

Simulation of Lévy Random Fields

Robert L. Wolpert
Katja Ickstadt

ABSTRACT

An efficient recently developed method, the Inverse Lévy Measure (ILM) algorithm, is presented for drawing random samples from gamma, skewed stable and other nonnegative independent-increment random fields, which we call Lévy random fields. The method is useful for computing posterior distributions in nonparametric hierarchical Bayesian statistical analysis. Algorithms are illustrated through prototype implementations in S-PLUS.

1 Introduction and Overview

Following the emergence of Gibbs sampling and more generally Markov chain Monte Carlo methods (Gelfand and Smith, 1990; Tierney, 1994; Gilks et al., 1996), Bayesian inference is now commonly based on simulations from the posterior distributions of model parameters. In nonparametric Bayesian analysis the “parameter” is typically a function or distribution so the simulation-based approach entails drawing random samples of functions (stochastic processes) or measures (random fields). Stationary independent-increment (i.e. homogeneous) random fields can serve as building blocks for Bayesian nonparametric priors: the Dirichlet process is simply the gamma random field normalized to have unit mass, for example, while the distributions called “neutral to the right” (Ferguson, 1973, 1974; Ferguson and Phadia, 1979) are simply those whose hazard functions have independent increments.

The extreme ease with which such samples may be drawn from the Dirichlet process distribution (Antoniak, 1974; Ferguson, 1973) is largely responsible for its frequent use in early Bayesian nonparametric analysis, despite its limitations (negative correlation for probabilities assigned to disjoint sets, absence of ways to express spatial correlation or varying precision, etc.).

Here we present a new and very efficient scheme, the Inverse Lévy Measure (ILM) algorithm based on an idea of (Lévy, 1937), for sampling from a wide class of random measures assigning positive, independent, infinitely divisible random variables to disjoint sets. The method is useful in Bayesian

hierarchical models, is sufficiently flexible to allow expression of a wide range of possible prior beliefs, and is nearly as simple to implement as are the Dirichlet process methods—our prototype implementations require only five lines of S-PLUS code (MathSoft, 1997). The new method’s flexibility allows a modeler to construct and make inference from hierarchical Bayesian nonparametric models without making unnecessary assumptions or approximations.

The underlying idea is presented in section 2. In section 3 it is illustrated for gamma random variables, processes, and fields; in section 4 it is presented in full generality for possibly inhomogeneous Lévy processes and fields. The method is compared with others in section 5 and discussed in section 6. Two appendices includes some S-PLUS display functions and the proof of Theorem 2.

2 Increasing Independent-Increment Processes: A New Look at an Old Idea

2.1 *Infinitely-Divisible Random Variables*

The celebrated Lévy-Khinchine formula (Jacod and Shiryaev, 1987, p. 75) asserts that every infinitely-divisible distribution $\mu(dx)$ has a characteristic function $\hat{\mu}(t) \equiv \int_{\mathbb{R}} e^{itx} \mu(dx)$ of the form

$$\hat{\mu}(t) = \exp \left\{ ita - t^2 \sigma^2 / 2 + \int_{\mathbb{R}} [e^{itu} - 1 - it h(u)] \nu(du) \right\}$$

for some numbers $a \in \mathbb{R}$, $\sigma^2 \in \mathbb{R}_+ \equiv [0, \infty)$, and some positive measure $\nu(du)$ (called the “Lévy measure”) on the real line satisfying the integrability condition $\int_{\mathbb{R}} (1 \wedge u^2) \nu(du) < \infty$, where the notation $1 \wedge u^2$ denotes the minimum of 1 and u^2 and $h(u)$ denotes an arbitrary bounded continuous function satisfying $h(u) = u + \mathcal{O}(u^2)$ near $u = 0$. For distributions of positive random variables, necessarily $a \geq 0$, $\sigma^2 = 0$, $\nu(\mathbb{R}_-) = 0$ and $\int_{\mathbb{R}_+} (1 \wedge u) \nu(du) < \infty$, so we may take $h(u) \equiv 0$ (absorbing its effect into a) and simplify this to the form

$$\hat{\mu}(t) = \exp \left[ita + \int_0^\infty (e^{itu} - 1) \nu(du) \right].$$

For example, the log characteristic functions of the Poisson (with mean λ), gamma (with shape α and scale β) and fully skewed stable (with exponent ξ) distributions are given by

$$\begin{aligned} \log \hat{\mu}_P(t) &= (e^{it} - 1)\lambda &= \int (e^{itu} - 1) \lambda \delta_1(du) \\ \log \hat{\mu}_\Gamma(t) &= -\log(1 - i\beta t)\alpha &= \int (e^{itu} - 1) u^{-1} e^{-u/\beta} \alpha du \\ \log \hat{\mu}_S(t) &= -\alpha |t|^\xi \Gamma(1 - \xi) e^{-it\xi\pi/2|t|} &= \int (e^{itu} - 1) u^{-1-\xi} \xi \alpha du \end{aligned}$$

with Lévy measures on \mathbb{R}_+ given by $\nu_P(du) = \lambda\delta_1(du)$ (where $\delta_x(du)$ denotes the unit point mass at locus $u = x$), $\nu_\Gamma(du) = u^{-1}e^{-u/\beta}\alpha du$ and $\nu_S(du) = u^{-1-\xi}\xi\alpha du$, respectively.

2.2 Increasing Independent-Increment Processes

Each nonnegative infinitely-divisible distribution may be thought of as the distribution at time $t = 1$ of an increasing stochastic process X_t with stationary independent increments, or a *subordinator* (Rogers and Williams, 1994, p. 78); in this context the Lévy measure of a Borel set $E \subset \mathbb{R}_+$ may be interpreted as the Poisson arrival rate $\nu(E)$ of jumps for the process of sizes $[X_t - X_{t-}] \in E$. For example, the Poisson process has jumps only of size $u = 1$, arriving at constant rate λ , while for each $\epsilon > 0$ the gamma process has jumps of size $u \geq \epsilon$ arriving as a Poisson process with rate $\nu_\Gamma([\epsilon, \infty)) = \int_\epsilon^\infty u^{-1}e^{-u/\beta}\alpha du$ and the one-sided stable process has jumps of size $u \geq \epsilon$ arriving as a Poisson process with rate $\nu_S([\epsilon, \infty)) = \alpha\epsilon^{-\xi}$. Other independent-increment processes have similar representations: the beta process (Hjort, 1990), with $\nu_\beta(du) = u^{-1}(1-u)^{\beta-1}\beta\alpha du$ on $[0, 1]$, for example, or the simple homogeneous process (Ferguson and Phadia, 1979), with $\nu(du) = e^{-\beta u}(1-e^{-u})^{-1}\alpha du$ on \mathbb{R}_+ .

2.3 The Key Idea

Implicit in the Lévy-Khinchine formula (and in Lévy (1937, ch. VII)) is the representation of a subordinator X_t as a stochastic integral $X_t = \iint_{\mathbb{R}_+ \times [0, t]} uH(du ds)$ of a Poisson measure on \mathbb{R}_+^2 with mean measure $E[H(du ds)] = \nu(du) ds$. This perspective suggests an exact method for simulating a subordinator X_t on the time interval $0 \leq s \leq T$: generate independent jump times σ_m from the uniform distribution on $[0, T]$; generate successive jumps τ_m of a standard unit-rate Poisson process (perhaps as partial sums of independent standard exponential variates) and set $v_m \equiv \inf\{u \geq 0 : T\nu([u, \infty)) \leq \tau_m\}$; then set $X_t \equiv \sum[v_m : \sigma_m \leq t]$. The sum will have finitely many terms if $\nu(\mathbb{R}_+) < \infty$, but will have infinitely many if $\nu(\mathbb{R}_+) = \infty$, as in the stable and gamma examples above. If we set $L(u) \equiv \nu([u, \infty))$ we can write $v_m = L^{-1}(\tau_m/T)$, justifying the name Inverse Lévy Measure (ILM) for this algorithm.

3 Example: Gamma Variates, Processes and Fields

For the gamma process the jump sizes $v_m \equiv \inf\{u \geq 0 : T\nu([u, \infty)) \leq \tau_m\}$ are found by solving for $u = x$ the equation $\tau/T = \nu([x, \infty)) = \int_x^\infty u^{-1}e^{-u/\beta}\alpha du = \alpha E_1(x/\beta)$, where $E_1(x) \equiv \int_x^\infty e^{-u} u^{-1} du$ is the exponential integral function (Abramowitz and Stegun, 1964, p. 228). This func-

tion, included in many special function subroutine libraries (IMSL, NSWG, etc.) and mathematical environments (Mathematica, MatLab, etc.), is the limit as $a \rightarrow 0$ of the complementary incomplete gamma function

$$E_1(x) = \lim_{a \rightarrow 0} \left[\Gamma(a, x) \equiv \int_x^\infty e^{-t} t^{a-1} dt \right]$$

(Abramowitz and Stegun, 1964, p. 260) and so can also be computed as the limit of rescaled χ^2 -tail probabilities as the degrees of freedom $d = 2a$ tend to zero,

$$E_1(x) = \lim_{d \rightarrow 0} \Gamma(d/2) \Pr[\chi_d^2 > 2x] = \lim_{d \rightarrow 0} (2/d) \Pr[\chi_d^2 > 2x].$$

Thus $E_1(x)$ and its inverse $E_1^{-1}(y)$ can be approximated in S-PLUS by the functions

```
"e1" <- function(x, d = 1e-9) {(2/d) * (1-pchisq(2*x, d))}
"e1inv" <- function(y, d = 1e-9) {qchisq(1-(d/2)*y, d)/2}
```

3.1 Gamma Random Variates

The first M jumps of a standard Poisson process can also be simulated easily in S-PLUS, as the partial sum `cumsum(rexp(M))`, so a gamma random variate with shape and scale parameters a and b can be constructed approximately using the ILM algorithm with the function

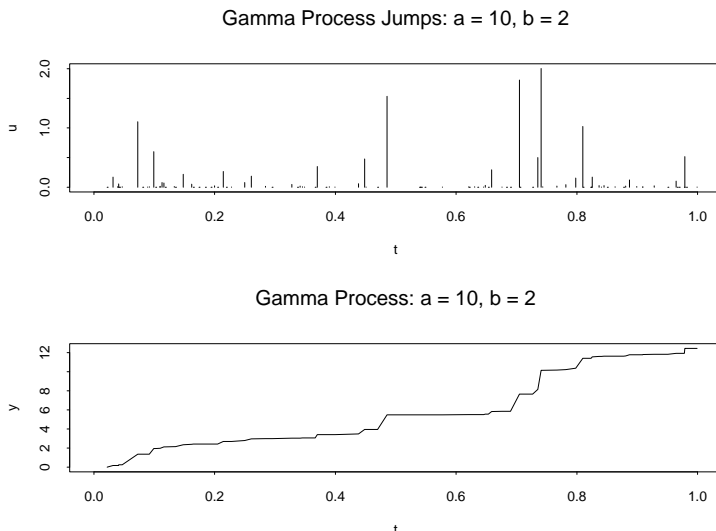
```
"gv" <- function(a, b, M = 1000) {
  return(sum(b * e1inv(cumsum(rexp(M)/a)))); }
```

Of course gamma random variates may be generated directly in S-PLUS with the command `b*rgamma(1,a)`; the strength of the present method lies in its extensibility to arbitrary infinitely-divisible distributions and to independent increment random processes and fields.

3.2 Gamma Random Processes

The ILM method may also be used for generating a gamma process (Ferguson and Klass, 1972); now we must additionally generate the jump times uniformly on the interval $[0, T]$, and sort the jumps in increasing temporal order. Figure 1, generated by the command `disp.1d(gp(10,2,M=250), "Gamma", c("a","b"), c(10,2))`; illustrates this with plots of both the gamma process and its jumps, using the display function `disp.1d()` given in Appendix 1 and the following function implementing the ILM algorithm for a gamma process:

```
"gp" <- function(a, b, tlim=1, M = 1000) {
  x <- runif(M)*tlim; # times x
```

FIGURE 1. Gamma Process realization with $\alpha = 10$, $\beta = 2$.


```
u <- b * elinv(cumsum(rexp(M)/(tlim*a))); # jumps u
ox<- order(x); x <- x[ox]; u <- u[ox]; # sort x,u on x
invisible(list(M=M,x=x,u=u,y=cumsum(u),t=tlim)); }
```

3.3 Gamma Random Fields

A gamma random field is a continuous (in probability) linear assignment of random variables $X[\phi]$ to bounded Borel functions $\phi(\cdot)$ on some set \mathcal{S} in Euclidean space such that indicator functions $\phi = 1_A$ will be assigned gamma-distributed random variables $X[1_A] \sim \text{Ga}(\alpha|A|, \beta)$ with shape parameter proportional to volume. With the ILM approach we can generate gamma random fields in any number of dimensions. Figure 2, produced with the command `disp.2d(grf(a=10,b=2), "Gamma", c("a","b"), c(10,2))`, shows representations of the gamma random field and its jumps on the unit square \mathcal{S} in two dimensions:

```
"grf"<- function(a, b, M=10000, g=51) {
  x <- runif(M); y <- runif(M); # jump locations
  u <- b * elinv(cumsum(rexp(M)/a)); # jump sizes
  xx <- seq(0,1,length=g); yy <- seq(0,1,length=g);
  zz <- matrix(0,nrow=g,ncol=g); # g*g grid
  for(i in 2:g) {
    xok <- x < xx[i];
    for(j in 2:g) {
```

Gamma Field Jumps:
a = 10, b = 2



Gamma Random Field:
a = 10, b = 2




FIGURE 2. Gamma Random Field realization with $\alpha = 10$, $\beta = 2$.

```

yok <- y < yy[j];
zz[i,j] <- sum(u[xok&yok]);      # sum of jumps SW
}                                # of (x[i],y[j])
}
invisible(list(x=x,y=y,u=u, xx=xx,yy=yy,zz=zz)); }

```

The ILM simulation method is applicable to more general gamma random fields: the shape measure may be any measure $\alpha(ds)$, not necessarily a constant times Lebesgue measure; the scale parameter $\beta(s)$ need not be constant; and the jump locations need not be drawn from a uniform distribution, but may be drawn from any convenient sampling distribution $\Pi(ds)$; see (Wolpert and Ickstadt, 1997) for the proof of:

Theorem 1 *Let $\alpha(s) \geq 0$ and $\beta(s) > 0$ be measurable functions on a space \mathcal{S} . Let $\{\sigma_m\}$ be independent identically distributed draws from any probability distribution $\Pi(ds)$ on \mathcal{S} , and let $\tau_m \geq 0$ be the successive jump times of a standard Poisson process. Set $\tau(u, s) \equiv E_1(u/\beta(s))\alpha(s)$ and $v_m \equiv \inf\{u \geq 0 : \tau(u, \sigma_m) \leq \tau_m\}$, that is,*

$$v_m \equiv E_1^{-1}(\tau_m/\alpha(\sigma_m)) \beta(\sigma_m),$$

or $v_m \equiv 0$ if $\alpha(\sigma_m) = 0$. Then the random field $\Gamma[\phi] \equiv \sum_{m < \infty} v_m \phi(\sigma_m)$ has the gamma process distribution $\Gamma(ds) \sim \text{Ga}(\alpha(ds), \beta(s))$ for the measure $\alpha(ds) \equiv \alpha(s)\Pi(ds)$.

4 Inhomogeneous Lévy Random Fields

Independent increment processes and fields that are *not* homogeneous nevertheless have a representation similar to that of Lévy and Khinchine. From (Jacod and Shiryaev, 1987, Chapter II, §4c) it follows that a nonnegative independent-increment random field must have a characteristic functional of the form

$$\hat{\mu}[\phi] = \mathbb{E} \left\{ e^{iX[\phi]} \right\} = \exp \left\{ i\eta[\phi] + \iint_{\mathbb{R}_+ \times \mathcal{S}} [e^{iu\phi(s)} - 1] \nu(du ds) \right\}, \quad (1.1)$$

where now $\eta(ds)$ is a finite positive measure on \mathcal{S} , $\eta[\phi] \equiv \int_{\mathcal{S}} \phi(s) \eta(ds)$, and the inhomogeneous Lévy measure $\nu(du ds)$ on $\mathbb{R}_+ \times \mathcal{S}$ satisfies the integrability condition $\iint_{\mathbb{R}_+ \times \mathcal{S}} (1 \wedge u) \nu(du ds) < \infty$. The inhomogeneous gamma random field of Theorem 1 is of this form with $\eta(ds) = 0$ and

$$\nu(du ds) = \exp(-u/\beta(s)) u^{-1} \alpha(s) du \Pi(ds),$$

and inhomogeneous versions of the Poisson, stable, and other processes may be generated in the same way. The heart of the idea is the same as before: represent the Lévy field in the form $X[\phi] = \eta[\phi] + \iint_{\mathbb{R}_+ \times \mathcal{S}} \phi(s) u H(du ds)$, where now the Poisson measure $H(du ds)$ on $\mathbb{R}_+ \times \mathcal{S}$ has mean $\mathbb{E}[H(du ds)] = \nu(du ds)$, leading to the following algorithm (in which we use the same letter to denote both the Lévy measure $\nu(du ds)$ and its density $\nu(u, s)$ with respect to the measure $du \Pi(ds)$):

Algorithm ILM:

- Choose any probability measure $\Pi(ds)$ on a Borel set \mathcal{S} from which samples may be drawn, any finite drift measure $\eta(ds)$ on \mathcal{S} and any nonnegative Borel function $\nu(u, s)$ on $\mathbb{R}_+ \times \mathcal{S}$ that satisfies the condition $\iint_{\mathbb{R}_+ \times \mathcal{S}} (1 \wedge u) \nu(u, s) du \Pi(ds) < \infty$.
- Set $\nu(du ds) \equiv \nu(u, s) du \Pi(ds)$ and $\tau(u, s) \equiv \int_u^\infty \nu(x, s) dx$.
- Generate independent draws σ_m from $\Pi(ds)$ and construct the event times of a standard Poisson process by generating independent unit-mean exponential random variables T_i and setting $\tau_m \equiv \sum_{i \leq m} T_i$.
- Set $v_m \equiv \inf[u \geq 0 : \tau(u, \sigma_m) \leq \tau_m]$ (typically by setting v_m to the inverse Lévy measure $\tau(\cdot, \sigma_m)^{-1}(\tau_m)$).
- Set $X[\phi] = \eta[\phi] + \sum_m v_m \phi(\sigma_m)$ for bounded Borel $\phi(s)$ on \mathcal{S} .

Of course in practice only a finite number M of jumps (v_m, σ_m) can be sampled, but the convergence in distribution of the truncated version $X_M[\phi]$ as $M \rightarrow \infty$ is an immediate consequence of the following theorem, whose proof appears in Appendix 2; truncation error bounds and suggestions for choosing M appear in (Wolpert and Ickstadt, 1997).

Theorem 2 *Let $\Pi(ds)$ be any Borel probability measure on \mathcal{S} and $\nu(u, s)$ a nonnegative Borel function on $\mathbb{R}_+ \times \mathcal{S}$ for which the measure $\nu(du ds) = \nu(u, s) du \Pi(ds)$ satisfies the condition $\iint_{\mathbb{R}_+ \times \mathcal{S}} (1 \wedge u) \nu(du ds) < \infty$, and let $\eta(ds)$ be a finite measure on \mathcal{S} . Let $\{\sigma_m\}$ be independent identically distributed draws from $\Pi(ds)$ and let $\{\tau_m\}$ be the successive jump times of a standard Poisson process. Set $\tau(u, s) \equiv \int_u^\infty \nu(x, s) dx$ and, for $m \in \mathbb{N}_+$, set $v_m \equiv \inf[u \geq 0 : \tau(u, \sigma_m) \leq \tau_m]$. Then the random field defined by $X[\phi] \equiv \eta[\phi] + \sum_{m < \infty} v_m \phi(\sigma_m)$ for bounded measurable $\phi(s)$ has characteristic functional (1.1).*



Note that there is no condition prohibiting fixed points of discontinuity—if $\Pi(ds)$ assigns positive mass to some point $s_0 \in \mathcal{S}$, as typically happens for posterior distributions in Bayesian analysis, then $X[\cdot]$ will be discontinuous at s_0 . It is usually more efficient numerically to sample the discontinuity points separately, however, and use the ILM algorithm only for the continuous portion of the random process or field.

4.1 Examples: Stable, Beta, Simple Homogeneous Fields

Lévy processes are of emerging importance in Bayesian hierarchical non-parametric and semiparametric modeling (Sinha and Dey, 1997). The one-sided stable processes have been used as time series innovations (Qiou and Ravishanker, 1997) and as survival analysis frailties (Qiou et al., 1997), the beta for survival analysis (Hjort, 1990), and the gamma process for point process intensities (Wolpert and Ickstadt, 1997), for example.

The ILM algorithm can be used to construct inhomogeneous analogues of all these Lévy processes: the one-sided stable distribution, with inhomogeneous Lévy measure $\nu(du ds) = u^{-1-\xi(s)} \xi(s) du \alpha(s) \Pi(ds)$ whose exponent $0 < \xi(s) < 1$ and intensity $\alpha(s)$ may vary with $s \in \mathcal{S}$, for example, or of the beta process with Lévy measure $\nu_\beta(du ds) = u^{-1} \beta(s) (1-u)^{\beta(s)-1} du \alpha(s) \Pi(ds)$ whose at-risk parameter $\beta(s)$ and mean $\alpha(s) \Pi(ds)$ need not be uniform. Evidently $\tau(u, s) = u^{-\xi(s)} \alpha(s)$ for the stable process and so the ILM jumps are available in closed form as $v_m \equiv \inf[u \geq 0 : \tau(u, \sigma_m) \leq \tau_m] = (\alpha(\sigma_m)/\tau_m)^{1/\xi(\sigma_m)}$, while the jumps for the beta process can be computed from the incomplete beta function (Abramowitz and Stegun, 1964, p. 263) using $\tau(u, s) = \alpha(s) \beta(s) B_{1-x}(\beta, 0)$.

The ILM approach is illustrated below for a stable process with constant exponent ξ and intensity α and for a beta process with constant at-risk parameter β and mean α with the prototype S-PLUS functions `sp()` and `bp()`. Figures 3 and 4 give the outputs, generated using the commands `disp.1d(sp(a=3,xi=0.9), "Stable", c("a","xi"), c(3,0.9));` and `disp.1d(bp(a=3, b=0.9), "Beta", c("a","b"), c(3,0.9));`.

Stable Process: $a = 3$, $\xi = 0.9$ FIGURE 3. Stable Process realization with $\alpha = 3$, $\xi = 0.9$.

```

"sp"<- function(a, xi, tlim=1, M = 1000) {
  x <- runif(M)*tlim;           # times x
  u <- (tlim*a/cumsum(rexp(M)))^(1/xi);   # jumps u
  ox<- order(x); x <- x[ox]; u <- u[ox]; # sort x,u on x
  invisible(list(M=M,x=x,u=u,y=cumsum(u),t=tlim)); }

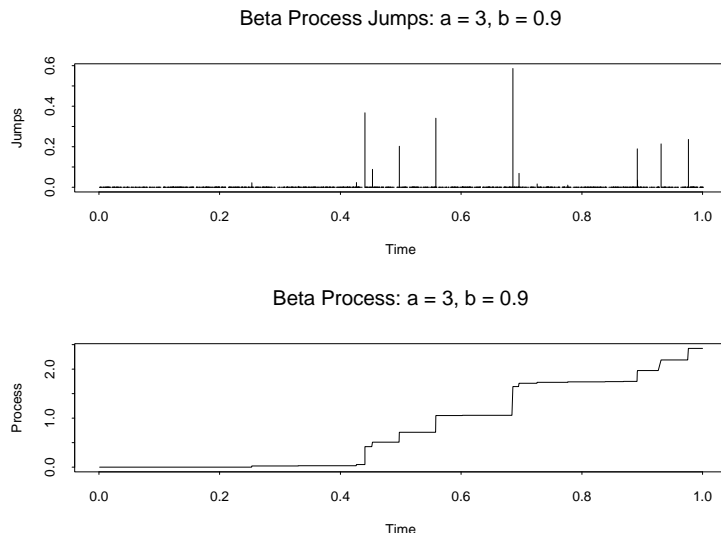
"bp"<- function(a, b, tlim=1, M = 1000, eps=1e-9) {
  x <- runif(M)*tlim;           # times x
  u <- 1-qbeta(eps*cumsum(rexp(M))/(a*b),b,eps);
  ox<- order(x); x <- x[ox]; u <- u[ox]; # sort x,u on x
  invisible(list(M=M,x=x,u=u,y=cumsum(u),t=tlim)); }

"sh"<- function(a, b, tlim=1, M = 1000, eps=1e-9) {
  x <- runif(M)*tlim;           # times x
  u <- -log(qbeta(eps*cumsum(rexp(M))/a,b,eps));
  ox<- order(x); x <- x[ox]; u <- u[ox]; # sort x,u on x
  invisible(list(M=M,x=x,u=u,y=cumsum(u),t=tlim)); }

```

The simple homogeneous processes (Ferguson and Phadia, 1979) has a similar construction:

In all cases the random field versions offer no additional difficulties, in any number of dimensions.

FIGURE 4. Beta Process realization with $\alpha = 3$, $\beta = 0.9$.

5 Comparisons with Other Methods

5.1 Method of Bondesson

For the gamma process in one dimension with $\alpha(\mathcal{S}) < \infty$ and constant $\beta(s) \equiv \beta$ the choice $\Pi(ds) = \alpha(ds)/\alpha(\mathcal{S})$ makes $\alpha(s) = \alpha(\mathcal{S})$ constant too, and the ILM algorithm will generate the jumps v_m in strictly decreasing order starting at the largest mass point of $\Gamma(ds)$. In this special case, in one dimension, $\mathcal{S} = [0, 1] \subset \mathbb{R}^1$, with Lebesgue measure for $\alpha(ds)$ and $\Pi(ds)$, the ILM method reduces to that of Bondesson (1982).

5.2 Method of Damien, Laud, Smith

A related method for generating Lévy processes was described in (Damien et al., 1995). These authors generate a finite number n of jump levels u_i from some distribution $\gamma(du)$ (they use $\gamma(du) \propto \frac{u}{1+u} \nu(du)$), take n Poisson processes P_i with rates $\lambda_i = \frac{1}{n} \frac{d\nu}{d\gamma}(u_i)$, and approximate the Lévy process X_t with Lévy measure $\nu(du)$ by $X_t \approx \sum_i u_i P_i(t)$. In our context this amounts to replacing the exact Poisson measure $H(du ds)$ on \mathbb{R}_+^2 with mean $E[H(du ds)] = \nu(du ds)$ by an approximate one concentrated on the n lines $\{(u, s) : u \in \{u_i\}\}$, with mean $E[H(du ds)] = \frac{1}{n} \sum \delta_{u_i}(du) \lambda_i ds$. For sufficiently large n (or small T) the restriction to a finite set of jump levels is scarcely noticeable; figure 5 illustrates this aspect of the method by requesting $n = 100$ levels drawn from an exponential distribution $\gamma(du)$ and drawing dotted horizontal lines at the 13 distinct levels that actually appear

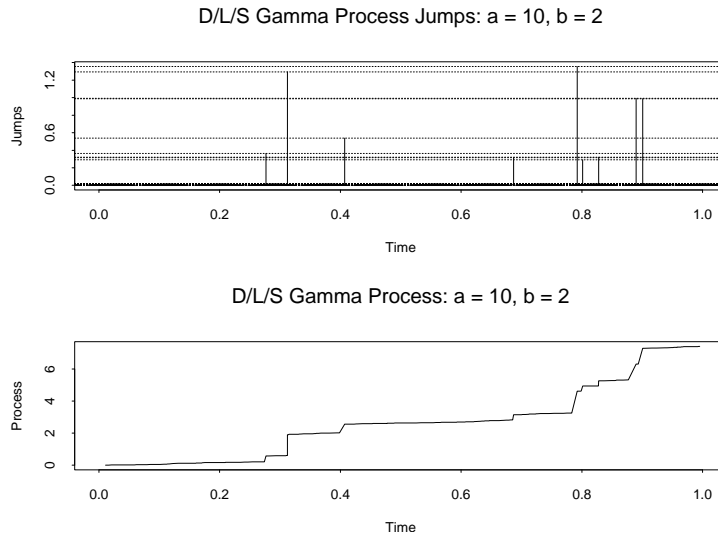


FIGURE 5. Method of Damien, Laud, and Smith.

in the realization (87 of the $n = 100$ Poisson processes had no jump during the time interval $[0, 1]$). This realization generated approximately the same number of mass points (245) as used by our ILM method illustrated in Figure 1 (250). Figure 5 was generated in S-PLUS with the commands

```
gam <- dls(a=10,b=2,n=100);
disp.1d(gam,"D/L/S Gamma",c("a","b"), c(10,2),h=gam$u);
```

using the function `dls()` given by

```
"dls"<- function(a, b, tlim=1, n = 100, aa=1, bb=b) {
  uvec <- rgamma(n,aa)*bb;
  lam <- a*bb/(n*uvec*dgamma(uvec,aa)*exp(b*uvec));
  M <- sum(mvec <- rpois(n,lam*tlim));
  x <- runif(M)*tlim; # times x
  u <- rep(uvec,mvec); # jumps u
  ox<- order(x); x <- x[ox]; u <- u[ox]; # sort on x
  invisible(list(n=n,x=x,u=u,y=cumsum(u),M=M,t=tlim)); }
```

The ILM method is more flexible in that it allows inhomogeneous processes (in the gamma example, both $\alpha(s)$ and $\beta(s)$ may be nonconstant), applies in any number of dimensions, does not require sampling from an inconvenient distribution (but when sampling from a distribution $\Pi(ds)$ that makes $\nu(u, s)$ constant in s will offer high efficiency by generating the biggest jumps v_m first) and will generate samples from the exact target distribution without any approximation error other than the unavoidable truncation error that arises from taking only finitely many jumps v_m .

5.3 Dirichlet Processes

Dirichlet processes and their mixtures frequently are chosen to represent nonparametric prior information in Bayesian analysis due to the extreme simplicity with which they can be simulated and manipulated (MacEachern, 1998). The present method offers an alternative that is almost as simple to implement: generate a gamma random process $\Gamma(s)$ with shape measure $\alpha(ds)$ and constant scale $\beta(s) \equiv 1$ and set $\mathcal{D}(s) = \Gamma(s)/\Gamma(1)$ on the unit interval $s \in \mathcal{S} = [0, 1]$. The appeal of the ILM approach is the ease with which the Dirichlet process limitations may be overcome: prior opinion can be more informed at some locations $s \in \mathcal{S}$ than at others, by allowing $\beta(s)$ to be nonconstant; spatial and spatio-temporal problems can be handled without any additional difficulty, using random fields $\mathcal{D}[\phi] \equiv \Gamma[\phi]/\Gamma[1]$ instead of processes; more robustness against prior misspecification can be obtained, by using a flatter-tailed distribution such as a stable process rather than the gamma process; mixtures are still available, just as for the Dirichlet process (a gamma random field mixture example appears in Ickstadt et al. (1998)).

6 Conclusions

The ILM algorithm enables the efficient generation of both homogeneous and inhomogeneous random processes and fields, in any number of dimensions, without any approximation other than truncation. It only requires sampling from convenient distributions (the exponential and the investigator's choice of an arbitrary distribution $\Pi(ds)$ on \mathcal{S}). It is nearly as simple to implement as Dirichlet process methods, but offers far more flexibility to enable an investigator to express a wide range of possible prior beliefs in simulation-based approaches to Bayesian nonparametric statistical analysis, especially in spatial and spatio-temporal applications where random fields are required.

The ILM method has been applied to spatial hierarchical mixture models in which observed point processes are treated as Poisson random fields whose latent intensity is given by a kernel mixture of Lévy random fields in a variety of applications areas: bioabundance (Ickstadt and Wolpert, 1997), transportation engineering (Ickstadt et al., 1998), and, in ongoing work, to problems in epidemiology and ecology where the use of inhomogeneous fields is critical to allow prior specifications that reflect spatial demographic and environmental collateral information.

Appendix 1: Display Functions

The algorithms described in this paper are straightforward to implement in any numerical programming language, and for production work should be implemented in a fast compiled language like C, FORTRAN, or JAVA. We include prototype implementations in the widely available interpreted language S-PLUS (MathSoft, 1997) in order to illustrate them and to encourage readers to experiment with them. The short functions needed to generate the processes and fields are included with the text discussion; here are the one- and two-dimensional display routines `disp.1d()` and `disp.2d()` used to generate the figures:

```
"disp.1d" <- function(p=list(), name="", pn=c("", ""),
                    pv=c(NA,NA), h=NULL) {
  oldpar <- par(mfrow=c(2,1));
  plot(p$x, p$u, type = "h", xlim=c(0,p$t),
       xlab="Time", ylab="Jumps");
  if(length(h)) {abline(h=h, lty=2)}
  title(paste(name, " Process Jumps: ", pn[1], " = ", pv[1],
             ", ", pn[2], " = ", pv[2], sep=""));
  plot(p$x, p$y, type = "l", xlim=c(0,p$t),
       xlab="Time", ylab="Process");
  title(paste(name, " Process: ", pn[1], " = ", pv[1],
             ", ", pn[2], " = ", pv[2], sep=""));
  par(mfrow=oldpar$mfrow);
  invisible(); }

"disp.2d" <- function(p=list(), name="", pn=c("", ""),
                    pv=c(NA,NA)) {
  par(mfrow=c(1,2));
  pincushion(p); # see note below
  title(paste(name, " Field Jumps:\n", pn[1], " = ", pv[1],
             ", ", pn[2], " = ", pv[2], sep=""));
  persp(p$xx, p$yy, p$zz);
  title(paste(name, " Random Field:\n", pn[1], " = ", pv[1],
             ", ", pn[2], " = ", pv[2], sep=""));
  invisible(); }
```

The following function `pincushion()` is called from within `disp.2d()` to generate a plot of the Lévy random field jumps as impulses (see Figure 2):

```
"pincushion" <- function(p=stop("Need list of x,y,u"),
                        n=length(x), screen=list(z=40, x=-60, y=0), hori=F) {
  co <- cos((screen$z * pi)/180);
  si <- sin((screen$z * pi)/180);
  xmat <- matrix(c(co,si,0, -si,co,0, 0, 0,1),
```

```

                                nr = 3);          # Transformation matrix
co <- cos((screen$x * pi)/180);
si <- sin((screen$x * pi)/180);
xmat <- matrix(c(1,0,0, 0,co,si, 0,-si,co),
              nr = 3) %*% xmat;
co <- cos((screen$y * pi)/180);
si <- sin((screen$y * pi)/180);
xmat <- matrix(c(co,0,-si, 0,1,0, si,0,co),
              nr = 3) %*% xmat;
n <- min(n,length(p$x));
or <- rev(order(p$u))[1:n];          # Sort by jump sizes:
x <- p$x[or]; y <- p$y[or]; z <- p$u[or];
zmax <- max(z); z <- 0.8*z/zmax; # Rescale jumps
box3d <- xmat %*% matrix(c(0,0,0, 0,1,0, 0,1,1, 1,1,1,
                        1,0,1, 0,0,1, 0,0,0, 1,0,0, 1,0,1, 0,0,1, 0,1,1,
                        0,1,0, 1,1,0, 1,1,1, 1,1,0, 1,0,0), nrow = 3)
box2d <- xmat %*% matrix(c(0,0,0, 0,1,0, 1,1,0,
                        1,0,0, 0,0,0), nrow = 3)
xl <- c(min(box3d[1, ]), max(box3d[1, ]));
yl <- c(min(box3d[2, ])-.75, 0.75 + max(box3d[2, ]));
plot(x = c(0,0), type = "n", xlim = xl, ylim = yl,
     axes = F, xlab = "", ylab = "");
lines(t(box2d[-3,]),lty=1); lines(t(box3d[-3,]),lty=2);
xy0 <- xmat %*% rbind(x, y, 0); # Bases of pins
xyz <- xmat %*% rbind(x, y, z); # Heads of pins
segments(xy0[1,], xy0[2,], xyz[1,], xyz[2,]); }

```

Appendix 2: Proof of Theorem 2.

Let $\epsilon > 0$, and, for $i \in \mathbb{N}$, set $a_i \equiv i\epsilon$, $A_i = [a_i, a_{i+1})$ and, for $u \geq 0$, set $A^\epsilon(u) \equiv \sum a_i 1_{A_i}(u) = \epsilon \lfloor u/\epsilon \rfloor$, u rounded down to the nearest integral multiple of ϵ . Let $\{B_j\}$ be any Borel partition of \mathcal{S} and take $\phi(s) = \sum \phi_j 1_{B_j}(s)$ and $\nu(u, s) = \sum \nu_j(u) 1_{B_j}(s)$ to be nonnegative and constant on partition elements, with $\phi(s) \leq M < \infty$ bounded. Set $\pi_j \equiv \Pi(B_j)$ and denote by N_{ij} the cardinality of the set $\{m : v_m \in A_i, \sigma_m \in B_j\}$. Then for $i \geq 1$ the N_{ij} have independent Poisson distributions with means $E[N_{ij}] = \nu(A_i \times B_j) = \pi_j \int_{A_i} \nu_j(u) du$, and the approximation $X_\epsilon[\phi] \equiv \eta[\phi] + \sum_m A^\epsilon(v_m) \phi(\sigma_m)$ has characteristic function

$$\begin{aligned}
E(e^{iX_\epsilon[\phi]}) &= e^{i\eta[\phi]} \prod_{ij} E[e^{ia_i \phi_j N_{ij}}] \\
&= \exp\left(i\eta[\phi] + \sum_{ij} [e^{ia_i \phi_j} - 1] \nu(A_i \times B_j)\right) \\
&= \exp\left(i\eta[\phi] + \iint_{\mathbb{R}_+ \times \mathcal{S}} [e^{iA^\epsilon(u)\phi(s)} - 1] \nu(du ds)\right).
\end{aligned}$$

By the dominated convergence theorem we can take $\epsilon \rightarrow 0$ to find $E[e^{iX[\phi]}] = \exp(i\eta[\phi] + \iint_{\mathbb{R}_+ \times \mathcal{S}} [e^{iu\phi(s)} - 1]\nu(du ds))$, as claimed. Now refine the partition $\{B_j\}$ to complete the proof for any bounded Borel function $\phi(s)$ and $(1 \wedge u)du\Pi(ds)$ -integrable $\nu(u, s)$. \square

Acknowledgments

This analysis was supported by U.S. E.P.A. grant CR822047-01-0, U.S. N.S.F. grant DMS-9626829, and the Deutsche Forschungsgemeinschaft.

Bibliography

- Abramowitz, M. and Stegun, I. A., eds., (1964), *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables*, volume 55 of *Applied Mathematics Series*, Washington, D.C.: National Bureau of Standards.
- Antoniak, C. E., (1974), "Mixtures of Dirichlet Processes with Applications to Bayesian Nonparametric Problems," *Annals of Statistics*, 2, 1152–1174.
- Bondesson, L., (1982), "On simulation from infinitely divisible distributions," *Advances in Applied Probability*, 14, 855–869.
- Damien, P., Laud, P. W., and Smith, A. F. M., (1995), "Approximate random variate generation from infinitely divisible distributions with applications to Bayesian inference," *Journal of the Royal Statistical Society, Series B (Methodological)*, 57, 547–563.
- Ferguson, T. S., (1973), "A Bayesian analysis of some nonparametric problems," *Annals of Statistics*, 1, 209–230.
- , (1974), "Prior distributions on spaces of probability measures," *Annals of Statistics*, 2, 615–629.
- Ferguson, T. S. and Klass, M. J., (1972), "A representation of independent increment processes without Gaussian components," *Annals of Mathematical Statistics*, 43, 1634–1643.
- Ferguson, T. S. and Phadia, E. G., (1979), "Bayesian nonparametric estimation based on censored data," *Annals of Statistics*, 7, 163–186.
- Gelfand, A. E. and Smith, A. F. M., (1990), "Sampling-based approaches to calculating marginal densities," *Journal of the American Statistical Association*, 85, 398–409.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. J., eds., (1996), *Markov Chain Monte Carlo in Practice*, New York, NY, USA: Chapman & Hall.
- Hjort, N. L., (1990), "Nonparametric Bayes estimators based on beta processes in models for life history data," *Annals of Statistics*, 18, 1259–1294.

- Ickstadt, K. and Wolpert, R. L., (1997), "Multiresolution Assessment of Forest Inhomogeneity," in *Case Studies in Bayesian Statistics, Volume III*, eds. C. Gatsonis, J. S. Hodges, R. E. Kass, R. E. McCulloch, P. Rossi, and N. D. Singpurwalla, volume 121 of *Lecture Notes in Statistics*, 371–386, New York, NY, USA: Springer-Verlag.
- Ickstadt, K., Wolpert, R. L., and Lu, X., (1998), "Modeling travel demand in Portland, Oregon," in Müller et al. (1998), 0–0.
- Jacod, J. and Shiryaev, A. N., (1987), *Limit Theorems for Stochastic Processes*, volume 288 of *Grundlehren der mathematischen Wissenschaften*, Berlin: Springer-Verlag.
- Lévy, P., (1937), *Théorie de L'Addition des Variables Aléatoires*, Paris: Gauthier-Villars.
- MacEachern, S. N., (1998), "Computational Methods for Mixture of Dirichlet Process Models," in Müller et al. (1998), 0–0.
- MathSoft, D. A. P. D., (1997), *S-Plus User's Guide*, Seattle: MathSoft.
- Müller, P., Dey, D., and Sinha, D., eds., (1998), *Practical Nonparametric and Semiparametric Bayesian Statistics*, Lecture Notes in Statistics, New York, NY, USA: Springer-Verlag.
- Qiou, Z. and Ravishanker, N., (1997), "Bayesian inference for time series with stable innovations," To appear in *Journal of Time Series Analysis*.
- Qiou, Z., Ravishanker, N., and Dey, D. K., (1997), "Multivariate survival analysis with positive stable frailties," Under revision for *Biometrics*.
- Rogers, L. and Williams, D., (1994), *Diffusions, Markov Processes, and Martingales*, volume 1, New York, NY, USA: John Wiley & Sons, second edition.
- Sinha, D. and Dey, D., (1997), "Semiparametric Bayesian Analysis of Survival Data," *Journal of the American Statistical Association*, 92, 1195–1212.
- Tierney, L., (1994), "Markov chains for exploring posterior distributions," *Annals of Statistics*, 22, 1701–1762, (With discussion).
- Wolpert, R. L. and Ickstadt, K., (1997), "Poisson/gamma random field models for spatial statistics," Discussion Paper 95-43, Duke University ISDS, USA, to appear in *Biometrika*.