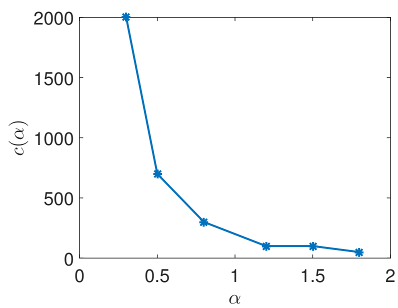

Fig. 2. Values of c that enable multivariate Normality of $\mathbf{r}_d$, as a function of α, for the considered CAR(2) process.

5.1. Validation of the Gaussian approximation

As anticipated, the quality of the Gaussian approximation depends on the series truncation limit c. In terms of the computational cost of the simulations from the process, we are interested in keeping c low. In fact, c is the average number of latent variables that we need to know exactly for the first part of the series in either approximation approach. On the other hand, the cost of generating the Gaussian residuals is independent of c, while this operation corrects the bias on the moments of the mere truncation, for any value of c. However, we cannot arbitrarily decrease this threshold, because this would degrade the Gaussian approximation, since the first terms assign the heavy-tailed behaviour to the series.

To test the multivariate Gaussianity of $\mathbf{r}_d$, we generate 10^3 samples from the CAR(2) model, with eigenvalues $\{-0.2, -0.3\}$, corresponding to $a_2 = 0.06$ and $a_1 = 0.5$. Furthermore we choose $\mu_W = 1, \sigma_W = 1, \delta = 1$. We perform the Royston's test [26] and we examine the contours of the sample histogram of the bivariate data, as well as at the marginal distributions of the components of $\mathbf{r}_d$ (the latter being only a necessary, but not sufficient condition). An example of contour plot of the sample distribution of $\mathbf{r}_d$ for $\alpha = 1.5$, $c = 100, d = 10^4$ can be seen in Figure 1. Figure 2 shows how c varies as a function of α in order for the Royston's test to be satisfied (over a trial set of c values).

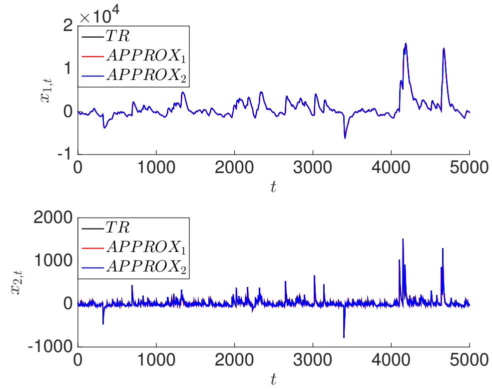

Fig. 3. Simulated process from $\mathbf{A}_2$.

5.2. Simulations from the process

Here we show simulations from the process (1), comparing the truncated series, and the two proposed residual approximations for a model that could represent the behaviour of a financial time series, as in [24, 25]. In Figure 3, we use the transition matrix $\mathbf{A}_2$, with $\theta_1 = -0.025, \theta_2 = -0.09, \alpha = 1.2, c = 100, \mu_W = 1, \sigma_W = 1$, and we simulate for 5000 irregularly sampled time steps. The process is initialised at $\mathbf{x}_0 = \mathbf{0}$, and it presents extreme values, while reverting to its zero mean value in both the components, as expected. The three realizations in Figure 3 (truncated series and truncation with added residuals, as in Lemma 4.1 and 4.2) appear very similar. However, analyses performed in [9] for the case $P = 1$, where an exact benchmark is available, demonstrate that the mere truncation introduces bias in the moments, significant especially for values of α close to 1 and 2. Simulations form the multivariate process show an analogous behaviour, indicating that residual approximations are necessary.

Other simulations not reported here illustrate that it is possible to represent different situations by changing the transition matrix and increasing P. For example, marginally stable systems can be obtained by setting one of the eigenvalues to zero, while complex eigenvalues produce oscillatory states[1] .

6. CONCLUSION

In this paper the expressions for a Gaussian approximation of the conditionally Gaussian structure of the MPSR of stochastic integrals of a vector valued function driven by α-stable Lévy noise are provided. The quality of the given approximation is also validated, including simulation from the process. The results obtained encourage future research into the formulation of Bayesian state and parameter inference techniques.

[1] A Matlab source code for the showed methods is available on `http://www-sigproc.eng.cam.ac.uk/Main/MR622`.

7. REFERENCES

[1] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, Springer, 2003.

[2] A. C. Harvey, *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, 1990.

[3] J. Nolan, "Stablebibliography," `http://fs2.american.edu/jpnolan/www/stable/stable.html`, 2016 (accessed January 9, 2017).

[4] S. G. Kou, "A jump-diffusion model for option pricing," *Manage. Sci.*, vol. 48, no. 8, pp. 1086–1101, 2002.

[5] A. Janicki and A. Weron, "Can one see α-Stable variables and processes?," *Statistical Science*, vol. 9, no. 1, pp. 109–126, 1994.

[6] G. Samoradnitsky and S. Taqqu, *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*, Stochastic Modeling Series. Taylor & Francis, 1994.

[7] P. Lévy, *Calcul des Probabilités*, PCMI collection. Gauthier-Villars, 1925.

[8] T. Lemke and S.J. Godsill, "Enhanced Poisson sum representation for alpha-stable processes," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, May 2011, pp. 4100–4103.

[9] T. Lemke, *Poisson series Approaches to Bayesian Monte Carlo Inference for Skewed α-Stable Distributions and Stochastic Processes*, Ph.D. thesis, Kaiserslautern University of Technology, 2014.

[10] S.J. Godsill and E.E. Kuruoğlu, "Bayesian inference for time series with heavy-tailed symmetric α-stable noise processes," in *Applications of Heavy Tailed Distributions in Economics, Engineering and Statistics*, Washington DC, USA, June 1999.

[11] S. J. Godsill, "Inference in symmetric alpha-stable noise using MCMC and the slice sampler," in *Acoustics, Speech, and Signal Processing (ICASSP), 2000 IEEE International Conference on*, June 2000, pp. 3806–3809.

[12] E. G. Tsionas, "Monte Carlo inference in econometric models with symmetric Stable disturbances ," *Journal of Econometrics*, vol. 88, no. 2, pp. 365 – 401, 1999.

[13] C. Robert and G. Casella, *Monte Carlo Statistical Methods*, Springer Texts in Statistics. Springer, 2004.

[14] A. Doucet, A. Smith, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, Information Science and Statistics. Springer New York, 2001.

[15] S. Asmussen and J. Rosinski, "Approximations of small jumps of Lévy processes with a view towards simulation," *Journal of Applied Probability*, vol. 38, no. 2, pp. 482–493, 2001.

[16] J. Rosinski, "Series representations of Lévy processes from the perspective of point processes," in *Lévy Processes*, O.E. Barndorff-Nielsen, S. Resnick, and T. Mikosch, Eds., pp. 401–415. Birkhauser Boston, 2001.

[17] J. Rosinski, *Simulation of Lévy Processes*, John Wiley & Sons, Ltd, 2014.

[18] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, vol. 23, Springer, 1992.

[19] J. P. Nolan, "Numerical calculation of Stable densities and distribution functions," *Comm. Statist. Stochastic Models*, vol. 13, no. 4, pp. 759–774, 1997.

[20] R. LePage, M. Woodroofe, and J. Zinn, "Convergence to a Stable distribution via order statistics," *The Annals of Probability*, vol. 9, no. 4, pp. 624–632, 1981.

[21] C. Van Loan, "Computing integrals involving the matrix exponential," Tech. Rep., Ithaca, NY, USA, 1976.

[22] C. Moler and C. Van Loan, "Nineteen dubious ways to compute the exponential of a matrix," *SIAM Review*, vol. 20, pp. 801–836, 1978.

[23] P.J. Brockwell, "Lévy-driven CARMA processes," *Annals of the Institute of Statistical Mathematics*, vol. 53, no. 1, pp. 113–124, 2001.

[24] H.L. Christensen, J. Murphy, and S.J. Godsill, "Forecasting high-frequency futures returns using online Langevin dynamics," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 6, no. 4, pp. 366–380, 2012.

[25] J. Murphy, *Hidden states, hidden structures: Bayesian learning in time series models*, Ph.D. thesis, University of Cambridge, Engineering Department, 2013.

[26] J. P. Royston, "Some techniques for assessing multivarate Normality based on the Shapiro- Wilk W," *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 32, no. 2, pp. 121–133, 1983.