

# Package ‘bssm’

September 21, 2021

**Type** Package

**Title** Bayesian Inference of Non-Linear and Non-Gaussian State Space Models

**Version** 1.1.7-1

**Description** Efficient methods for Bayesian inference of state space models via particle Markov chain Monte Carlo (MCMC) and MCMC based on parallel importance sampling type weighted estimators (Vihola, Helske, and Franks, 2020, <[doi:10.1111/sjos.12492](https://doi.org/10.1111/sjos.12492)>). Gaussian, Poisson, binomial, negative binomial, and Gamma observation densities and basic stochastic volatility models with linear-Gaussian state dynamics, as well as general non-linear Gaussian models and discretised diffusion models are supported.

**License** GPL (>= 2)

**Depends** R (>= 3.5.0)

**Suggests** covr, dplyr, ggplot2 (>= 2.0.0), Hmisc, KFAS (>= 1.2.1), knitr (>= 1.11), MASS, posterior, rmarkdown (>= 0.8.1), ramcmc, sde, sitmo, testthat

**Imports** checkmate, coda (>= 0.18-1), diagis, Rcpp (>= 0.12.3)

**LinkingTo** Rcpp, RcppArmadillo, ramcmc, sitmo

**SystemRequirements** C++11, pandoc (>= 1.12.3, needed for vignettes)

**RoxygenNote** 7.1.2

**VignetteBuilder** knitr

**BugReports** <https://github.com/helske/bssm/issues>

**URL** <https://github.com/helske/bssm>

**ByteCompile** true

**Encoding** UTF-8

**NeedsCompilation** yes

**Author** Jouni Helske [aut, cre] (<<https://orcid.org/0000-0001-7130-793X>>),  
Matti Vihola [aut] (<<https://orcid.org/0000-0002-8041-7222>>)

**Maintainer** Jouni Helske <jouni.helske@iki.fi>

**Repository** CRAN

**Date/Publication** 2021-09-21 13:10:45 UTC

## R topics documented:

ar1_lg . . . . .	3
ar1_ng . . . . .	4
as.data.frame.mcmc_output . . . . .	5
as_bssm . . . . .	6
as_draws_df.mcmc_output . . . . .	7
bootstrap_filter . . . . .	8
bsm_lg . . . . .	10
bsm_ng . . . . .	12
bssm . . . . .	14
cpp_example_model . . . . .	15
drownings . . . . .	15
ekf . . . . .	16
ekf_smoother . . . . .	17
ekpf_filter . . . . .	18
exchange . . . . .	19
expand_sample . . . . .	20
fast_smoother . . . . .	21
gaussian_approx . . . . .	22
importance_sample . . . . .	23
kfilter . . . . .	24
logLik.gaussian . . . . .	25
particle_smoother . . . . .	27
poisson_series . . . . .	29
post_correct . . . . .	30
predict.mcmc_output . . . . .	32
print.mcmc_output . . . . .	35
run_mcmc . . . . .	35
sim_smoother . . . . .	41
ssm_mlg . . . . .	43
ssm_mng . . . . .	45
ssm_nlg . . . . .	47
ssm_sde . . . . .	49
ssm_ulg . . . . .	51
ssm_ung . . . . .	55
suggest_N . . . . .	57
summary.mcmc_output . . . . .	59
svm . . . . .	59
ukf . . . . .	61
uniform_prior . . . . .	62

**Index**

**65**

ar1\_lg

*Univariate Gaussian model with AR(1) latent process***Description**

Constructs a simple Gaussian model where the state dynamics follow an AR(1) process.

**Usage**

```
ar1_lg(y, rho, sigma, mu, sd_y, beta, xreg = NULL)
```

**Arguments**

y	Vector or a ts object of observations.
rho	Prior for autoregressive coefficient. Should be an object of class bssm_prior.
sigma	Prior for the standard deviation of noise of the AR-process. Should be an object of class bssm_prior
mu	A fixed value or a prior for the stationary mean of the latent AR(1) process. Should be an object of class bssm_prior or scalar value defining a fixed mean such as 0.
sd_y	A prior for the standard deviation of observation equation.
beta	Prior for the regression coefficients. Should be an object of class bssm_prior or bssm_prior_list (in case of multiple coefficients) or missing in case of no covariates.
xreg	Matrix containing covariates with number of rows matching the length of y.

**Value**

Object of class ar1\_lg.

**Examples**

```
set.seed(1)
mu <- 2
rho <- 0.7
sd_y <- 0.1
sigma <- 0.5
beta <- -1
x <- rnorm(30)
z <- y <- numeric(30)
z[1] <- rnorm(1, mu, sigma / sqrt(1 - rho^2))
y[1] <- rnorm(1, beta * x[1] + z[1], sd_y)
for(i in 2:30) {
  z[i] <- rnorm(1, mu * (1 - rho) + rho * z[i - 1], sigma)
  y[i] <- rnorm(1, beta * x[i] + z[i], sd_y)
}
model <- ar1_lg(y, rho = uniform(0.5, -1, 1),
```

```

sigma = halfnormal(1, 10), mu = normal(0, 0, 1),
sd_y = halfnormal(1, 10),
xreg = x, beta = normal(0, 0, 1))
out <- run_mcmc(model, iter = 2e4)
summary(out, return_se = TRUE)

```

ar1\_ng

*Non-Gaussian model with AR(1) latent process***Description**

Constructs a simple non-Gaussian model where the state dynamics follow an AR(1) process.

**Usage**

```
ar1_ng(y, rho, sigma, mu, distribution, phi, u, beta, xreg = NULL)
```

**Arguments**

y	Vector or a ts object of observations.
rho	Prior for autoregressive coefficient. Should be an object of class <code>bssm_prior</code> .
sigma	Prior for the standard deviation of noise of the AR-process. Should be an object of class <code>bssm_prior</code>
mu	A fixed value or a prior for the stationary mean of the latent AR(1) process. Should be an object of class <code>bssm_prior</code> or scalar value defining a fixed mean such as 0.
distribution	Distribution of the observed time series. Possible choices are "poisson", "binomial", "gamma", and "negative binomial".
phi	Additional parameter relating to the non-Gaussian distribution. For negative binomial distribution this is the dispersion term, for gamma distribution this is the shape parameter, and for other distributions this is ignored. Should be an object of class <code>bssm_prior</code> or a positive scalar.
u	Vector of positive constants for non-Gaussian models. For Poisson, gamma, and negative binomial distribution, this corresponds to the offset term. For binomial, this is the number of trials.
beta	Prior for the regression coefficients. Should be an object of class <code>bssm_prior</code> or <code>bssm_prior_list</code> (in case of multiple coefficients) or missing in case of no covariates.
xreg	Matrix containing covariates with number of rows matching the length of y.

**Value**

Object of class `ar1_ng`.

**Examples**

```

model <- ar1_ng(discoveries, rho = uniform(0.5,-1,1),
  sigma = halfnormal(0.1, 1), mu = normal(0, 0, 1),
  distribution = "poisson")
out <- run_mcmc(model, iter = 1e4, mcmc_type = "approx",
  output_type = "summary")

ts.plot(cbind(discoveries, exp(out$alphahat)), col = 1:2)

set.seed(1)
n <- 30
phi <- 2
rho <- 0.9
sigma <- 0.1
beta <- 0.5
u <- rexp(n, 0.1)
x <- rnorm(n)
z <- y <- numeric(n)
z[1] <- rnorm(1, 0, sigma / sqrt(1 - rho^2))
y[1] <- rbinom(1, mu = u * exp(beta * x[1] + z[1]), size = phi)
for(i in 2:n) {
  z[i] <- rnorm(1, rho * z[i - 1], sigma)
  y[i] <- rbinom(1, mu = u * exp(beta * x[i] + z[i]), size = phi)
}

model <- ar1_ng(y, rho = uniform_prior(0.9, 0, 1),
  sigma = gamma_prior(0.1, 2, 10), mu = 0.,
  phi = gamma_prior(2, 2, 1), distribution = "negative binomial",
  xreg = x, beta = normal_prior(0.5, 0, 1), u = u)

```

---

as.data.frame.mcmc\_output

*Convert MCMC chain to data.frame*


---

**Description**

Converts the MCMC chain output of `run_mcmc` to data.frame.

**Usage**

```

## S3 method for class 'mcmc_output'
as.data.frame(
  x,
  row.names,
  optional,
  variable = c("theta", "states"),
  times,
  states,

```

```

    expand = !(x$mcmc_type %in% paste0("is", 1:3)),
    ...
  )

```

### Arguments

x	Output from <code>run_mcmc</code> .
row.names	Ignored.
optional	Ignored.
variable	Return samples of "theta" (default) or "states"?
times	Vector of indices. In case of states, what time points to return? Default is all.
states	Vector of indices. In case of states, what states to return? Default is all.
expand	Should the jump-chain be expanded? Defaults to TRUE for non-IS-MCMC, and FALSE for IS-MCMC. For <code>expand = FALSE</code> and always for IS-MCMC, the resulting data.frame contains variable weight (= counts * IS-weights).
...	Ignored.

### Examples

```

data("poisson_series")
model <- bsm_ng(y = poisson_series,
sd_slope = halfnormal(0.1, 0.1),
sd_level = halfnormal(0.1, 1),
distribution = "poisson")

out <- run_mcmc(model, iter = 2000, particles = 10)
head(as.data.frame(out, variable = "theta"))
head(as.data.frame(out, variable = "state"))

# don't expand the jump chain:
head(as.data.frame(out, variable = "theta", expand = FALSE))

# IS-weighted version:
out_is <- run_mcmc(model, iter = 2000, particles = 10, mcmc_type = "is2")
head(as.data.frame(out_is, variable = "theta"))

```

---

as\_bssm

---

*Convert KFAS Model to bssm Model*


---

### Description

Converts `SSModel` object of KFAS package to general bssm model of type `ssm_ulg`, `ssm_mlg`, `ssm_ung` or `ssm_mng`.

### Usage

```
as_bssm(model, kappa = 100, ...)
```

**Arguments**

model	Object of class SSMModel.
kappa	For SSMModel object, a prior variance for initial state used to replace exact diffuse elements of the original model.
...	Additional arguments to model building functions of bssm (such as prior and updating functions).

**Value**

Object of class ssm\_ulg, ssm\_mlg, ssm\_ung or ssm\_mng.

**Examples**

```
library("KFAS")
model_KFAS <- SSMModel(Nile ~
  SSMtrend(1, Q = 2, P1 = 1e4), H = 2)
model_bssm <- as_bssm(model_KFAS)
logLik(model_KFAS)
logLik(model_bssm)
```

---

as\_draws\_df.mcmc\_output

*Convert run\_mcmc output to draws\_df format*

---

**Description**

Converts MCMC output from run\_mcmc call to a draws\_df format of the posterior package. This enables the use of diagnostics and plotting methods of posterior and bayesplot packages. Note though that if run\_mcmc used IS-MCMC method, the resulting weight column of the output is ignored by the aforementioned packages, i.e. the results correspond to approximate MCMC.

**Usage**

```
as_draws_df.mcmc_output(x)
```

```
as_draws.mcmc_output(x)
```

**Arguments**

x                    An object of class mcmc\_output

**Value**

A draws\_df object.

**Examples**

```

model <- bsm_lg(Nile,
  sd_y = tnormal(init = 100, mean = 100, sd = 100, min = 0),
  sd_level = tnormal(init = 50, mean = 50, sd = 100, min = 0),
  a1 = 1000, P1 = 500^2)

fit1 <- run_mcmc(model, iter = 2000)
library("posterior")
draws <- as_draws(fit1)
head(draws, 4)
ess_bulk(draws$sd_y)
summary(fit1, return_se = TRUE)

# More chains:
model$theta[] <- c(50, 150) # change initial value
fit2 <- run_mcmc(model, iter = 2000)
model$theta[] <- c(150, 50) # change initial value
fit3 <- run_mcmc(model, iter = 2000)

draws <- bind_draws(as_draws(fit1),
  as_draws(fit2), as_draws(fit3), along = "chain")
# it is actually enough to transform first mcmc_output to draws object,
# rest are transformed automatically inside bind_draws
rhat(draws$sd_y)
ess_bulk(draws$sd_y)
ess_tail(draws$sd_y)

```

---

bootstrap\_filter      *Bootstrap Filtering*

---

**Description**

Function `bootstrap_filter` performs a bootstrap filtering with stratification resampling.

**Usage**

```

bootstrap_filter(model, particles, ...)

## S3 method for class 'gaussian'
bootstrap_filter(
  model,
  particles,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

## S3 method for class 'nongaussian'

```



```

bootstrap_filter(
  model,
  particles,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

## S3 method for class 'ssm_nlg'
bootstrap_filter(
  model,
  particles,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

## S3 method for class 'ssm_sde'
bootstrap_filter(
  model,
  particles,
  L,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

```

### Arguments

model	A model object of class <code>bssm_model</code> .
particles	Number of particles as a positive integer.
...	Ignored.
seed	Seed for RNG (non-negative integer).
L	Positive integer defining the discretization level for SDE models.

### Value

List with samples (`alpha`) from the filtering distribution and corresponding weights (`weights`), as well as filtered and predicted states and corresponding covariances (`at`, `att`, `Pt`, `Ptt`), and estimated log-likelihood (`logLik`).

### References

Gordon, NJ, Salmond, DJ, Smith, AFM (1993) Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings F, 140(2), p. 107-113.

### Examples

```

set.seed(1)
x <- cumsum(rnorm(50))
y <- rnorm(50, x, 0.5)

```

```

model <- bsm_lg(y, sd_y = 0.5, sd_level = 1, P1 = 1)

out <- bootstrap_filter(model, particles = 1000)
ts.plot(cbind(y, x, out$att), col = 1:3)
ts.plot(cbind(kfilter(model)$att, out$att), col = 1:3)

data("poisson_series")
model <- bsm_ng(poisson_series, sd_level = 0.1, sd_slope = 0.01,
  P1 = diag(1, 2), distribution = "poisson")

out <- bootstrap_filter(model, particles = 100)
ts.plot(cbind(poisson_series, exp(out$att[, 1])), col = 1:2)

```

---

bsm\_lg

*Basic Structural (Time Series) Model*


---

## Description

Constructs a basic structural model with local level or local trend component and seasonal component.

## Usage

```

bsm_lg(
  y,
  sd_y,
  sd_level,
  sd_slope,
  sd_seasonal,
  beta,
  xreg = NULL,
  period = frequency(y),
  a1 = NULL,
  P1 = NULL,
  D = NULL,
  C = NULL
)

```

## Arguments

<code>y</code>	Vector or a ts object of observations.
<code>sd_y</code>	Standard deviation of the noise of observation equation. Should be an object of class <code>bssm_prior</code> or scalar.
<code>sd_level</code>	Standard deviation of the noise of level equation. Should be an object of class <code>bssm_prior</code> or scalar value defining a known value such as 0.

sd_slope	Standard deviation of the noise of slope equation. Should be an object of class <code>bssm_prior</code> , scalar value defining a known value such as 0, or missing, in which case the slope term is omitted from the model.
sd_seasonal	Standard deviation of the noise of seasonal equation. Should be an object of class <code>bssm_prior</code> , scalar value defining a known value such as 0, or missing, in which case the seasonal term is omitted from the model.
beta	Prior for the regression coefficients. Should be an object of class <code>bssm_prior</code> or <code>bssm_prior_list</code> (in case of multiple coefficients) or missing in case of no covariates.
xreg	Matrix containing covariates with number of rows matching the length of <code>y</code> .
period	Length of the seasonal pattern. Default is <code>frequency(y)</code> . Must be a positive integer greater than 2 and less than the length of the input time series.
a1	Prior means for the initial states (level, slope, seasonals). Defaults to vector of zeros.
P1	Prior covariance for the initial states (level, slope, seasonals). Default is diagonal matrix with 1000 on the diagonal.
D, C	Intercept terms for observation and state equations, given as a length <code>n</code> vector and <code>m</code> times <code>n</code> matrix respectively (or scalar and <code>m</code> times 1 matrix).

### Value

Object of class `bsm_lg`.

### Examples

```
prior <- uniform(0.1 * sd(log10(UKgas)), 0, 1)
# period here is redundant as frequency(UKgas) = 4
model <- bsm_lg(log10(UKgas), sd_y = prior, sd_level = prior,
  sd_slope = prior, sd_seasonal = prior, period = 4)

mcmc_out <- run_mcmc(model, iter = 5000)
summary(expand_sample(mcmc_out, "theta"))$stat
mcmc_out$theta[which.max(mcmc_out$posterior), ]
sqrt((fit <- StructTS(log10(UKgas), type = "BSM"))$coef)[c(4, 1:3)]

set.seed(1)
n <- 10
x <- rnorm(n)
level <- numeric(n)
level[1] <- rnorm(1)
for (i in 2:n) level[i] <- rnorm(1, -0.2 + level[i-1], sd = 0.1)
y <- rnorm(n, 2.1 + x + level)
model <- bsm_lg(y, sd_y = halfnormal(1, 5), sd_level = 0.1, a1 = level[1],
  P1 = matrix(0, 1, 1), xreg = x, beta = normal(1, 0, 1),
  D = 2.1, C = matrix(-0.2, 1, 1))

ts.plot(cbind(fast_smoother(model), level), col = 1:2)
```

bsm\_ng

*Non-Gaussian Basic Structural (Time Series) Model***Description**

Constructs a non-Gaussian basic structural model with local level or local trend component, a seasonal component, and regression component (or subset of these components).

**Usage**

```
bsm_ng(
  y,
  sd_level,
  sd_slope,
  sd_seasonal,
  sd_noise,
  distribution,
  phi,
  u,
  beta,
  xreg = NULL,
  period = frequency(y),
  a1 = NULL,
  P1 = NULL,
  C = NULL
)
```

**Arguments**

<code>y</code>	Vector or a ts object of observations.
<code>sd_level</code>	Standard deviation of the noise of level equation. Should be an object of class <code>bssm_prior</code> or scalar value defining a known value such as 0.
<code>sd_slope</code>	Standard deviation of the noise of slope equation. Should be an object of class <code>bssm_prior</code> , scalar value defining a known value such as 0, or missing, in which case the slope term is omitted from the model.
<code>sd_seasonal</code>	Standard deviation of the noise of seasonal equation. Should be an object of class <code>bssm_prior</code> , scalar value defining a known value such as 0, or missing, in which case the seasonal term is omitted from the model.
<code>sd_noise</code>	Prior for the standard deviation of the additional noise term to be added to linear predictor, defined as an object of class <code>bssm_prior</code> . If missing, no additional noise term is used.
<code>distribution</code>	Distribution of the observed time series. Possible choices are "poisson", "binomial", "gamma", and "negative binomial".

phi	Additional parameter relating to the non-Gaussian distribution. For negative binomial distribution this is the dispersion term, for gamma distribution this is the shape parameter, and for other distributions this is ignored. Should be an object of class <code>bssm_prior</code> or a positive scalar.
u	Vector of positive constants for non-Gaussian models. For Poisson, gamma, and negative binomial distribution, this corresponds to the offset term. For binomial, this is the number of trials.
beta	Prior for the regression coefficients. Should be an object of class <code>bssm_prior</code> or <code>bssm_prior_list</code> (in case of multiple coefficients) or missing in case of no covariates.
xreg	Matrix containing covariates with number of rows matching the length of <code>y</code> .
period	Length of the seasonal pattern. Default is <code>frequency(y)</code> . Must be a positive integer greater than 2 and less than the length of the input time series.
a1	Prior means for the initial states (level, slope, seasonals). Defaults to vector of zeros.
P1	Prior covariance for the initial states (level, slope, seasonals). Default is diagonal matrix with 1000 on the diagonal.
C	Intercept terms for state equation, given as a $m \times n$ or $m \times 1$ matrix.

### Value

Object of class `bsm_ng`.

### Examples

```

model <- bsm_ng(Seatbelts[, "VanKilled"], distribution = "poisson",
  sd_level = halfnormal(0.01, 1),
  sd_seasonal = halfnormal(0.01, 1),
  beta = normal(0, 0, 10),
  xreg = Seatbelts[, "law"],
  # default values, just for illustration
  period = 12,
  a1 = rep(0, 1 + 11), # level + period - 1 seasonal states
  P1 = diag(1, 12),
  C = matrix(0, 12, 1),
  u = rep(1, nrow(Seatbelts)))

set.seed(123)
mcmc_out <- run_mcmc(model, iter = 5000, particles = 10, mcmc_type = "da")
mcmc_out$acceptance_rate
theta <- expand_sample(mcmc_out, "theta")
plot(theta)
summary(theta)

library("ggplot2")
ggplot(as.data.frame(theta[,1:2]), aes(x = sd_level, y = sd_seasonal)) +
  geom_point() + stat_density2d(aes(fill = ..level.., alpha = ..level..),
  geom = "polygon") + scale_fill_continuous(low = "green", high = "blue") +

```

```

    guides(alpha = "none")

# Traceplot using as.data.frame method for MCMC output
library("dplyr")
as.data.frame(mcmc_out) %>%
  filter(variable == "sd_level") %>%
  ggplot(aes(y = value, x = iter)) + geom_line()

# Model with slope term and additional noise to linear predictor to capture
# excess variation
model2 <- bsm_ng(Seatbelts[, "VanKilled"], distribution = "poisson",
  sd_level = halfnormal(0.01, 1),
  sd_seasonal = halfnormal(0.01, 1),
  beta = normal(0, 0, 10),
  xreg = Seatbelts[, "law"],
  sd_slope = halfnormal(0.01, 0.1),
  sd_noise = halfnormal(0.01, 1))

# instead of extra noise term, model using negative binomial distribution:
model3 <- bsm_ng(Seatbelts[, "VanKilled"],
  distribution = "negative binomial",
  sd_level = halfnormal(0.01, 1),
  sd_seasonal = halfnormal(0.01, 1),
  beta = normal(0, 0, 10),
  xreg = Seatbelts[, "law"],
  sd_slope = halfnormal(0.01, 0.1),
  phi = gamma_prior(1, 5, 5))

```

**Description**

This package contains functions for efficient Bayesian inference of state space models, where model is assumed to be either

**Details**

- \* Exponential family state space models, where the state equation is linear Gaussian, and the conditional observation density is either Gaussian, Poisson, binomial, negative binomial or Gamma density.
- \* Basic stochastic volatility model.
- \* General non-linear model with Gaussian noise terms.
- \* Model with continuous SDE dynamics.

For formal definition of the currently supported models and methods, as well as some theory behind the IS-MCMC and  $\psi$ -APF, see Helske and Vihola (2021), Vihola, Helske, Franks (2020) and the package vignettes.

**References**

Helske J, Vihola M (2021). bssm: Bayesian Inference of Non-linear and Non-Gaussian State Space Models in R. ArXiv 2101.08492, <URL: <https://arxiv.org/abs/2101.08492>>.

Vihola, M, Helske, J, Franks, J. (2020). Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. Scand J Statist. 1-38. <https://doi.org/10.1111/sjos.12492>

---

cpp\_example\_model      *Example C++ Codes for Non-Linear and SDE Models*

---

**Description**

Example C++ Codes for Non-Linear and SDE Models

**Usage**

```
cpp_example_model(example, return_code = FALSE)
```

**Arguments**

example	Name of the example model. Run <code>cpp_example_model("abc")</code> to get the names of possible models.
return_code	If TRUE, will not compile the model but only returns the corresponding code.

**Value**

Returns pointers to the C++ snippets defining the model, or in case of `return_code = TRUE`, returns the example code without compiling.

**Examples**

```
cpp_example_model("sde_poisson_OU", return_code = TRUE)
```

---

drownings      *Deaths by drowning in Finland in 1969-2019*

---

**Description**

Dataset containing number of deaths by drowning in Finland in 1969-2019, corresponding population sizes (in hundreds of thousands), and yearly average summer temperatures (June to August), based on simple unweighted average of three weather stations: Helsinki (Southern Finland), Jyväskylä (Central Finland), and Sodankylä (Northern Finland).

**Format**

A time series object containing 51 observations.

**Source**

Statistics Finland <https://pxnet2.stat.fi/PXWeb/pxweb/en/StatFin/>.

**Examples**

```
data("drownings")
model <- bsm_ng(drownings[, "deaths"], u = drownings[, "population"],
  xreg = drownings[, "summer_temp"], distribution = "poisson",
  beta = normal(0, 0, 1),
  sd_level = gamma_prior(0.1, 2, 10), sd_slope = gamma_prior(0, 2, 10))

fit <- run_mcmc(model, iter = 5000,
  output_type = "summary", mcmc_type = "approx")
fit
ts.plot(model$y/model$u, exp(fit$alphahat[, 1]), col = 1:2)
```

ekf

*(Iterated) Extended Kalman Filtering***Description**

Function `ekf` runs the (iterated) extended Kalman filter for the given non-linear Gaussian model of class `ssm_nlg`, and returns the filtered estimates and one-step-ahead predictions of the states  $\alpha_t$  given the data up to time  $t$ .

**Usage**

```
ekf(model, iekf_iter = 0)
```

**Arguments**

`model` Model of class `ssm_nlg`.  
`iekf_iter` Non-negative integer. The default zero corresponds to normal EKF, whereas `iekf_iter > 0` corresponds to iterated EKF with `iekf_iter` iterations.

**Value**

List containing the log-likelihood, one-step-ahead predictions at and filtered estimates at `t` of states, and the corresponding variances `Pt` and `Ptt`.

**Examples**

```
# Takes a while on CRAN
set.seed(1)
mu <- -0.2
rho <- 0.7
sigma_y <- 0.1
sigma_x <- 1
```



```

x <- numeric(50)
x[1] <- rnorm(1, mu, sigma_x / sqrt(1 - rho^2))
for(i in 2:length(x)) {
  x[i] <- rnorm(1, mu * (1 - rho) + rho * x[i - 1], sigma_x)
}
y <- rnorm(50, exp(x), sigma_y)

pntrs <- cpp_example_model("nlg_ar_exp")

model_nlg <- ssm_nlg(y = y, a1 = pntrs$a1, P1 = pntrs$P1,
  Z = pntrs$Z_fn, H = pntrs$H_fn, T = pntrs$T_fn, R = pntrs$R_fn,
  Z_gn = pntrs$Z_gn, T_gn = pntrs$T_gn,
  theta = c(mu = mu, rho = rho,
    log_sigma_x = log(sigma_x), log_sigma_y = log(sigma_y)),
  log_prior_pdf = pntrs$log_prior_pdf,
  n_states = 1, n_etas = 1, state_names = "state")

out_ekf <- ekf(model_nlg, iekf_iter = 0)
out_iekf <- ekf(model_nlg, iekf_iter = 5)
ts.plot(cbind(x, out_ekf$att, out_iekf$att), col = 1:3)

```

---

ekf\_smoother

*Extended Kalman Smoothing*


---

### Description

Function `ekf_smoother` runs the (iterated) extended Kalman smoother for the given non-linear Gaussian model of class `ssm_nlg`, and returns the smoothed estimates of the states and the corresponding variances. Function `ekf_fast_smoother` computes only smoothed estimates of the states.

### Usage

```
ekf_smoother(model, iekf_iter = 0)
```

```
ekf_fast_smoother(model, iekf_iter = 0)
```

### Arguments

<code>model</code>	Model of class <code>ssm_nlg</code> .
<code>iekf_iter</code>	Non-negative integer. The default zero corresponds to normal EKF, whereas <code>iekf_iter &gt; 0</code> corresponds to iterated EKF with <code>iekf_iter</code> iterations.

### Value

List containing the log-likelihood, smoothed state estimates  $\hat{\alpha}_t$ , and the corresponding variances  $V_t$  and  $P_{tt}$ .

**Examples**

```

# Takes a while on CRAN
set.seed(1)
mu <- -0.2
rho <- 0.7
sigma_y <- 0.1
sigma_x <- 1
x <- numeric(50)
x[1] <- rnorm(1, mu, sigma_x / sqrt(1 - rho^2))
for(i in 2:length(x)) {
  x[i] <- rnorm(1, mu * (1 - rho) + rho * x[i - 1], sigma_x)
}
y <- rnorm(length(x), exp(x), sigma_y)

pntrs <- cpp_example_model("nlg_ar_exp")

model_nlg <- ssm_nlg(y = y, a1 = pntrs$a1, P1 = pntrs$P1,
  Z = pntrs$Z_fn, H = pntrs$H_fn, T = pntrs$T_fn, R = pntrs$R_fn,
  Z_gn = pntrs$Z_gn, T_gn = pntrs$T_gn,
  theta = c(mu = mu, rho = rho,
    log_sigma_x = log(sigma_x), log_sigma_y = log(sigma_y)),
  log_prior_pdf = pntrs$log_prior_pdf,
  n_states = 1, n_etas = 1, state_names = "state")

out_ekf <- ekf_smoother(model_nlg, iekf_iter = 0)
out_iekf <- ekf_smoother(model_nlg, iekf_iter = 1)
ts.plot(cbind(x, out_ekf$alphahat, out_iekf$alphahat), col = 1:3)

```

---

ekpf\_filter

*Extended Kalman Particle Filtering*


---

**Description**

Function `ekpf_filter` performs an extended Kalman particle filtering with stratification resampling, based on Van Der Merwe et al (2001).

**Usage**

```

ekpf_filter(object, particles, ...)

## S3 method for class 'ssm_nlg'
ekpf_filter(
  object,
  particles,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

```

**Arguments**

object	Model of class <code>ssm_nlg</code> .
particles	Number of particles as a positive integer.
...	Ignored.
seed	Seed for RNG (positive integer).

**Value**

A list containing samples, filtered estimates and the corresponding covariances, weights, and an estimate of log-likelihood.

**References**

Van Der Merwe, R., Doucet, A., De Freitas, N., & Wan, E. A. (2001). The unscented particle filter. In *Advances in neural information processing systems* (pp. 584-590).

**Examples**

```
# Takes a while
set.seed(1)
n <- 50
x <- y <- numeric(n)
y[1] <- rnorm(1, exp(x[1]), 0.1)
for(i in 1:(n-1)) {
  x[i+1] <- rnorm(1, sin(x[i]), 0.1)
  y[i+1] <- rnorm(1, exp(x[i+1]), 0.1)
}

pntrs <- cpp_example_model("nlg_sin_exp")

model_nlg <- ssm_nlg(y = y, a1 = pntrs$a1, P1 = pntrs$P1,
  Z = pntrs$Z_fn, H = pntrs$H_fn, T = pntrs$T_fn, R = pntrs$R_fn,
  Z_gn = pntrs$Z_gn, T_gn = pntrs$T_gn,
  theta = c(log_H = log(0.1), log_R = log(0.1)),
  log_prior_pdf = pntrs$log_prior_pdf,
  n_states = 1, n_etas = 1, state_names = "state")

out <- ekpf_filter(model_nlg, particles = 100)
ts.plot(cbind(x, out$at[1:n], out$att[1:n]), col = 1:3)
```

---

exchange

*Pound/Dollar daily exchange rates*

---

**Description**

Dataset containing daily log-returns from 1/10/81-28/6/85 as in [1]

**Format**

A vector of length 945.

**Source**

<http://www.ssfpack.com/DKbook.html>.

**References**

James Durbin, Siem Jan Koopman (2012). Time Series Analysis by State Space Methods. Oxford University Press.

**Examples**

```
data("exchange")
model <- svm(exchange, rho = uniform(0.97, -0.999, 0.999),
  sd_ar = halfnormal(0.175, 2), mu = normal(-0.87, 0, 2))

out <- particle_smoother(model, particles = 500)
plot.ts(cbind(model$y, exp(out$alphahat)))
```

---

expand\_sample

*Expand the Jump Chain representation*

---

**Description**

The MCMC algorithms of `bssm` use a jump chain representation where we store the accepted values and the number of times we stayed in the current value. Although this saves bit memory and is especially convenient for IS-corrected MCMC, sometimes we want to have the usual sample paths. Function `expand_sample` returns the expanded sample based on the counts. Note that for IS-corrected output the expanded sample corresponds to the approximate posterior.

**Usage**

```
expand_sample(x, variable = "theta", times, states, by_states = TRUE)
```

**Arguments**

<code>x</code>	Output from <code>run_mcmc</code> .
<code>variable</code>	Expand parameters "theta" or states "states".
<code>times</code>	Vector of indices. In case of states, what time points to expand? Default is all.
<code>states</code>	Vector of indices. In case of states, what states to expand? Default is all.
<code>by_states</code>	If TRUE (default), return list by states. Otherwise by time.

---

fast_smoother	<i>Kalman Smoothing</i>
---------------	-------------------------

---

### Description

Methods for Kalman smoothing of the states. Function `fast_smoother` computes only smoothed estimates of the states, and function `smoother` computes also smoothed variances.

### Usage

```
fast_smoother(model, ...)  
  
## S3 method for class 'gaussian'  
fast_smoother(model, ...)  
  
smoother(model, ...)  
  
## S3 method for class 'gaussian'  
smoother(model, ...)
```

### Arguments

<code>model</code>	Model to be approximated. Should be of class <code>bsm_ng</code> , <code>ar1_ng_svm</code> , <code>ssm_ung</code> , or <code>ssm_mng</code> , or <code>ssm_nlg</code> , i.e. non-gaussian or non-linear <code>bssm_model</code> .
<code>...</code>	Ignored.

### Details

For non-Gaussian models, the smoothing is based on the approximate Gaussian model.

### Value

Matrix containing the smoothed estimates of states, or a list with the smoothed states and the variances.

### Examples

```
model <- bsm_lg(Nile,  
  sd_level = tnormal(120, 100, 20, min = 0),  
  sd_y = tnormal(50, 50, 25, min = 0),  
  a1 = 1000, P1 = 200)  
ts.plot(cbind(Nile, fast_smoother(model)), col = 1:2)  
model <- bsm_lg(Nile,  
  sd_y = tnormal(120, 100, 20, min = 0),  
  sd_level = tnormal(50, 50, 25, min = 0),  
  a1 = 1000, P1 = 500^2)  
  
out <- smoother(model)
```

```
ts.plot(cbind(Nile, out$alphahat), col = 1:2)
ts.plot(sqrt(out$Vt[1, 1, ]))
```

---

gaussian_approx	<i>Gaussian Approximation of Non-Gaussian/Non-linear State Space Model</i>
-----------------	----------------------------------------------------------------------------

---

### Description

Returns the approximating Gaussian model which has the same conditional mode of  $p(\text{alphaly}, \text{theta})$  as the original model. This function is rarely needed itself, and is mainly available for testing and debugging purposes.

### Usage

```
gaussian_approx(model, max_iter, conv_tol, ...)

## S3 method for class 'nongaussian'
gaussian_approx(model, max_iter = 100, conv_tol = 1e-08, ...)

## S3 method for class 'ssm_nlg'
gaussian_approx(model, max_iter = 100, conv_tol = 1e-08, iekf_iter = 0, ...)
```

### Arguments

model	Model to be approximated. Should be of class <code>bsm_ng</code> , <code>ar1_ng_svm</code> , <code>ssm_ung</code> , or <code>ssm_mng</code> , or <code>ssm_nlg</code> , i.e. non-gaussian or non-linear <code>bssm_model</code> .
max_iter	Maximum number of iterations as a positive integer. Default is 100 (although typically only few iterations are needed).
conv_tol	Positive tolerance parameter. Default is 1e-8. Approximation is claimed to be converged when the mean squared difference of the modes of is less than <code>conv_tol</code> .
...	Ignored.
iekf_iter	For non-linear models, non-negative number of iterations in iterated EKF (defaults to 0, i.e. normal EKF). Used only for models of class <code>ssm_nlg</code> .

### Value

Returns linear-Gaussian SSM of class `ssm_ulg` or `ssm_mlg` which has the same conditional mode of  $p(\text{alphaly}, \text{theta})$  as the original model.

### References

Koopman, SJ and Durbin J (2012). Time Series Analysis by State Space Methods. Second edition. Oxford: Oxford University Press.

Vihola, M, Helske, J, Franks, J. (2020). Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. Scand J Statist. 1-38. <https://doi.org/10.1111/sjos.12492>

**Examples**

```

data("poisson_series")
model <- bsm_ng(y = poisson_series, sd_slope = 0.01, sd_level = 0.1,
  distribution = "poisson")
out <- gaussian_approx(model)
for(i in 1:7)
  cat("Number of iterations used: ", i, ", y[1] = ",
    gaussian_approx(model, max_iter = i, conv_tol = 0)$y[1], "\n", sep = "")

```

---

importance\_sample      *Importance Sampling from non-Gaussian State Space Model*

---

**Description**

Returns `nsim` samples from the approximating Gaussian model with corresponding (scaled) importance weights. Probably mostly useful for comparing KFAS and bssm packages.

**Usage**

```
importance_sample(model, nsim, use_antithetic, max_iter, conv_tol, seed, ...)
```

```

## S3 method for class 'nongaussian'
importance_sample(
  model,
  nsim,
  use_antithetic = TRUE,
  max_iter = 100,
  conv_tol = 1e-08,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

```

**Arguments**

<code>model</code>	Model of class <code>bsm_ng</code> , <code>ar1_ng_svm</code> , <code>ssm_ung</code> , or <code>ssm_mng</code> .
<code>nsim</code>	Number of samples (positive integer).
<code>use_antithetic</code>	Logical. If TRUE (default), use antithetic variable for location in simulation smoothing. Ignored for <code>ssm_mng</code> models.
<code>max_iter</code>	Maximum number of iterations as a positive integer. Default is 100 (although typically only few iterations are needed).
<code>conv_tol</code>	Positive tolerance parameter. Default is 1e-8. Approximation is claimed to be converged when the mean squared difference of the modes of is less than <code>conv_tol</code> .
<code>seed</code>	Seed for the random number generator (positive integer).
<code>...</code>	Ignored.

**Examples**

```

data("sexratio", package = "KFAS")
model <- bsm_ng(sexratio[, "Male"], sd_level = 0.001,
  u = sexratio[, "Total"],
  distribution = "binomial")

imp <- importance_sample(model, nsim = 1000)

est <- matrix(NA, 3, nrow(sexratio))
for(i in 1:ncol(est)) {
  est[, i] <- Hmisc::wtd.quantile(exp(imp$alpha[i, 1, ]), imp$weights,
    prob = c(0.05,0.5,0.95), normwt=TRUE)
}

ts.plot(t(est),lty = c(2,1,2))

```

---

kfilter

*Kalman Filtering*


---

**Description**

Function `kfilter` runs the Kalman filter for the given model, and returns the filtered estimates and one-step-ahead predictions of the states  $\alpha_t$  given the data up to time  $t$ .

**Usage**

```

kfilter(model, ...)

## S3 method for class 'gaussian'
kfilter(model, ...)

## S3 method for class 'nongaussian'
kfilter(model, ...)

```

**Arguments**

```

model      Model of class gaussian, nongaussian or ssm_nlg.
...        Ignored.

```

**Details**

For non-Gaussian models, the filtering is based on the approximate Gaussian model.

**Value**

List containing the log-likelihood (approximate in non-Gaussian case), one-step-ahead predictions at and filtered estimates att of states, and the corresponding variances Pt and Ptt.



**See Also**[bootstrap\\_filter](#)**Examples**

```
x <- cumsum(rnorm(20))
y <- x + rnorm(20, sd = 0.1)
model <- bsm_lg(y, sd_level = 1, sd_y = 0.1)
ts.plot(cbind(y, x, kfilter(model)$att), col = 1:3)
```

---

logLik.gaussian	<i>Extract Log-likelihood of a State Space Model of class bssm_model</i>
-----------------	--------------------------------------------------------------------------

---

**Description**

Computes the log-likelihood of a state space model defined by bssm package.

**Usage**

```
## S3 method for class 'gaussian'
logLik(object, ...)

## S3 method for class 'nongaussian'
logLik(
  object,
  particles,
  method = "psi",
  max_iter = 100,
  conv_tol = 1e-08,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

## S3 method for class 'ssm_nlg'
logLik(
  object,
  particles,
  method = "bsf",
  max_iter = 100,
  conv_tol = 1e-08,
  iekf_iter = 0,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

## S3 method for class 'ssm_sde'
logLik(
```

```

    object,
    particles,
    L,
    seed = sample(.Machine$integer.max, size = 1),
    ...
)

```

### Arguments

object	Model of class <code>bssm_model</code> .
...	Ignored.
particles	Number of samples for particle filter. If 0, approximate log-likelihood is returned either based on the Gaussian approximation or EKF, depending on the method argument.
method	Sampling method. For Gaussian and non-Gaussian models with linear dynamics, options are "bsf" (bootstrap particle filter, default for non-linear models) and "psi" ( $\psi$ -APF, the default for other models). For non-linear models option "ekf" uses EKF/IEKF-based particle filter (or just EKF/IEKF approximation in the case of <code>particles = 0</code> ).
max_iter	Maximum number of iterations used in Gaussian approximation, as a positive integer. Default is 100 (although typically only few iterations are needed).
conv_tol	Positive tolerance parameter used in Gaussian approximation. Default is $1e-8$ .
seed	Seed for RNG (non-negative integer).
iekf_iter	Non-negative integer. If zero (default), first approximation for non-linear Gaussian models is obtained from extended Kalman filter. If <code>iekf_iter &gt; 0</code> , iterated extended Kalman filter is used with <code>iekf_iter</code> iterations.
L	Integer defining the discretization level defined as $(2^L)$ .

### References

- Durbin, J., & Koopman, S. (2002). A Simple and Efficient Simulation Smoother for State Space Time Series Analysis. *Biometrika*, 89(3), 603-615.
- Shephard, N., & Pitt, M. (1997). Likelihood Analysis of Non-Gaussian Measurement Time Series. *Biometrika*, 84(3), 653-667.
- Gordon, NJ, Salmond, DJ, Smith, AFM (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings-F*, 140, 107-113.
- Vihola, M, Helske, J, Franks, J. Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scand J Statist.* 2020; 1-38. <https://doi.org/10.1111/sjos.12492>
- Van Der Merwe, R, Doucet, A, De Freitas, N, Wan, EA (2001). The unscented particle filter. In *Advances in neural information processing systems*, p 584-590.
- Jazwinski, A 1970. *Stochastic Processes and Filtering Theory*. Academic Press.
- Kitagawa, G (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5, 1-25.

**See Also**

particle\_smoother

**Examples**

```

model <- ssm_ulg(y = c(1,4,3), Z = 1, H = 1, T = 1, R = 1)
logLik(model)
model <- ssm_ung(y = c(1,4,3), Z = 1, T = 1, R = 0.5, P1 = 2,
  distribution = "poisson")

model2 <- bsm_ng(y = c(1,4,3), sd_level = 0.5, P1 = 2,
  distribution = "poisson")

logLik(model, particles = 0)
logLik(model2, particles = 0)
logLik(model, particles = 10, seed = 1)
logLik(model2, particles = 10, seed = 1)

```

---

particle\_smoother      *Particle Smoothing*

---

**Description**

Function `particle_smoother` performs particle smoothing based on either bootstrap particle filter [1],  $\psi$ -auxiliary particle filter ( $\psi$ -APF) [2], or extended Kalman particle filter [3] (or its iterated version [4]). The smoothing phase is based on the filter-smoother algorithm by [5].

**Usage**

```

particle_smoother(model, particles, ...)

## S3 method for class 'gaussian'
particle_smoother(
  model,
  particles,
  method = "psi",
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

## S3 method for class 'nongaussian'
particle_smoother(
  model,
  particles,
  method = "psi",
  seed = sample(.Machine$integer.max, size = 1),
  max_iter = 100,
)

```

```

    conv_tol = 1e-08,
    ...
)

## S3 method for class 'ssm_nlg'
particle_smoother(
  model,
  particles,
  method = "bsf",
  seed = sample(.Machine$integer.max, size = 1),
  max_iter = 100,
  conv_tol = 1e-08,
  iekf_iter = 0,
  ...
)

## S3 method for class 'ssm_sde'
particle_smoother(
  model,
  particles,
  L,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

```

### Arguments

model	A model object of class <code>bssm_model</code> .
particles	Number of particles as a positive integer.
...	Ignored.
method	Choice of particle filter algorithm. For Gaussian and non-Gaussian models with linear dynamics, options are "bsf" (bootstrap particle filter, default for non-linear models) and "psi" ( $\psi$ -APF, the default for other models), and for non-linear models option "ekf" (extended Kalman particle filter) is also available.
seed	Seed for RNG (non-negative integer).
max_iter	Maximum number of iterations used in Gaussian approximation, as a positive integer. Default is 100 (although typically only few iterations are needed).
conv_tol	Positive tolerance parameter used in Gaussian approximation. Default is 1e-8.
iekf_iter	Non-negative integer. If zero (default), first approximation for non-linear Gaussian models is obtained from extended Kalman filter. If <code>iekf_iter &gt; 0</code> , iterated extended Kalman filter is used with <code>iekf_iter</code> iterations.
L	Positive integer defining the discretization level for SDE model.

### Details

See one of the vignettes for  $\psi$ -APF in case of nonlinear models.

**Value**

List with samples (alpha) from the smoothing distribution and corresponding weights (weights), as well as smoothed means and covariances (alphahat and Vt) of the states and estimated log-likelihood (logLik).

**References**

- [1] Gordon, NJ, Salmond, DJ, Smith, AFM (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings-F, 140, 107-113.
- [2] Vihola, M, Helske, J, Franks, J. Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. Scand J Statist. 2020; 1-38. <https://doi.org/10.1111/sjos.12492>
- [3] Van Der Merwe, R, Doucet, A, De Freitas, N, Wan, EA (2001). The unscented particle filter. In Advances in neural information processing systems, p 584-590.
- [4] Jazwinski, A 1970. Stochastic Processes and Filtering Theory. Academic Press.
- [5] Kitagawa, G (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. Journal of Computational and Graphical Statistics, 5, 1-25.

**Examples**

```
set.seed(1)
x <- cumsum(rnorm(100))
y <- rnorm(100, x)
model <- ssm_ulg(y, Z = 1, T = 1, R = 1, H = 1, P1 = 1)
system.time(out <- particle_smoother(model, particles = 1000))
# same with simulation smoother:
system.time(out2 <- sim_smoother(model, particles = 1000,
  use_antithetic = TRUE))
ts.plot(out$alphahat, rowMeans(out2), col = 1:2)
```

---

poisson\_series

*Simulated Poisson time series data*

---

**Description**

See example for code for reproducing the data.

**Format**

A vector of length 100

**Examples**

```
# The data was generated as follows:
set.seed(321)
slope <- cumsum(c(0, rnorm(99, sd = 0.01)))
y <- rpois(100, exp(cumsum(slope + c(0, rnorm(99, sd = 0.1)))))
```

---

 post\_correct

*Run Post-correction for Approximate MCMC using  $\psi$ -APF*


---

### Description

Function `post_correct` updates previously obtained approximate MCMC output with post-correction weights leading to asymptotically exact weighted posterior, and returns updated MCMC output where components `weights`, `posterior`, `alpha`, `alphahat`, and `Vt` are updated (depending on the original output type).

### Usage

```
post_correct(
  model,
  mcmc_output,
  particles,
  threads = 1L,
  is_type = "is2",
  seed = sample(.Machine$integer.max, size = 1)
)
```

### Arguments

<code>model</code>	Model of class <code>nongaussian</code> or <code>ssm_nlg</code> .
<code>mcmc_output</code>	An output from <code>run_mcmc</code> used to compute the MAP estimate of $\theta$ . While the intended use assumes this is from approximate MCMC, it is not actually checked, i.e., it is also possible to input previous (asymptotically) exact output.
<code>particles</code>	Number of particles for $\psi$ -APF (positive integer).
<code>threads</code>	Number of parallel threads (positive integer, default is 1).
<code>is_type</code>	Type of IS-correction. Possible choices are <code>"is3"</code> for simple importance sampling (weight is computed for each MCMC iteration independently), <code>"is2"</code> for jump chain importance sampling type weighting (default), or <code>"is1"</code> for importance sampling type weighting where the number of particles used for weight computations is proportional to the length of the jump chain block.
<code>seed</code>	Seed for the random number generator (positive integer).

### Value

List with suggested number of particles `N` and matrix containing estimated standard deviations of the log-weights and corresponding number of particles.

### References

A. Doucet, M. K. Pitt, G. Deligiannidis, R. Kohn (2018). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102, 2, 295-313, <https://doi.org/10.1093/biomet/asu075>

Vihola, M, Helske, J, Franks, J (2020). Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scand J Statist.* 1-38. <https://doi.org/10.1111/sjos.12492>

## Examples

```

set.seed(1)
n <- 300
x1 <- sin((2 * pi / 12) * 1:n)
x2 <- cos((2 * pi / 12) * 1:n)
alpha <- numeric(n)
alpha[1] <- 0
rho <- 0.7
sigma <- 2
mu <- 1
for(i in 2:n) {
  alpha[i] <- rnorm(1, mu * (1 - rho) + rho * alpha[i-1], sigma)
}
u <- rpois(n, 50)
y <- rbinom(n, size = u, plogis(0.5 * x1 + x2 + alpha))

ts.plot(y / u)

model <- ar1_ng(y, distribution = "binomial",
  rho = uniform(0.5, -1, 1), sigma = gamma_prior(1, 2, 0.001),
  mu = normal(0, 0, 10),
  xreg = cbind(x1,x2), beta = normal(c(0, 0), 0, 5),
  u = u)

out_approx <- run_mcmc(model, mcmc_type = "approx",
  local_approx = FALSE, iter = 50000)

out_is2 <- post_correct(model, out_approx, particles = 30,
  threads = 2)
out_is2$time

summary(out_approx, return_se = TRUE)
summary(out_is2, return_se = TRUE)

# latent state
library("dplyr")
library("ggplot2")
state_approx <- as.data.frame(out_approx, variable = "states") %>%
  group_by(time) %>%
  summarise(mean = mean(value))

state_exact <- as.data.frame(out_is2, variable = "states") %>%
  group_by(time) %>%
  summarise(mean = weighted.mean(value, weight))

dplyr::bind_rows(approx = state_approx,
  exact = state_exact, .id = "method") %>%
  filter(time > 200) %>%

```

```

ggplot(aes(time, mean, colour = method)) +
  geom_line() +
  theme_bw()

# posterior means
p_approx <- predict(out_approx, model, type = "mean",
  nsim = 1000, future = FALSE) %>%
  group_by(time) %>%
  summarise(mean = mean(value))
p_exact <- predict(out_is2, model, type = "mean",
  nsim = 1000, future = FALSE) %>%
  group_by(time) %>%
  summarise(mean = mean(value))

dplyr::bind_rows(approx = p_approx,
  exact = p_exact, .id = "method") %>%
  filter(time > 200) %>%
  ggplot(aes(time, mean, colour = method)) +
  geom_line() +
  theme_bw()

```

---

predict.mcmc\_output     *Predictions for State Space Models*

---

### Description

Draw samples from the posterior predictive distribution for future time points given the posterior draws of hyperparameters  $\theta$  and latent state  $\alpha_{n+1}$ . Function can also be used to draw samples from the posterior predictive distribution  $p(\tilde{y}_1, \dots, \tilde{y}_n | y_1, \dots, y_n)$ .

### Usage

```

## S3 method for class 'mcmc_output'
predict(
  object,
  model,
  nsim,
  type = "response",
  future = TRUE,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

```

### Arguments

object                    Results object of class mcmc\_output from [run\\_mcmc](#)



model	A <code>bssm_model</code> object.. Should have same structure and class as the original model which was used in <code>run_mcmc</code> , in order to plug the posterior samples of the model parameters to the right places. It is also possible to input the original model for obtaining predictions for past time points. In this case, set argument <code>future</code> to <code>FALSE</code> .
nsim	Positive integer defining number of samples to draw.
type	Type of predictions. Possible choices are "mean" "response", or "state" level.
future	Default is <code>TRUE</code> , in which case predictions are for the future, using posterior samples of $(\theta, \alpha_{T+1})$ i.e. the posterior samples of hyperparameters and latest states. Otherwise it is assumed that <code>model</code> corresponds to the original model.
seed	Seed for RNG (positive integer). Note that this affects only the C++ side, and <code>predict</code> also uses R side RNG for subsampling, so for replicable results you should call <code>set.seed</code> before <code>predict</code> .
...	Ignored.

### Value

A `data.frame` consisting of samples from the predictive posterior distribution.

### Examples

```
library("graphics")
y <- log10(JohnsonJohnson)
prior <- uniform(0.01, 0, 1)
model <- bsm_lg(window(y, end = c(1974, 4)), sd_y = prior,
  sd_level = prior, sd_slope = prior, sd_seasonal = prior)

mcmc_results <- run_mcmc(model, iter = 5000)
future_model <- model
future_model$y <- ts(rep(NA, 25),
  start = tsp(model$y)[2] + 2 * deltat(model$y),
  frequency = frequency(model$y))
# use "state" for illustrative purposes, we could use type = "mean" directly
pred <- predict(mcmc_results, future_model, type = "state",
  nsim = 1000)

library("dplyr")
sumr_fit <- as.data.frame(mcmc_results, variable = "states") %>%
  group_by(time, iter) %>%
  mutate(signal =
    value[variable == "level"] +
    value[variable == "seasonal_1"]) %>%
  group_by(time) %>%
  summarise(mean = mean(signal),
    lwr = quantile(signal, 0.025),
    upr = quantile(signal, 0.975))

sumr_pred <- pred %>%
```

```

group_by(time, sample) %>%
mutate(signal =
  value[variable == "level"] +
  value[variable == "seasonal_1"]) %>%
group_by(time) %>%
summarise(mean = mean(signal),
  lwr = quantile(signal, 0.025),
  upr = quantile(signal, 0.975))

# If we used type = "mean", we could do
# sumr_pred <- pred %>%
#   group_by(time) %>%
#   summarise(mean = mean(value),
#     lwr = quantile(value, 0.025),
#     upr = quantile(value, 0.975))

library("ggplot2")
rbind(sumr_fit, sumr_pred) %>%
  ggplot(aes(x = time, y = mean)) +
  geom_ribbon(aes(ymin = lwr, ymax = upr),
    fill = "#92f0a8", alpha = 0.25) +
  geom_line(colour = "#92f0a8") +
  theme_bw() +
  geom_point(data = data.frame(
    mean = log10(JohnsonJohnson),
    time = time(JohnsonJohnson)))

# Posterior predictions for past observations:
yrep <- predict(mcmc_results, model, type = "response",
  future = FALSE, nsim = 1000)
meanrep <- predict(mcmc_results, model, type = "mean",
  future = FALSE, nsim = 1000)

sumr_yrep <- yrep %>%
  group_by(time) %>%
  summarise(earnings = mean(value),
    lwr = quantile(value, 0.025),
    upr = quantile(value, 0.975)) %>%
  mutate(interval = "Observations")

sumr_meanrep <- meanrep %>%
  group_by(time) %>%
  summarise(earnings = mean(value),
    lwr = quantile(value, 0.025),
    upr = quantile(value, 0.975)) %>%
  mutate(interval = "Mean")

rbind(sumr_meanrep, sumr_yrep) %>%
  mutate(interval = factor(interval, levels = c("Observations", "Mean"))) %>%
  ggplot(aes(x = time, y = earnings)) +
  geom_ribbon(aes(ymin = lwr, ymax = upr, fill = interval),
    alpha = 0.75) +
  theme_bw() +

```

```
geom_point(data = data.frame(
  earnings = model$y,
  time = time(model$y)))
```

---

```
print.mcmc_output      Print Results from MCMC Run
```

---

### Description

Prints some basic summaries from the MCMC run by `run_mcmc`.

### Usage

```
## S3 method for class 'mcmc_output'
print(x, ...)
```

### Arguments

x	Output from <code>run_mcmc</code> .
...	Ignored.

---

```
run_mcmc              Bayesian Inference of State Space Models
```

---

### Description

Adaptive Markov chain Monte Carlo simulation for SSMs using Robust Adaptive Metropolis algorithm by Vihola (2012). Several different MCMC sampling schemes are implemented, see parameter arguments, package vignette, Vihola, Helske, Franks (2020) and Helske and Vihola (2021) for details.

### Usage

```
run_mcmc(model, ...)

## S3 method for class 'gaussian'
run_mcmc(
  model,
  iter,
  output_type = "full",
  burnin = floor(iter/2),
  thin = 1,
  gamma = 2/3,
  target_acceptance = 0.234,
```

```
S,  
end_adaptive_phase = FALSE,  
threads = 1,  
seed = sample(.Machine$integer.max, size = 1),  
...  
)  
  
## S3 method for class 'nongaussian'  
run_mcmc(  
  model,  
  iter,  
  particles,  
  output_type = "full",  
  mcmc_type = "is2",  
  sampling_method = "psi",  
  burnin = floor(iter/2),  
  thin = 1,  
  gamma = 2/3,  
  target_acceptance = 0.234,  
  S,  
  end_adaptive_phase = FALSE,  
  local_approx = TRUE,  
  threads = 1,  
  seed = sample(.Machine$integer.max, size = 1),  
  max_iter = 100,  
  conv_tol = 1e-08,  
  ...  
)  
  
## S3 method for class 'ssm_nlg'  
run_mcmc(  
  model,  
  iter,  
  particles,  
  output_type = "full",  
  mcmc_type = "is2",  
  sampling_method = "bsf",  
  burnin = floor(iter/2),  
  thin = 1,  
  gamma = 2/3,  
  target_acceptance = 0.234,  
  S,  
  end_adaptive_phase = FALSE,  
  threads = 1,  
  seed = sample(.Machine$integer.max, size = 1),  
  max_iter = 100,  
  conv_tol = 1e-08,  
  iekf_iter = 0,
```

```

    ...
)

## S3 method for class 'ssm_sde'
run_mcmc(
  model,
  iter,
  particles,
  output_type = "full",
  mcmc_type = "is2",
  L_c,
  L_f,
  burnin = floor(iter/2),
  thin = 1,
  gamma = 2/3,
  target_acceptance = 0.234,
  S,
  end_adaptive_phase = FALSE,
  threads = 1,
  seed = sample(.Machine$integer.max, size = 1),
  ...
)

```

### Arguments

model	Model of class <code>bssm_model</code> .
...	Ignored.
iter	Positive integer defining the total number of MCMC iterations.
output_type	Either "full" (default, returns posterior samples from the posterior $p(\alpha, \theta y)$ ), "theta" (for marginal posterior of theta), or "summary" (return the mean and variance estimates of the states and posterior samples of theta). See details.
burnin	Positive integer defining the length of the burn-in period which is disregarded from the results. Defaults to <code>iter / 2</code> . Note that all MCMC algorithms of <code>bssm</code> use adaptive MCMC during the burn-in period in order to find good proposal distribution.
thin	Positive integer defining the thinning rate. All MCMC algorithms in <code>bssm</code> use the jump chain representation (see ref [2]), and the thinning is applied to these blocks. Defaults to 1. For IS-corrected methods, larger value can also be statistically more effective. Note: With <code>output_type = "summary"</code> , the thinning does not affect the computations of the summary statistics in case of pseudo-marginal methods.
gamma	Tuning parameter for the adaptation of RAM algorithm. Must be between 0 and 1.
target_acceptance	Target acceptance rate for MCMC. Defaults to 0.234. Must be between 0 and 1.
S	Matrix defining the initial value for the lower triangular matrix of the RAM algorithm, so that the covariance matrix of the Gaussian proposal distribution is $SS'$ .

Note that for some parameters (currently the standard deviation, dispersion, and autoregressive parameters of the BSM and AR(1) models) the sampling is done in unconstrained parameter space, i.e.  $\text{internal\_theta} = \log(\text{theta})$  (and  $\text{logit}(\text{rho})$  or AR coefficient).

end_adaptive_phase	Logical, if TRUE, S is held fixed after the burnin period. Default is FALSE.
threads	Number of threads for state simulation. Positive integer (default is 1). Note that parallel computing is only used in the post-correction phase of IS-MCMC and when sampling the states in case of (approximate) Gaussian models.
seed	Seed for the random number generator (positive integer).
particles	Number of state samples per MCMC iteration for models other than linear-Gaussian models. Ignored if mcmc_type is "approx" or "ekf".
mcmc_type	What type of MCMC algorithm should be used for models other than linear-Gaussian models? Possible choices are "pm" for pseudo-marginal MCMC, "da" for delayed acceptance version of PMCMC, "approx" for approximate inference based on the Gaussian approximation of the model, "ekf" for approximate inference using extended Kalman filter (for ssm_nlg), or one of the three importance sampling type weighting schemes: "is3" for simple importance sampling (weight is computed for each MCMC iteration independently), "is2" for jump chain importance sampling type weighting (default), or "is1" for importance sampling type weighting where the number of particles used for weight computations is proportional to the length of the jump chain block.
sampling_method	Method for state sampling when for models other than linear-Gaussian models. If "psi", $\psi$ -APF is used (default). If "spdk", non-sequential importance sampling based on Gaussian approximation is used. If "bsf", bootstrap filter is used. If "ekf", particle filter based on EKF-proposals are used (only for ssm_nlg models).
local_approx	If TRUE (default), Gaussian approximation needed for some of the methods is performed at each iteration. If FALSE, approximation is updated only once at the start of the MCMC using the initial model.
max_iter	Maximum number of iterations used in Gaussian approximation, as a positive integer. Default is 100 (although typically only few iterations are needed).
conv_tol	Positive tolerance parameter used in Gaussian approximation.
iekf_iter	Non-negative integer. The default zero corresponds to normal EKF, whereas $\text{iekf\_iter} > 0$ corresponds to iterated EKF with $\text{iekf\_iter}$ iterations. Used only for models of class ssm_nlg.
L_c, L_f	For ssm_sde models, Positive integer values defining the discretization levels for first and second stages (defined as $2^L$ ). For pseudo-marginal methods ("pm"), maximum of these is used.

## Details

For linear-Gaussian models, option "summary" does not simulate states directly but computes the posterior means and variances of states using fast Kalman smoothing. This is slightly faster, more memory efficient and more accurate than calculations based on simulation smoother. In other cases,

the means and covariances are computed using the full output of particle filter instead of subsampling one of these as in case of `output_type = "full"`.

## References

- [1] Vihola M (2012). Robust adaptive Metropolis algorithm with coerced acceptance rate. *Statistics and Computing*, 22(5), p 997-1008.
  - [2] Vihola, M, Helske, J, Franks, J (2020). Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scand J Statist.* 1-38. <https://doi.org/10.1111/sjos.12492>
  - [3] Helske, J, Vihola, M (2019). bssm: Bayesian Inference of Non-linear and Non-Gaussian State Space Models in R. Arxiv preprint 2101.08492.
- Vihola, M, Helske, J, Franks, J. Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scand J Statist.* 2020; 1-38. <https://doi.org/10.1111/sjos.12492>

## Examples

```
model <- ar1_lg(LakeHuron, rho = uniform(0.5,-1,1),
  sigma = halfnormal(1, 10), mu = normal(500, 500, 500),
  sd_y = halfnormal(1, 10))

mcmc_results <- run_mcmc(model, iter = 2e4)
summary(mcmc_results, return_se = TRUE)

library("dplyr")
sumr <- as.data.frame(mcmc_results, variable = "states") %>%
  group_by(time) %>%
  summarise(mean = mean(value),
    lwr = quantile(value, 0.025),
    upr = quantile(value, 0.975))
library("ggplot2")
sumr %>% ggplot(aes(time, mean)) +
  geom_ribbon(aes(ymin = lwr, ymax = upr),alpha=0.25) +
  geom_line() + theme_bw() +
  geom_point(data = data.frame(mean = LakeHuron, time = time(LakeHuron)),
    col = 2)
set.seed(1)
n <- 50
slope <- cumsum(c(0, rnorm(n - 1, sd = 0.001)))
level <- cumsum(slope + c(0, rnorm(n - 1, sd = 0.2)))
y <- rpois(n, exp(level))
poisson_model <- bsm_ng(y,
  sd_level = halfnormal(0.01, 1),
  sd_slope = halfnormal(0.01, 0.1),
  P1 = diag(c(10, 0.1)), distribution = "poisson")

# Note small number of iterations for CRAN checks
mcmc_out <- run_mcmc(poisson_model, iter = 1000, particles = 10,
  mcmc_type = "da")
summary(mcmc_out, what = "theta", return_se = TRUE)

set.seed(123)
```

```

n <- 50
sd_level <- 0.1
drift <- 0.01
beta <- -0.9
phi <- 5

level <- cumsum(c(5, drift + rnorm(n - 1, sd = sd_level)))
x <- 3 + (1:n) * drift + sin(1:n + runif(n, -1, 1))
y <- rbinom(n, size = phi, mu = exp(beta * x + level))

model <- bsm_ng(y, xreg = x,
  beta = normal(0, 0, 10),
  phi = halfnormal(1, 10),
  sd_level = halfnormal(0.1, 1),
  sd_slope = halfnormal(0.01, 0.1),
  a1 = c(0, 0), P1 = diag(c(10, 0.1)^2),
  distribution = "negative binomial")

# run IS-MCMC
# Note small number of iterations for CRAN checks
fit <- run_mcmc(model, iter = 5000,
  particles = 10, mcmc_type = "is2", seed = 1)

# extract states
d_states <- as.data.frame(fit, variable = "states", time = 1:n)

library("dplyr")
library("ggplot2")

# compute summary statistics
level_sumr <- d_states %>%
  filter(variable == "level") %>%
  group_by(time) %>%
  summarise(mean = Hmisc::wtd.mean(value, weight, normwt = TRUE),
    lwr = Hmisc::wtd.quantile(value, weight,
    0.025, normwt = TRUE),
    upr = Hmisc::wtd.quantile(value, weight,
    0.975, normwt = TRUE))

# visualize
level_sumr %>% ggplot(aes(x = time, y = mean)) +
  geom_line() +
  geom_line(aes(y = lwr), linetype = "dashed", na.rm = TRUE) +
  geom_line(aes(y = upr), linetype = "dashed", na.rm = TRUE) +
  theme_bw() +
  theme(legend.title = element_blank()) +
  xlab("Time") + ylab("Level")

# theta
d_theta <- as.data.frame(fit, variable = "theta")
ggplot(d_theta, aes(x = value)) +
  geom_density(aes(weight = weight), adjust = 2, fill = "#92f0a8") +
  facet_wrap(~ variable, scales = "free") +

```



```

theme_bw()

# Bivariate Poisson model:

set.seed(1)
x <- cumsum(c(3, rnorm(19, sd = 0.5)))
y <- cbind(
  rpois(20, exp(x)),
  rpois(20, exp(x)))

prior_fn <- function(theta) {
  # half-normal prior using transformation
  dnorm(exp(theta), 0, 1, log = TRUE) + theta # plus jacobian term
}

update_fn <- function(theta) {
  list(R = array(exp(theta), c(1, 1, 1)))
}

model <- ssm_mng(y = y, Z = matrix(1,2,1), T = 1,
  R = 0.1, P1 = 1, distribution = "poisson",
  init_theta = log(0.1),
  prior_fn = prior_fn, update_fn = update_fn)

# Note small number of iterations for CRAN checks
out <- run_mcmc(model, iter = 5000, mcmc_type = "approx")

sumr <- as.data.frame(out, variable = "states") %>%
  group_by(time) %>% mutate(value = exp(value)) %>%
  summarise(mean = mean(value),
    ymin = quantile(value, 0.05), ymax = quantile(value, 0.95))
ggplot(sumr, aes(time, mean)) +
  geom_ribbon(aes(ymin = ymin, ymax = ymax), alpha = 0.25) +
  geom_line() +
  geom_line(data = data.frame(mean = y[, 1], time = 1:20),
    colour = "tomato") +
  geom_line(data = data.frame(mean = y[, 2], time = 1:20),
    colour = "tomato") +
  theme_bw()

```

**Description**

Function `sim_smoother` performs simulation smoothing i.e. simulates the states from the conditional distribution  $p(\alpha|y, \theta)$  for linear-Gaussian models.

**Usage**

```

sim_smoother(model, nsim, seed, use_antithetic = TRUE, ...)

## S3 method for class 'gaussian'
sim_smoother(
  model,
  nsim = 1,
  seed = sample(.Machine$integer.max, size = 1),
  use_antithetic = TRUE,
  ...
)

## S3 method for class 'nongaussian'
sim_smoother(
  model,
  nsim = 1,
  seed = sample(.Machine$integer.max, size = 1),
  use_antithetic = TRUE,
  ...
)

```

**Arguments**

model	Model of class <code>bsm_lg</code> , <code>ar1_lg</code> , <code>ssm_ulg</code> , or <code>ssm_mlg</code> , or one of the non-gaussian models <code>bsm_ng</code> , <code>ar1_ng</code> , <code>svm</code> , <code>ssm_ung</code> , or <code>ssm_mng</code> .
nsim	Number of samples (positive integer).
seed	Seed for the random number generator (positive integer).
use_antithetic	Logical. If TRUE (default), use antithetic variable for location in simulation smoothing. Ignored for <code>ssm_mng</code> models.
...	Ignored.

**Details**

For non-Gaussian/non-linear models, the simulation is based on the approximating Gaussian model.

**Value**

An array containing the generated samples.

**Examples**

```

# only missing data, simulates from prior
model <- bsm_lg(rep(NA, 25), sd_level = 1,
  sd_y = 1)
# use antithetic variable for location
sim <- sim_smoother(model, nsim = 4, use_antithetic = TRUE, seed = 1)
ts.plot(sim[, 1, ])
cor(sim[, 1, ])

```

**Description**

Construct an object of class `ssm_mlg` by directly defining the corresponding terms of the model.

**Usage**

```
ssm_mlg(
  y,
  Z,
  H,
  T,
  R,
  a1 = NULL,
  P1 = NULL,
  init_theta = numeric(0),
  D = NULL,
  C = NULL,
  state_names,
  update_fn = default_update_fn,
  prior_fn = default_prior_fn
)
```

**Arguments**

<code>y</code>	Observations as multivariate time series or matrix with dimensions $n \times p$ .
<code>Z</code>	System matrix $Z$ of the observation equation as $p \times m$ matrix or $p \times m \times n$ array.
<code>H</code>	Lower triangular matrix $H$ of the observation. Either a scalar or a vector of length $n$ .
<code>T</code>	System matrix $T$ of the state equation. Either a $m \times m$ matrix or a $m \times m \times n$ array.
<code>R</code>	Lower triangular matrix $R$ the state equation. Either a $m \times k$ matrix or a $m \times k \times n$ array.
<code>a1</code>	Prior mean for the initial state as a vector of length $m$ .
<code>P1</code>	Prior covariance matrix for the initial state as $m \times m$ matrix.
<code>init_theta</code>	Initial values for the unknown hyperparameters $\theta$ .
<code>D</code>	Intercept terms for observation equation, given as a $p \times n$ matrix.
<code>C</code>	Intercept terms for state equation, given as $m \times n$ matrix.
<code>state_names</code>	Names for the states.

update_fn	Function which returns list of updated model components given input vector theta. This function should take only one vector argument which is used to create list with elements named as Z, H, T, R, a1, P1, D, and C, where each element matches the dimensions of the original model. It's best to check the internal dimensions with <code>str(model_object)</code> as the dimensions of input arguments can differ from the final dimensions. If any of these components is missing, it is assumed to be constant wrt. theta.
prior_fn	Function which returns log of prior density given input vector theta.

### Details

The general multivariate linear-Gaussian model is defined using the following observational and state equations:

$$y_t = D_t + Z_t\alpha_t + H_t\epsilon_t, \text{ (observation equation)}$$

$$\alpha_{t+1} = C_t + T_t\alpha_t + R_t\eta_t, \text{ (transition equation)}$$

where  $\epsilon_t \sim N(0, I_p)$ ,  $\eta_t \sim N(0, I_k)$  and  $\alpha_1 \sim N(a_1, P_1)$  independently of each other. Here  $p$  is the number of time series and  $k$  is the number of disturbance terms (which can be less than  $m$ , the number of states).

The `update_fn` function should take only one vector argument which is used to create list with elements named as Z, H, T, R, a1, P1, D, and C, where each element matches the dimensions of the original model. If any of these components is missing, it is assumed to be constant wrt. theta. Note that while you can input say R as  $m \times k$  matrix for `ssm_mlg`, `update_fn` should return R as  $m \times k \times 1$  in this case. It might be useful to first construct the model without updating function

### Value

Object of class `ssm_mlg`.

### Examples

```
data("GlobalTemp", package = "KFAS")
model_temp <- ssm_mlg(GlobalTemp, H = matrix(c(0.15,0.05,0, 0.05), 2, 2),
  R = 0.05, Z = matrix(1, 2, 1), T = 1, P1 = 10,
  state_names = "temperature",
  # using default values, but being explicit for testing purposes
  D = matrix(0, 2, 1), C = matrix(0, 1, 1))
ts.plot(cbind(model_temp$y, smoother(model_temp)$alphahat), col = 1:3)
```

---

ssm\_mng

*General Non-Gaussian State Space Model*


---

### Description

Construct an object of class `ssm_mng` by directly defining the corresponding terms of the model.

### Usage

```
ssm_mng(
  y,
  Z,
  T,
  R,
  a1 = NULL,
  P1 = NULL,
  distribution,
  phi = 1,
  u,
  init_theta = numeric(0),
  D = NULL,
  C = NULL,
  state_names,
  update_fn = default_update_fn,
  prior_fn = default_prior_fn
)
```

### Arguments

<code>y</code>	Observations as multivariate time series or matrix with dimensions $n \times p$ .
<code>Z</code>	System matrix $Z$ of the observation equation as $p \times m$ matrix or $p \times m \times n$ array.
<code>T</code>	System matrix $T$ of the state equation. Either a $m \times m$ matrix or a $m \times m \times n$ array.
<code>R</code>	Lower triangular matrix $R$ the state equation. Either a $m \times k$ matrix or a $m \times k \times n$ array.
<code>a1</code>	Prior mean for the initial state as a vector of length $m$ .
<code>P1</code>	Prior covariance matrix for the initial state as $m \times m$ matrix.
<code>distribution</code>	vector of distributions of the observed series. Possible choices are "poisson", "binomial", "negative binomial", "gamma", and "gaussian".
<code>phi</code>	Additional parameters relating to the non-Gaussian distributions. For negative binomial distribution this is the dispersion term, for gamma distribution this is the shape parameter, for Gaussian this is standard deviation, and for other distributions this is ignored.

<code>u</code>	Matrix of positive constants for non-Gaussian models (of same dimensions as $y$ ). For Poisson, gamma, and negative binomial distribution, this corresponds to the offset term. For binomial, this is the number of trials.
<code>init_theta</code>	Initial values for the unknown hyperparameters $\theta$ .
<code>D</code>	Intercept terms for observation equation, given as $p \times n$ matrix.
<code>C</code>	Intercept terms for state equation, given as $m \times n$ matrix.
<code>state_names</code>	Names for the states.
<code>update_fn</code>	Function which returns list of updated model components given input vector $\theta$ . This function should take only one vector argument which is used to create list with elements named as $Z$ , $T$ , $R$ , $a_1$ , $P_1$ , $D$ , $C$ , and $\phi$ , where each element matches the dimensions of the original model. If any of these components is missing, it is assumed to be constant wrt. $\theta$ . It's best to check the internal dimensions with <code>str(model_object)</code> as the dimensions of input arguments can differ from the final dimensions.
<code>prior_fn</code>	Function which returns log of prior density given input vector $\theta$ .

### Details

The general multivariate non-Gaussian model is defined using the following observational and state equations:

$$p^i(y_t^i | D_t + Z_t \alpha_t), \text{ (observation equation)}$$

$$\alpha_{t+1} = C_t + T_t \alpha_t + R_t \eta_t, \text{ (transition equation)}$$

where  $\eta_t \sim N(0, I_k)$  and  $\alpha_1 \sim N(a_1, P_1)$  independently of each other, and  $p^i(y_t | \cdot)$  is either Poisson, binomial, gamma, Gaussian, or negative binomial distribution for each observation series  $i = 1, \dots, p$ . Here  $k$  is the number of disturbance terms (which can be less than  $m$ , the number of states).

### Value

Object of class `ssm_mng`.

### Examples

```
set.seed(1)
n <- 20
x <- cumsum(rnorm(n, sd = 0.5))
phi <- 2
y <- cbind(
  rgamma(n, shape = phi, scale = exp(x) / phi),
  rbinom(n, 10, plogis(x)))

Z <- matrix(1, 2, 1)
T <- 1
R <- 0.5
```

```

a1 <- 0
P1 <- 1

update_fn <- function(theta) {
  list(R = array(theta[1], c(1, 1, 1)), phi = c(theta[2], 1))
}

prior_fn <- function(theta) {
  ifelse(all(theta > 0), sum(dnorm(theta, 0, 1, log = TRUE)), -Inf)
}

model <- ssm_mng(y, Z, T, R, a1, P1, phi = c(2, 1),
  init_theta = c(0.5, 2),
  distribution = c("gamma", "binomial"),
  u = cbind(1, rep(10, n)),
  update_fn = update_fn, prior_fn = prior_fn,
  state_names = "random_walk",
  # using default values, but being explicit for testing purposes
  D = matrix(0, 2, 1), C = matrix(0, 1, 1))

# smoothing based on approximating gaussian model
ts.plot(cbind(y, fast_smoother(model)),
  col = 1:3, lty = c(1, 1, 2))

```

---

ssm\_nlg

*General multivariate nonlinear Gaussian state space models*


---

### Description

Constructs an object of class `ssm_nlg` by defining the corresponding terms of the observation and state equation.

### Usage

```

ssm_nlg(
  y,
  Z,
  H,
  T,
  R,
  Z_gn,
  T_gn,
  a1,
  P1,
  theta,
  known_params = NA,
  known_tv_params = matrix(NA),
  n_states,

```

```

    n_etas,
    log_prior_pdf,
    time_varying = rep(TRUE, 4),
    state_names = paste0("state", 1:n_states)
  )

```

### Arguments

<code>y</code>	Observations as multivariate time series (or matrix) of length $n$ .
<code>Z, H, T, R</code>	An external pointers (object of class <code>externalptr</code> ) for the C++ functions which define the corresponding model functions.
<code>Z_gn, T_gn</code>	An external pointers (object of class <code>externalptr</code> ) for the C++ functions which define the gradients of the corresponding model functions.
<code>a1</code>	Prior mean for the initial state as object of class <code>externalptr</code>
<code>P1</code>	Prior covariance matrix for the initial state as object of class <code>externalptr</code>
<code>theta</code>	Parameter vector passed to all model functions.
<code>known_params</code>	Vector of known parameters passed to all model functions.
<code>known_tv_params</code>	Matrix of known parameters passed to all model functions.
<code>n_states</code>	Number of states in the model (positive integer).
<code>n_etas</code>	Dimension of the noise term of the transition equation (positive integer).
<code>log_prior_pdf</code>	An external pointer (object of class <code>externalptr</code> ) for the C++ function which computes the log-prior density given <code>theta</code> .
<code>time_varying</code>	Optional logical vector of length 4, denoting whether the values of <code>Z, H, T, and R</code> vary with respect to time variable (given identical states). If used, this can speed up some computations.
<code>state_names</code>	Vector containing names for the states.

### Details

The nonlinear Gaussian model is defined as

$$y_t = Z(t, \alpha_t, \theta) + H(t, \theta)\epsilon_t, \text{ (observation equation)}$$

$$\alpha_{t+1} = T(t, \alpha_t, \theta) + R(t, \theta)\eta_t, \text{ (transition equation)}$$

where  $\epsilon_t \sim N(0, I_p)$ ,  $\eta_t \sim N(0, I_m)$  and  $\alpha_1 \sim N(a_1, P_1)$  independently of each other, and functions  $Z, H, T, R$  can depend on  $\alpha_t$  and parameter vector  $\theta$ .

Compared to other models, these general models need a bit more effort from the user, as you must provide the several small C++ snippets which define the model structure. See examples in the vignette.

### Value

Object of class `ssm_nlg`.



**Examples**

```

# Takes a while on CRAN
set.seed(1)
n <- 50
x <- y <- numeric(n)
y[1] <- rnorm(1, exp(x[1]), 0.1)
for(i in 1:(n-1)) {
  x[i+1] <- rnorm(1, sin(x[i]), 0.1)
  y[i+1] <- rnorm(1, exp(x[i+1]), 0.1)
}

pntrs <- cpp_example_model("nlg_sin_exp")

model_nlg <- ssm_nlg(y = y, a1 = pntrs$a1, P1 = pntrs$P1,
  Z = pntrs$Z_fn, H = pntrs$H_fn, T = pntrs$T_fn, R = pntrs$R_fn,
  Z_gn = pntrs$Z_gn, T_gn = pntrs$T_gn,
  theta = c(log_H = log(0.1), log_R = log(0.1)),
  log_prior_pdf = pntrs$log_prior_pdf,
  n_states = 1, n_etas = 1, state_names = "state")

out <- ekf(model_nlg, iekf_iter = 100)
ts.plot(cbind(x, out$at[1:n], out$att[1:n]), col = 1:3)

```

ssm\_sde

*Univariate state space model with continuous SDE dynamics***Description**

Constructs an object of class `ssm_sde` by defining the functions for the drift, diffusion and derivative of diffusion terms of univariate SDE, as well as the log-density of observation equation. We assume that the observations are measured at integer times (missing values are allowed).

**Usage**

```

ssm_sde(
  y,
  drift,
  diffusion,
  ddiffusion,
  obs_pdf,
  prior_pdf,
  theta,
  x0,
  positive
)

```

**Arguments**

<code>y</code>	Observations as univariate time series (or vector) of length $n$ .
<code>drift, diffusion, ddiffusion</code>	An external pointers for the C++ functions which define the drift, diffusion and derivative of diffusion functions of SDE.
<code>obs_pdf</code>	An external pointer for the C++ function which computes the observational log-density given the the states and parameter vector <code>theta</code> .
<code>prior_pdf</code>	An external pointer for the C++ function which computes the prior log-density given the parameter vector <code>theta</code> .
<code>theta</code>	Parameter vector passed to all model functions.
<code>x0</code>	Fixed initial value for SDE at time 0.
<code>positive</code>	If TRUE, positivity constraint is forced by <code>abs</code> in Milstein scheme.

**Details**

As in case of `ssm_nlg` models, these general models need a bit more effort from the user, as you must provide the several small C++ snippets which define the model structure. See vignettes for an example.

**Value**

Object of class `ssm_sde`.

**Examples**

```
# Takes a while on CRAN
library("sde")
set.seed(1)
# theta_0 = rho = 0.5
# theta_1 = nu = 2
# theta_2 = sigma = 0.3
x <- sde.sim(t0 = 0, T = 50, X0 = 1, N = 50,
            drift = expression(0.5 * (2 - x)),
            sigma = expression(0.3),
            sigma.x = expression(0))
y <- rpois(50, exp(x[-1]))

# source c++ snippets
pntrs <- cpp_example_model("sde_poisson_OU")

sde_model <- ssm_sde(y, pntrs$drift, pntrs$diffusion,
                  pntrs$ddiffusion, pntrs$obs_density, pntrs$prior,
                  c(rho = 0.5, nu = 2, sigma = 0.3), 1, positive = FALSE)

est <- particle_smoother(sde_model, L = 12, particles = 500)

ts.plot(cbind(x, est$alphahat,
             est$alphahat - 2*sqrt(c(est$Vt))),
```

```

est$alphahat + 2*sqrt(c(est$Vt))),
col = c(2, 1, 1, 1), lty = c(1, 1, 2, 2))

# Takes time with finer mesh, parallelization with IS-MCMC helps a lot
out <- run_mcmc(sde_model, L_c = 4, L_f = 8,
  particles = 50, iter = 2e4,
  threads = 4L)

```

---

ssm\_ulg

*General univariate linear-Gaussian state space models*


---

### Description

Construct an object of class `ssm_ulg` by directly defining the corresponding terms of the model.

### Usage

```

ssm_ulg(
  y,
  Z,
  H,
  T,
  R,
  a1 = NULL,
  P1 = NULL,
  init_theta = numeric(0),
  D = NULL,
  C = NULL,
  state_names,
  update_fn = default_update_fn,
  prior_fn = default_prior_fn
)

```

### Arguments

<code>y</code>	Observations as time series (or vector) of length $n$ .
<code>Z</code>	System matrix $Z$ of the observation equation. Either a vector of length $m$ , a $m \times n$ matrix, or object which can be coerced to such.
<code>H</code>	Vector of standard deviations. Either a scalar or a vector of length $n$ .
<code>T</code>	System matrix $T$ of the state equation. Either a $m \times m$ matrix or a $m \times m \times n$ array, or object which can be coerced to such.
<code>R</code>	Lower triangular matrix $R$ the state equation. Either a $m \times k$ matrix or a $m \times k \times n$ array, or object which can be coerced to such.

a1	Prior mean for the initial state as a vector of length m.
P1	Prior covariance matrix for the initial state as m x m matrix.
init_theta	Initial values for the unknown hyperparameters theta.
D	Intercept terms $D_t$ for the observations equation, given as a scalar or vector of length n.
C	Intercept terms $C_t$ for the state equation, given as a m times 1 or m times n matrix.
state_names	Names for the states.
update_fn	Function which returns list of updated model components given input vector theta. This function should take only one vector argument which is used to create list with elements named as Z, H, T, R, a1, P1, D, and C, where each element matches the dimensions of the original model. It's best to check the internal dimensions with <code>str(model_object)</code> as the dimensions of input arguments can differ from the final dimensions. If any of these components is missing, it is assumed to be constant wrt. theta.
prior_fn	Function which returns log of prior density given input vector theta.

### Details

The general univariate linear-Gaussian model is defined using the following observational and state equations:

$$y_t = D_t + Z_t\alpha_t + H_t\epsilon_t, \text{ (observation equation)}$$

$$\alpha_{t+1} = C_t + T_t\alpha_t + R_t\eta_t, \text{ (transition equation)}$$

where  $\epsilon_t \sim N(0, 1)$ ,  $\eta_t \sim N(0, I_k)$  and  $\alpha_1 \sim N(a_1, P_1)$  independently of each other. Here k is the number of disturbance terms which can be less than m, the number of states.

The `update_fn` function should take only one vector argument which is used to create list with elements named as Z, H, T, R, a1, P1, D, and C, where each element matches the dimensions of the original model. If any of these components is missing, it is assumed to be constant wrt. theta. Note that while you can input say R as m x k matrix for `ssm_ulg`, `update_fn` should return R as m x k x 1 in this case. It might be useful to first construct the model without updating function and then check the expected structure of the model components from the output.

### Value

Object of class `ssm_ulg`.

### Examples

```
# Regression model with time-varying coefficients
set.seed(1)
n <- 100
x1 <- rnorm(n)
x2 <- rnorm(n)
```

```

b1 <- 1 + cumsum(rnorm(n, sd = 0.5))
b2 <- 2 + cumsum(rnorm(n, sd = 0.1))
y <- 1 + b1 * x1 + b2 * x2 + rnorm(n, sd = 0.1)

Z <- rbind(1, x1, x2)
H <- 0.1
T <- diag(3)
R <- diag(c(0, 1, 0.1))
a1 <- rep(0, 3)
P1 <- diag(10, 3)

# updates the model given the current values of the parameters
update_fn <- function(theta) {
  R <- diag(c(0, theta[1], theta[2]))
  dim(R) <- c(3, 3, 1)
  list(R = R, H = theta[3])
}
# prior for standard deviations as half-normal(1)
prior_fn <- function(theta) {
  if(any(theta < 0)) {
    log_p <- -Inf
  } else {
    log_p <- sum(dnorm(theta, 0, 1, log = TRUE))
  }
  log_p
}

model <- ssm_ulg(y, Z, H, T, R, a1, P1,
  init_theta = c(1, 0.1, 0.1),
  update_fn = update_fn, prior_fn = prior_fn,
  state_names = c("level", "b1", "b2"),
  # using default values, but being explicit for testing purposes
  C = matrix(0, 3, 1), D = numeric(1))

out <- run_mcmc(model, iter = 10000)
out
sumr <- summary(out, variable = "state")
ts.plot(sumr$Mean, col = 1:3)
lines(b1, col= 2, lty = 2)
lines(b2, col= 3, lty = 2)

# Perhaps easiest way to construct a general SSM for bssm is to use the
# model building functionality of KFAS:
library("KFAS")

model_kfas <- SSMModel(log(drivers) ~ SSMtrend(1, Q = 5e-4)+
  SSMseasonal(period = 12, sea.type = "trigonometric", Q = 0) +
  log(PetrolPrice) + law, data = Seatbelts, H = 0.005)

# use as_bssm function for conversion, kappa defines the
# prior variance for diffuse states
model_bssm <- as_bssm(model_kfas, kappa = 100)

```

```

# define updating function for parameter estimation
# we can use SSMModel and as_bssm functions here as well
# (for large model it is more efficient to do this
# "manually" by constructing only necessary matrices,
# i.e., in this case a list with H and Q)

updatefn <- function(theta) {

  model_kfas <- SSMModel(log(drivers) ~ SSMtrend(1, Q = theta[1]^2)+
    SSMseasonal(period = 12,
      sea.type = "trigonometric", Q = theta[2]^2) +
    log(PetrolPrice) + law, data = Seatbelts, H = theta[3]^2)

  as_bssm(model_kfas, kappa = 100)
}

prior <- function(theta) {
  if(any(theta < 0)) -Inf else sum(dnorm(theta, 0, 0.1, log = TRUE))
}
init_theta <- rep(1e-2, 3)
c("sd_level", "sd_seasonal", "sd_y")
model_bssm <- as_bssm(model_kfas, kappa = 100,
  init_theta = init_theta,
  prior_fn = prior, update_fn = updatefn)

out <- run_mcmc(model_bssm, iter = 10000, burnin = 5000)
out

# Above the regression coefficients are modelled as
# time-invariant latent states.
# Here is an alternative way where we use variable D so that the
# coefficients are part of parameter vector theta:

updatefn2 <- function(theta) {
  # note no PetrolPrice or law variables here
  model_kfas2 <- SSMModel(log(drivers) ~ SSMtrend(1, Q = theta[1]^2)+
    SSMseasonal(period = 12, sea.type = "trigonometric", Q = theta[2]^2),
    data = Seatbelts, H = theta[3]^2)

  X <- model.matrix(~ -1 + law + log(PetrolPrice), data = Seatbelts)
  D <- t(X %*% theta[4:5])
  as_bssm(model_kfas2, D = D, kappa = 100)
}
prior2 <- function(theta) {
  if(any(theta[1:3] < 0)) {
    -Inf
  } else {
    sum(dnorm(theta[1:3], 0, 0.1, log = TRUE)) +
    sum(dnorm(theta[4:5], 0, 10, log = TRUE))
  }
}
init_theta <- c(rep(1e-2, 3), 0, 0)

```

```

names(init_theta) <- c("sd_level", "sd_seasonal", "sd_y", "law", "Petrol")
model_bssm2 <- updatefn2(init_theta)
model_bssm2$theta <- init_theta
model_bssm2$prior_fn <- prior2
model_bssm2$update_fn <- updatefn2

out2 <- run_mcmc(model_bssm2, iter = 10000, burnin = 5000)
out2

```

---

ssm\_ung

*General univariate non-Gaussian state space model*


---

### Description

Construct an object of class `ssm_ung` by directly defining the corresponding terms of the model.

### Usage

```

ssm_ung(
  y,
  Z,
  T,
  R,
  a1 = NULL,
  P1 = NULL,
  distribution,
  phi = 1,
  u,
  init_theta = numeric(0),
  D = NULL,
  C = NULL,
  state_names,
  update_fn = default_update_fn,
  prior_fn = default_prior_fn
)

```

### Arguments

<code>y</code>	Observations as time series (or vector) of length $n$ .
<code>Z</code>	System matrix $Z$ of the observation equation. Either a vector of length $m$ , a $m \times n$ matrix, or object which can be coerced to such.
<code>T</code>	System matrix $T$ of the state equation. Either a $m \times m$ matrix or a $m \times m \times n$ array, or object which can be coerced to such.
<code>R</code>	Lower triangular matrix $R$ the state equation. Either a $m \times k$ matrix or a $m \times k \times n$ array, or object which can be coerced to such.
<code>a1</code>	Prior mean for the initial state as a vector of length $m$ .

P1	Prior covariance matrix for the initial state as m x m matrix.
distribution	Distribution of the observed time series. Possible choices are "poisson", "binomial", "gamma", and "negative binomial".
phi	Additional parameter relating to the non-Gaussian distribution. For negative binomial distribution this is the dispersion term, for gamma distribution this is the shape parameter, and for other distributions this is ignored. Should be an object of class <code>bssm_prior</code> or a positive scalar.
u	Vector of positive constants for non-Gaussian models. For Poisson, gamma, and negative binomial distribution, this corresponds to the offset term. For binomial, this is the number of trials.
init_theta	Initial values for the unknown hyperparameters theta.
D	Intercept terms $D_t$ for the observations equation, given as a scalar or vector of length n.
C	Intercept terms $C_t$ for the state equation, given as a m times 1 or m times n matrix.
state_names	Names for the states.
update_fn	Function which returns list of updated model components given input vector theta. This function should take only one vector argument which is used to create list with elements named as Z, T, R, a1, P1, D, C, and phi, where each element matches the dimensions of the original model. If any of these components is missing, it is assumed to be constant wrt. theta. It's best to check the internal dimensions with <code>str(model_object)</code> as the dimensions of input arguments can differ from the final dimensions.
prior_fn	Function which returns log of prior density given input vector theta.

### Details

The general univariate non-Gaussian model is defined using the following observational and state equations:

$$p(y_t | D_t + Z_t \alpha_t), \text{ (observation equation)}$$

$$\alpha_{t+1} = C_t + T_t \alpha_t + R_t \eta_t, \text{ (transition equation)}$$

where  $\eta_t \sim N(0, I_k)$  and  $\alpha_1 \sim N(a_1, P_1)$  independently of each other, and  $p(y_t | \cdot)$  is either Poisson, binomial, gamma, or negative binomial distribution. Here k is the number of disturbance terms which can be less than m, the number of states.

The `update_fn` function should take only one vector argument which is used to create list with elements named as Z, phi, T, R, a1, P1, D, and C, where each element matches the dimensions of the original model. If any of these components is missing, it is assumed to be constant wrt. theta. Note that while you can input say R as m x k matrix for `ssm_ung`, `update_fn` should return R as m x k x 1 in this case. It might be useful to first construct the model without updating function and then check the expected structure of the model components from the output.

### Value

Object of class `ssm_ung`.



**Examples**

```

data("drownings", package = "bssm")
model <- ssm_ung(drownings[, "deaths"], Z = 1, T = 1, R = 0.2,
  a1 = 0, P1 = 10, distribution = "poisson", u = drownings[, "population"])

# approximate results based on Gaussian approximation
out <- smoother(model)
ts.plot(cbind(model$y / model$u, exp(out$alphahat)), col = 1:2)

```

---

suggest\_N

*Suggest Number of Particles for  $\psi$ -APF Post-correction*


---

**Description**

Function `estimate_N` estimates suitable number particles needed for accurate post-correction of approximate MCMC.

**Usage**

```

suggest_N(
  model,
  theta,
  candidates = seq(10, 100, by = 10),
  replications = 100,
  seed = sample(.Machine$integer.max, size = 1)
)

```

**Arguments**

<code>model</code>	Model of class <code>nongaussian</code> or <code>ssm_nlg</code> .
<code>theta</code>	A vector of <code>theta</code> corresponding to the model, at which point the standard deviation of the log-likelihood is computed. Typically MAP estimate from the (approximate) MCMC run. Can also be an output from <code>run_mcmc</code> which is then used to compute the MAP estimate of <code>theta</code> .
<code>candidates</code>	Vector of positive integers containing the candidate number of particles to test. Default is <code>seq(10, 100, by = 10)</code> .
<code>replications</code>	Positive integer, how many replications should be used for computing the standard deviations? Default is 100.
<code>seed</code>	Seed for the random number generator (positive integer).

**Details**

Function `suggest_N` estimates the standard deviation of the logarithm of the post-correction weights at approximate MAP of `theta`, using various particle sizes and suggest smallest number of particles which still leads standard deviation less than 1. Similar approach was suggested in the context of pseudo-marginal MCMC by Doucet et al. (2015), but see also Section 10.3 in Vihola et al (2020).

**Value**

List with suggested number of particles  $N$  and matrix containing estimated standard deviations of the log-weights and corresponding number of particles.

**References**

Doucet, A, Pitt, MK, Deligiannidis, G, Kohn, R (2015). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator, *Biometrika*, 102(2) p. 295-313, <https://doi.org/10.1093/biomet/asu075>

Vihola, M, Helske, J, Franks, J (2020). Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scand J Statist.* 1-38. <https://doi.org/10.1111/sjos.12492>

**Examples**

```

set.seed(1)
n <- 300
x1 <- sin((2 * pi / 12) * 1:n)
x2 <- cos((2 * pi / 12) * 1:n)
alpha <- numeric(n)
alpha[1] <- 0
rho <- 0.7
sigma <- 1.2
mu <- 1
for(i in 2:n) {
  alpha[i] <- rnorm(1, mu * (1 - rho) + rho * alpha[i-1], sigma)
}
u <- rpois(n, 50)
y <- rbinom(n, size = u, plogis(0.5 * x1 + x2 + alpha))

ts.plot(y / u)

model <- ar1_ng(y, distribution = "binomial",
  rho = uniform(0.5, -1, 1), sigma = gamma_prior(1, 2, 0.001),
  mu = normal(0, 0, 10),
  xreg = cbind(x1,x2), beta = normal(c(0, 0), 0, 5),
  u = u)

# theta from earlier approximate MCMC run
# out_approx <- run_mcmc(model, mcmc_type = "approx",
#   iter = 5000)
# theta <- out_approx$theta[which.max(out_approx$posterior), ]

theta <- c(rho = 0.64, sigma = 1.16, mu = 1.1, x1 = 0.56, x2 = 1.28)

estN <- suggest_N(model, theta, candidates = seq(10, 50, by = 10),
  replications = 50, seed = 1)
plot(x = estN$results$N, y = estN$results$sd, type = "b")
estN$N

```

---

summary.mcmc\_output     *Summary of MCMC object*

---

### Description

This functions returns a list containing mean, standard deviations, standard errors, and effective sample size estimates for parameters and states.

### Usage

```
## S3 method for class 'mcmc_output'
summary(object, return_se = FALSE, variable = "theta", only_theta = FALSE, ...)
```

### Arguments

object	Output from run_mcmc
return_se	if FALSE (default), computation of standard errors and effective sample sizes is omitted.
variable	Are the summary statistics computed for either "theta" (default), "states", or "both"?
only_theta	Deprecated. If TRUE, summaries are computed only for hyperparameters theta.
...	Ignored.

### Details

For IS-MCMC two types of standard errors are reported. SE-IS can be regarded as the square root of independent IS variance, whereas SE corresponds to the square root of total asymptotic variance (see Remark 3 of Vihola et al. (2020)).

### References

Vihola, M, Helske, J, Franks, J. Importance sampling type estimators based on approximate marginal Markov chain Monte Carlo. *Scand J Statist.* 2020; 1-38. <https://doi.org/10.1111/sjos.12492>

---

svm                             *Stochastic Volatility Model*

---

### Description

Constructs a simple stochastic volatility model with Gaussian errors and first order autoregressive signal. See the main vignette for details.

### Usage

```
svm(y, mu, rho, sd_ar, sigma)
```

**Arguments**

<code>y</code>	Vector or a <code>ts</code> object of observations.
<code>mu</code>	Prior for <code>mu</code> parameter of transition equation. Should be an object of class <code>bssm_prior</code> .
<code>rho</code>	prior for autoregressive coefficient. Should be an object of class <code>bssm_prior</code> .
<code>sd_ar</code>	Prior for the standard deviation of noise of the AR-process. Should be an object of class <code>bssm_prior</code> .
<code>sigma</code>	Prior for <code>sigma</code> parameter of observation equation, internally denoted as <code>phi</code> . Should be an object of class <code>bssm_prior</code> . Ignored if <code>mu</code> is provided. Note that typically parametrization using <code>mu</code> is preferred due to better numerical properties and availability of better Gaussian approximation. Most notably the global approximation approach does not work with <code>sigma</code> parameterization as <code>sigma</code> is not a parameter of the resulting approximate model.

**Value**

Object of class `svm`.

**Examples**

```
data("exchange")
exchange <- exchange[1:100] # faster CRAN check
model <- svm(exchange, rho = uniform(0.98, -0.999, 0.999),
  sd_ar = halfnormal(0.15, 5), sigma = halfnormal(0.6, 2))

obj <- function(pars) {
  -logLik(svm(exchange,
    rho = uniform(pars[1], -0.999, 0.999),
    sd_ar = halfnormal(pars[2], 5),
    sigma = halfnormal(pars[3], 2)), particles = 0)
}
opt <- optim(c(0.98, 0.15, 0.6), obj,
  lower = c(-0.999, 1e-4, 1e-4),
  upper = c(0.999, 10, 10), method = "L-BFGS-B")
pars <- opt$par
model <- svm(exchange,
  rho = uniform(pars[1], -0.999, 0.999),
  sd_ar = halfnormal(pars[2], 5),
  sigma = halfnormal(pars[3], 2))

# alternative parameterization
model2 <- svm(exchange, rho = uniform(0.98, -0.999, 0.999),
  sd_ar = halfnormal(0.15, 5), mu = normal(0, 0, 1))

obj2 <- function(pars) {
  -logLik(svm(exchange,
    rho = uniform(pars[1], -0.999, 0.999),
    sd_ar = halfnormal(pars[2], 5),
    mu = normal(pars[3], 0, 1)), particles = 0)
```

```

}
opt2 <- optim(c(0.98, 0.15, 0), obj2, lower = c(-0.999, 1e-4, -Inf),
  upper = c(0.999, 10, Inf), method = "L-BFGS-B")
pars2 <- opt2$par
model2 <- svm(exchange,
  rho = uniform(pars2[1], -0.999, 0.999),
  sd_ar = halfnormal(pars2[2], 5),
  mu = normal(pars2[3], 0, 1))

# sigma is internally stored in phi
ts.plot(cbind(model$phi * exp(0.5 * fast_smoother(model)),
  exp(0.5 * fast_smoother(model2))), col = 1:2)

```

---

ukf

*Unscented Kalman Filtering*


---

### Description

Function `ukf` runs the unscented Kalman filter for the given non-linear Gaussian model of class `ssm_nlg`, and returns the filtered estimates and one-step-ahead predictions of the states  $\alpha_t$  given the data up to time  $t$ .

### Usage

```
ukf(model, alpha = 0.001, beta = 2, kappa = 0)
```

### Arguments

<code>model</code>	Model of class <code>ssm_nlg</code> .
<code>alpha</code>	Positive tuning parameter of the UKF. Default is 0.001. Smaller the value, closer the sigma point are to the mean of the state.
<code>beta</code>	Non-negative tuning parameter of the UKF. The default value is 2, which is optimal for Gaussian states.
<code>kappa</code>	Non-negative tuning parameter of the UKF, which also affects the spread of sigma points. Default value is 0.

### Value

List containing the log-likelihood, one-step-ahead predictions `at` and filtered estimates `att` of states, and the corresponding variances `Pt` and `Ptt`.

**Examples**

```

# Takes a while on CRAN
set.seed(1)
mu <- -0.2
rho <- 0.7
sigma_y <- 0.1
sigma_x <- 1
x <- numeric(50)
x[1] <- rnorm(1, mu, sigma_x / sqrt(1 - rho^2))
for(i in 2:length(x)) {
  x[i] <- rnorm(1, mu * (1 - rho) + rho * x[i - 1], sigma_x)
}
y <- rnorm(50, exp(x), sigma_y)

pntrs <- cpp_example_model("nlg_ar_exp")

model_nlg <- ssm_nlg(y = y, a1 = pntrs$a1, P1 = pntrs$P1,
  Z = pntrs$Z_fn, H = pntrs$H_fn, T = pntrs$T_fn, R = pntrs$R_fn,
  Z_gn = pntrs$Z_gn, T_gn = pntrs$T_gn,
  theta = c(mu = mu, rho = rho,
    log_sigma_x = log(sigma_x), log_sigma_y = log(sigma_y)),
  log_prior_pdf = pntrs$log_prior_pdf,
  n_states = 1, n_etas = 1, state_names = "state")

out_iekf <- ekf(model_nlg, iekf_iter = 5)
out_ukf <- ukf(model_nlg, alpha = 0.01, beta = 2, kappa = 1)
ts.plot(cbind(x, out_iekf$att, out_ukf$att), col = 1:3)

```

---

uniform\_prior

*Prior objects for bssm models*


---

**Description**

These simple objects of class `bssm_prior` are used to construct a prior distributions for the some of the model objects of `bssm` package. Currently supported priors are uniform (`uniform()`), half-normal (`halfnormal()`), normal (`normal()`), gamma (`gamma`), and truncated normal distribution (`tnormal()`). All parameters are vectorized so for regression coefficient vector `beta` you can define prior for example as `normal(0,0,c(10,20))`.

**Usage**

```
uniform_prior(init, min, max)
```

```
uniform(init, min, max)
```

```
halfnormal_prior(init, sd)
```

```
halfnormal(init, sd)
```

```

normal_prior(init, mean, sd)

normal(init, mean, sd)

tnormal_prior(init, mean, sd, min = -Inf, max = Inf)

tnormal(init, mean, sd, min = -Inf, max = Inf)

gamma_prior(init, shape, rate)

gamma(init, shape, rate)

```

### Arguments

<code>init</code>	Initial value for the parameter, used in initializing the model components and as a starting values in MCMC.
<code>min</code>	Lower bound of the uniform and truncated normal prior.
<code>max</code>	Upper bound of the uniform and truncated normal prior.
<code>sd</code>	Positive value defining the standard deviation of the (underlying i.e. non-truncated) Normal distribution.
<code>mean</code>	Mean of the Normal prior.
<code>shape</code>	Positive shape parameter of the Gamma prior.
<code>rate</code>	Positive rate parameter of the Gamma prior.

### Details

The longer name versions of the prior functions with `_prior` ending are identical with shorter versions and they are available only to avoid clash with R's primitive function `gamma` (other long prior names are just for consistent naming).

### Value

object of class `bssm_prior` or `bssm_prior_list` in case of multiple priors (i.e. multiple regression coefficients).

### Examples

```

# create uniform prior on [-1, 1] for one parameter with initial value 0.2:
uniform(init = 0.2, min = -1.0, max = 1.0)
# two normal priors at once i.e. for coefficients beta:
normal(init = c(0.1, 2.5), mean = 0.1, sd = c(1.5, 2.8))
# Gamma prior
gamma(init = 0.1, shape = 2.5, rate = 1.1)
# Same as
gamma_prior(init = 0.1, shape = 2.5, rate = 1.1)
# Half-normal
halfnormal(init = 0.01, sd = 0.1)

```

```
# Truncated normal  
tnormal(init = 5.2, mean = 5.0, sd = 3.0, min = 0.5, max = 9.5)
```



# Index

- \* **datasets**
  - drownings, 15
  - exchange, 19
  - poisson\_series, 29
- ar1\_lg, 3
- ar1\_ng, 4
- as.data.frame.mcmc\_output, 5
- as\_bssm, 6
- as\_draws.mcmc\_output
  - (as\_draws\_df.mcmc\_output), 7
- as\_draws\_df.mcmc\_output, 7
  
- bootstrap\_filter, 8, 25
- bsm\_lg, 10
- bsm\_ng, 12
- bssm, 14
- bssm\_prior (uniform\_prior), 62
- bssm\_prior\_list (uniform\_prior), 62
  
- cpp\_example\_model, 15
  
- drownings, 15
  
- ekf, 16
- ekf\_fast\_smoother (ekf\_smoother), 17
- ekf\_smoother, 17
- ekpf\_filter, 18
- exchange, 19
- expand\_sample, 20
  
- fast\_smoother, 21
  
- gamma (uniform\_prior), 62
- gamma\_prior (uniform\_prior), 62
- gaussian\_approx, 22
  
- halfnormal (uniform\_prior), 62
- halfnormal\_prior (uniform\_prior), 62
  
- importance\_sample, 23
  
- kfilter, 24
  
- logLik.gaussian, 25
- logLik.nongaussian (logLik.gaussian), 25
- logLik.ssm\_nlg (logLik.gaussian), 25
- logLik.ssm\_sde (logLik.gaussian), 25
  
- normal (uniform\_prior), 62
- normal\_prior (uniform\_prior), 62
  
- particle\_smoother, 27
- poisson\_series, 29
- post\_correct, 30
- predict (predict.mcmc\_output), 32
- predict.mcmc\_output, 32
- print.mcmc\_output, 35
  
- run\_mcmc, 5, 6, 20, 32, 35, 35
  
- sim\_smoother, 41
- smoother (fast\_smoother), 21
- ssm\_mlg, 43
- ssm\_mng, 45
- ssm\_nlg, 47
- ssm\_sde, 49
- ssm\_ulg, 51
- ssm\_ung, 55
- suggest\_N, 57
- summary.mcmc\_output, 59
- svm, 59
  
- tnormal (uniform\_prior), 62
- tnormal\_prior (uniform\_prior), 62
- ts, 60
  
- ukf, 61
- uniform (uniform\_prior), 62
- uniform\_prior, 62