

# Package ‘nmathopencl’

June 21, 2026

**Type** Package

**Title** 'OpenCL'-Ported R 'Mathlib' for GPU-Accelerated Packages

**Version** 0.8.2

**Date** 2026-06-11

**Description** Ships statistical and mathematical routines from R internal 'nmath' ('Mathlib') as 'OpenCL' C sources under directory 'inst/cl/', with R wrappers that use the GPU when 'OpenCL' is available at compile time and fall back to 'stats' equivalents otherwise. Aimed at package developers building custom kernels (for example Bayesian GLMs via suggested package 'glmbytes') using 'opencltools' kernel loaders and related helpers. Contains translated shims, an illustrative GLM-related kernel subsystem, vignettes, and optional GPU acceleration. The ported routines are translated from the 'nmath' ('Mathlib') and 'Rmath' sources of R Core Team (2026) ``R: A Language and Environment for Statistical Computing" <doi:10.32614/R.manuals>. 'OpenCL' GPU execution follows the standard described in Stone, Gohara, and Shi (2010) <doi:10.1109/MCSE.2010.69>. The likelihood subgradient simulation methodology implemented by the illustrative GLM kernel subsystem is described in Nygren and Nygren (2006) <doi:10.1198/016214506000000357>.

**License** GPL (>= 2)

**URL** <https://github.com/knygren/nmathopencl>,  
<https://knygren.r-universe.dev/nmathopencl>

**BugReports** <https://github.com/knygren/nmathopencl/issues>

**Imports** stats, Rcpp (>= 1.1.1), RcppParallel, Rdpack (>= 0.11-0),  
opencltools (>= 0.8.1)

**RdMacros** Rdpack

**LinkingTo** Rcpp, RcppArmadillo, RcppParallel, opencltools

**Depends** MASS, R (>= 3.5.0)

**Suggests** glmbytes (>= 0.9.3), knitr, rmarkdown, testthat (>= 3.0.0),  
spelling

**SystemRequirements** Optional 'OpenCL' support. If available, GPU acceleration will be used; otherwise, computation runs on CPU.

**Encoding** UTF-8

**RoxygenNote** 7.3.3

**VignetteBuilder** knitr

**Config/testthat/edition** 3

**Language** en-US

**NeedsCompilation** yes

**Author** Kjell Nygren [aut, cre],

The R Core Team [ctb, cph] (R 'Mathlib' sources and derived/adapted routines),

The R Foundation [cph] (Portions of R 'Mathlib' and R source code),

Ross Ihaka [ctb, cph] (R 'Mathlib'),

Robert Gentleman [ctb, cph] (Portions of R 'Mathlib'),

Morten Welinder [ctb, cph] (Portions of R 'Mathlib' (pgamma, phyper, ebd0)),

Martin Maechler [ctb] (Portions of R 'Mathlib'),

Catherine Loader [ctb] (Author of the dbinom/bd0/stirlerr density routines in R 'Mathlib' ported here),

Claus Ekstrøm [ctb] (Author of the noncentral t density (dnt) in R 'Mathlib' ported here),

Peter Ruckdeschel [ctb] (Author of the noncentral F density (dnf) in R 'Mathlib' ported here),

Alfred H. Morris, Jr. [ctb] (ACM TOMS 708 incomplete beta code (toms708) ported here),

Armido R. Didonato [ctb] (ACM TOMS 708 incomplete beta code (toms708) ported here),

The Khronos Group Inc [cph] ('OpenCL' API headers in inst/include/CL (Apache License 2.0))

**Maintainer** Kjell Nygren <kjell.a.nygren@gmail.com>

**Repository** CRAN

**Date/Publication** 2026-06-21 15:20:08 UTC

## Contents

nmathopencl-package . . . . .	3
attach_cross_library_tags . . . . .	5
attach_kernel_call_tags . . . . .	8
attach_kernel_dependency_tags . . . . .	10
dbeta_opencl . . . . .	11
dbinom_raw_opencl . . . . .	14
dcauchy_opencl . . . . .	16
dchisq_opencl . . . . .	18
dexp_opencl . . . . .	20
df_opencl . . . . .	22
dgamma_opencl . . . . .	24
dgeom_opencl . . . . .	27

dhyper_opencl . . . . .	28
dlnorm_opencl . . . . .	31
dlogis_opencl . . . . .	33
dnbinom_opencl . . . . .	35
dnorm_opencl . . . . .	37
dpois_raw_opencl . . . . .	40
dt_opencl . . . . .	42
dunif_opencl . . . . .	44
dweibull_opencl . . . . .	46
extract_library_subset . . . . .	48
Ex_EnvelopeEval . . . . .	51
Ex_EnvelopeSize . . . . .	54
Ex_glmbfamfunc . . . . .	56
Ex_glmb_Standardize_Model . . . . .	57
gammafn_opencl . . . . .	58
glmbayesEnvelopeExample . . . . .	60
imax2_opencl . . . . .	60
load_library_for_kernel . . . . .	62
nmathopencl_has_opencl . . . . .	64
norm_rand_opencl . . . . .	66
port_to_opencl_configure . . . . .	67
ptukey_opencl . . . . .	69
rmultinom_opencl . . . . .	70
r_check_user_interrupt_opencl . . . . .	71
r_pow_opencl . . . . .	72
stage_kernel_dependency_sort . . . . .	74
use_opencl_configure . . . . .	75
write_kernel_dependency_index . . . . .	76

<b>Index</b>	<b>79</b>
--------------	-----------

---

nmathopencl-package	<i>nmathopencl: OpenCL-Ported R Math Library for GPU-Accelerated Packages</i>
---------------------	---

---

## Description

**nmathopencl** provides OpenCL-ported versions of R's internal `nmath` and `R_ext` math routines, enabling downstream R packages to build custom GPU-accelerated kernels that call the same statistical distribution functions available in base R. The package is intended as a **developer library**: users install it to gain access to the ported `.c1` source files, then write their own OpenCL kernels that `#include` those sources as needed.

## Details

The core deliverable is a collection of `.cl` files installed under `inst/cl/nmath/` that mirror the R `nmath` library (density, distribution, quantile, and random-variate functions). Downstream packages locate these files at runtime with `system.file("cl", package = "nmathopencl")` and assemble them into an OpenCL program using `opencltools::load_kernel_library(..., package = "nmathopencl")`.

The package also ships `Ex_EnvelopeEval` and its supporting functions (`Ex_glmfamfunc`, `Ex_glmb_Standardize_Model`, `Ex_EnvelopeSize`) as a worked example of how a downstream package—here the **glmbayes** Bayesian GLM sampler—builds a custom kernel on top of the ported `nmath` routines. See `system.file("examples", "Ex_EnvelopeEval.R", package = "nmathopencl")` and `vignette("Chapter-10")` for a complete walkthrough.

Optional GPU acceleration is available wherever an OpenCL runtime is installed. Use `nmathopencl_has_opencl` to query compile-time OpenCL support, `nmathopencl_opencl_fp64_available` / `nmathopencl_opencl_device_info` for double-precision device selection used by kernels. Host/runtime diagnostics use `opencltools::diagnose_glmbayes()`.

The simulation theory underlying the envelope construction is described in (Nygren and Nygren 2006), with implementation details in (Nygren 2025, 2025). OpenCL GPU execution follows (Stone et al. 2010); package vignettes also discuss GPU workflows ((Nygren 2025, 2025)).

## C-callable API (C++ package developers)

GPU `*_opencl` routines are registered for `R_GetCCallable` on load. Downstream packages should `LinkingTo: nmathopencl, Imports: nmathopencl, opencltools, and #include <nmathopencl/nmathopencl_capi.h>` (see `system.file("include/nmathopencl/README.md", package = "nmathopencl")`). Call `nmathopencl_api_version` for ABI compatibility. Kernel loading remains on **opencltools** (`opencltools_capi.h`).

## OpenCL startup checks

In interactive sessions, attaching the package with `library(nmathopencl)` may emit a `packageStartupMessage` comparing compile-time OpenCL support in **nmathopencl** and **opencltools**, noting that CPU fallbacks remain available, and summarizing whether an OpenCL runtime appears available on the host. Messages point to `?gpu_diagnostics`, `vignette("Chapter-01")` (OpenCL enablement), `vignette("Chapter-12")` (packaged `*_opencl` API), and this help page. Host-side OpenCL diagnostics use **opencltools**. Set `options(nmathopencl.quiet_opencl_startup = TRUE)` to suppress these notes (recommended for CI and R CMD check).

## Author(s)

Kjell Nygren

## References

Nygren K (2025). “Chapter A05: Simulation Methods – Likelihood Subgradient Densities.” Vignette in the `glmbayes` R package. R vignette name: `Chapter-A05`.

Nygren K (2025). “Chapter A08: Overview of Envelope Related Functions.” Vignette in the `glmbayes` R package. R vignette name: `Chapter-A08`.

Nygren K (2025). “Chapter 12: Large Models: GPU Acceleration using OpenCL.” Vignette in

the glmbayes R package. R vignette name: Chapter-12.

Nygren K (2025). “Chapter A10: Accelerated EnvelopeBuild Implementation using OpenCL.” Vignette in the glmbayes R package. R vignette name: Chapter-A10.

Nygren K~N, Nygren L~M (2006). “Likelihood Subgradient Densities.” *Journal of the American Statistical Association*, **101**(475), 1144–1156. doi:10.1198/016214506000000357.

Stone JE, Gohara D, Shi G (2010). “OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems.” *Computing in Science & Engineering*, **12**(3), 66–72. doi:10.1109/MCSE.2010.69.

### See Also

Key developer entry points:

- `opencltools::load_kernel_library` — assemble the nmath .cl sources (package = "nmathopencl").
- `nmathopencl_has_opencl` — check whether an OpenCL runtime is present.
- `nmathopencl_opencl_device_info` — which OpenCL device is used for fp64 kernels.
- `Ex_EnvelopeEval` — worked example of a custom kernel built on the ported nmath routines.

Useful links:

- GitHub: <https://github.com/knygren/nmathopencl>
- R-Universe: <https://knygren.r-universe.dev/nmathopencl>
- Related sampler package (Suggests): **glmbayes**

---

attach\_cross\_library\_tags

*Attach Cross-Library Dependency Tags to Kernel Files*

---

### Description

Given a set of user-facing kernel .cl files and a library directory, computes the full transitive dependency list for each kernel (using the library’s pre-built dependency index) and writes the results back into the kernel files as annotation tags.

### Usage

```
attach_cross_library_tags(  
  kernel_paths,  
  library_dir,  
  depends_tag = "depends_nmath",  
  index = NULL,  
  dry_run = FALSE  
)
```

**Arguments**

kernel_paths	Character vector of paths to kernel .cl files.
library_dir	Path to the library directory containing kernel_dependency_index.rds and the library .cl files.
depends_tag	Name of the annotation tag in the kernel files that lists the direct library entry-point stems (e.g. "depends_nmath"). The function reads @{depends_tag} and writes @all_{depends_tag} and @all_{depends_tag}_count.
index	Optional dependency index ( <a href="#">write_kernel_dependency_index</a> , load via readRDS). If NULL, reads file.path(library_dir, "kernel_dependency_index.rds") with message().
dry_run	Logical; if TRUE, compute tags but do not write any files.

**Details**

Cross-library analogue of [attach\\_kernel\\_dependency\\_tags](#): it targets kernels that depend on a library through a @{depends\_tag} tag (entry-point stems), instead of expanding @depends purely inside one library tree.

Typical usage for kernels that call nmath functions:

```

nmath_dir <- system.file(
  "cl", "nmath", package = "opencltools")
idx <- readRDS(file.path(nmath_dir, "kernel_dependency_index.rds"))

attach_cross_library_tags(
  kernel_paths = list.files("inst/cl/src", "\\.*\\.cl$", full.names = TRUE),
  library_dir  = nmath_dir,
  depends_tag  = "depends_nmath",
  index        = idx
)

```

This writes @all\_depends\_nmath\_count and @all\_depends\_nmath into each kernel file that carries a @depends\_nmath annotation.

**Value**

A data frame (returned invisibly) with one row per kernel file and columns:

file	Basename of the kernel file.
direct_stems	Comma-separated direct entry-point stems read from @{depends_tag}.
all_depends_count	Number of library files in the full transitive closure.
all_depends	Comma-separated full transitive dependency list in load order.
changed	TRUE if the file was (or would be, under dry_run) modified.

**See Also**

[attach\\_kernel\\_dependency\\_tags](#) [write\\_kernel\\_dependency\\_index](#) [load\\_library\\_for\\_kernel](#)

**Examples**

```
##### Start of kernel_tagging_workflow example #####

lib_dir <- system.file("cl/ex_glmbyes_nmath", package = "opencltools")
kernels <- list.files(
  system.file("cl/ex_glmbyes_draft_src", package = "opencltools"),
  pattern = "\\\\.cl$", full.names = TRUE
)

# Step 1: scan draft kernels for library calls (read-only dry run)
step1 <- attach_kernel_call_tags(
  kernel_paths = kernels,
  library_dir = lib_dir,
  library_tag = "nmath",
  dry_run = TRUE
)
step1

# Step 2: expand transitive closure (small nmath library; runs on CRAN check)
nmath_small <- system.file("cl/nmath_small", package = "opencltools")
tagged <- system.file("cl/src/dnorm_kernel.cl", package = "opencltools")
idx_small <- write_kernel_dependency_index(library_dir = nmath_small, write = FALSE)

step2_small <- attach_cross_library_tags(
  kernel_paths = tagged,
  library_dir = nmath_small,
  depends_tag = "depends_nmath",
  index = idx_small,
  dry_run = TRUE
)
nrow(step2_small)

# Step 2: full nmath (slow)
nmath_dir <- system.file("cl/nmath", package = "opencltools")
idx <- write_kernel_dependency_index(library_dir = nmath_dir, write = FALSE)

step2 <- attach_cross_library_tags(
  kernel_paths = tagged,
  library_dir = nmath_dir,
  depends_tag = "depends_nmath",
  index = idx,
  dry_run = TRUE
)
step2

#####
## End of kernel_tagging_workflow example
#####
```

---

 attach\_kernel\_call\_tags

*Attach Library Call Tags to Kernel Files*


---

### Description

Scans kernel .cl source files for calls to functions provided by a pre-annotated library, then writes the discovered dependencies as annotation tags directly into the kernel files.

### Usage

```
attach_kernel_call_tags(
    kernel_paths,
    library_dir,
    library_tag,
    overwrite_existing = FALSE,
    dry_run = FALSE
)
```

### Arguments

kernel_paths	Character vector of paths to kernel .cl files.
library_dir	Path to the pre-annotated library directory. Each .cl file in this directory must carry a @provides annotation listing the symbols it exports.
library_tag	String tag suffix used for annotation names, e.g. "nmath". Must not contain spaces or regex special characters.
overwrite_existing	Logical; if FALSE (default), skip files that already carry a @calls_<library_tag> annotation. Set to TRUE to re-scan and overwrite (also clears any existing @all_depends_<tag> so <a href="#">attach_cross_library_tags</a> re-computes it cleanly).
dry_run	Logical; if TRUE, compute tags but do not write any files.

### Details

For a library such as **nmathopencl**, each .cl shard carries a @provides annotation listing the symbols it defines. This function builds a **provides map** from those annotations (symbol → shard stem), scans each kernel's source (with comments and string literals stripped) for matching function calls, and writes four annotation tags at the top of each kernel file:

@library\_deps: <library\_tag> Library name (written if absent).

@calls\_<library\_tag>: sym1, sym2 Symbols from the library actually called in this kernel.

@depends\_<library\_tag>: stem1, stem2 Library shard stems that define the called symbols.

@calls\_opencl\_builtin: sym | (none) Detected OpenCL work-item and synchronization builtins (standard math builtins excluded).

### Two-step tagging workflow:

```

# Step 1 - this function: infer direct library calls from source
nmath_dir <- system.file("cl/nmath", package = "nmathopencl")
attach_kernel_call_tags(
  kernel_paths = list.files("inst/cl/src", "\\\\.cl$", full.names = TRUE),
  library_dir = nmath_dir,
  library_tag = "nmath"
)

# Step 2 - expand to full transitive closure
attach_cross_library_tags(
  kernel_paths = list.files("inst/cl/src", "\\\\.cl$", full.names = TRUE),
  library_dir = nmath_dir,
  depends_tag = "depends_nmath"
)

```

**Value**

A data frame (returned invisibly) with one row per kernel file and columns:

file Basename of the kernel file.

calls Comma-separated library symbols detected in source.

depends Comma-separated shard stems for the detected symbols.

opencl\_builtins Comma-separated OpenCL builtins detected, or empty string if none.

changed TRUE if the file was (or would be) modified. NA means the file was skipped (already tagged).

**See Also**

[attach\\_cross\\_library\\_tags](#), [attach\\_kernel\\_dependency\\_tags](#)

**Examples**

```

##### Start of kernel_tagging_workflow example #####

lib_dir <- system.file("cl/ex_glmbyes_nmath", package = "opencltools")
kernels <- list.files(
  system.file("cl/ex_glmbyes_draft_src", package = "opencltools"),
  pattern = "\\\\.cl$", full.names = TRUE
)

# Step 1: scan draft kernels for library calls (read-only dry run)
step1 <- attach_kernel_call_tags(
  kernel_paths = kernels,
  library_dir = lib_dir,
  library_tag = "nmath",
  dry_run = TRUE
)
step1

# Step 2: expand transitive closure (small nmath library; runs on CRAN check)

```

```

nmath_small <- system.file("cl/nmath_small", package = "opencltools")
tagged      <- system.file("cl/src/dnorm_kernel.cl", package = "opencltools")
idx_small   <- write_kernel_dependency_index(library_dir = nmath_small, write = FALSE)

step2_small <- attach_cross_library_tags(
  kernel_paths = tagged,
  library_dir  = nmath_small,
  depends_tag  = "depends_nmath",
  index        = idx_small,
  dry_run      = TRUE
)
nrow(step2_small)

# Step 2: full nmath (slow)
nmath_dir <- system.file("cl/nmath", package = "opencltools")
idx       <- write_kernel_dependency_index(library_dir = nmath_dir, write = FALSE)

step2 <- attach_cross_library_tags(
  kernel_paths = tagged,
  library_dir  = nmath_dir,
  depends_tag  = "depends_nmath",
  index        = idx,
  dry_run      = TRUE
)
step2

#####
## End of kernel_tagging_workflow example
#####

```

---

attach\_kernel\_dependency\_tags

*Attach Dependency Tags to a Dependency-Ordered Kernel Library*

---

### Description

Compute and attach derived tags to each .cl file in a library: @load\_order, @all\_depends, and @all\_depends\_count.

### Usage

```
attach_kernel_dependency_tags(library_dir, dry_run = FALSE)
```

### Arguments

library\_dir     Directory containing .cl files with @depends tags.  
dry\_run         Logical; if TRUE, compute tags and reports without writing.

**Details**

Tags are written only if dependency sorting fully succeeds. If unresolved files remain, no files are modified and a cycle report is returned so source refactoring can be prioritized.

**Value**

A list with:

**ok** Logical success flag.

**message** Human-readable summary.

**sorted** Sorted records data frame.

**unresolved** Unresolved records data frame (empty on success).

**cycles** Cycle report data frame (empty on success).

**tags** Per-file tag data frame with `source_origin`, `source_type`, `includes`, `depends`, `provides`, `load_order`, `all_depends`, and `all_depends_count` (present on success).

**header\_functions** Functions declared in header files without `attribute_hidden`, including declaring header, inferred definition file, declaration signature, `define_alias`, and `all_depends` for the definition file.

An S3 class "openc1\_dependency\_tags" is attached for use with `print()`.

**Examples**

```
##### Start of attach_kernel_dependency_tags example #####
lib_dir <- system.file("cl/ex_glmbayes_nmath", package = "openc1tools")
res <- attach_kernel_dependency_tags(lib_dir, dry_run = TRUE)
res$ok
print(res)

#####
## End of attach_kernel_dependency_tags example
#####
```

**Description**

OpenCL-backed density, distribution, quantile, and random generation wrappers for the beta distribution. These mirror the base `stats` beta family while adding OpenCL dispatch and optional CPU fallback behavior.

**Usage**

```
dbeta_openc1(
  x,
  shape1,
  shape2,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
dnbeta_openc1(
  x,
  shape1,
  shape2,
  ncp,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
pbeta_openc1(
  q,
  shape1,
  shape2,
  ncp = 0,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
rbeta_openc1(n, shape1, shape2, fallback = FALSE, verbose = FALSE)
```

**Arguments**

x	(0, 1) quantiles for central/non-central densities on this page.
shape1	First shape parameter (must be > 0).
shape2	Second shape parameter (must be > 0).
log	log density flags for density wrappers (stats semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports

	no OpenCL. All wrappers on this page default FALSE so OpenCL build/runtime faults surface unless you opt in ( <code>fallback = TRUE</code> ). Kernels overlapping <code>'inst/cl/nmath/pgamma_utils'</code> can still be brittle on some devices; see <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'</code> and <code>'inst/OPENCL_KERNEL_KNOWN_FAILURES.md'</code> .
<code>verbose</code>	Logical; print fallback/error diagnostics.
<code>ncp</code>	Non-centrality parameter (must be $\geq 0$ ). Used by <code>dnbeta_openc1()</code> and <code>pbeta_openc1()</code> .
<code>q</code>	Numeric vector of quantiles for <code>pbeta_openc1</code> ; recycled like <code>stats::pbeta</code> .
<code>lower.tail, log.p</code>	Tail/log- $p$ inputs ( <code>stats</code> meanings).
<code>n</code>	Number of observations (non-negative integer scalar). Used only by <code>rbeta_openc1</code> .

**Value**

Numeric vector of length `n`.

**Known OpenCL limitations**

**Quantile path** The beta quantile OpenCL kernel fails at device link (`Rf_lbeta` unresolved), so no `qbeta` wrapper is exported; use `stats::qbeta()`.  
Tracker: `'inst/OPENCL_KERNEL_KNOWN_FAILURES.md'`

**References**

- Abramowitz M, Stegun IA (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York. ISBN 0486612724.
- Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.
- Brown BW, Levy LB (1994). "Certification of Algorithm 708: Significant-Digit Computation of the Incomplete Beta." *ACM Transactions on Mathematical Software*, **20**(3), 393–397. doi:10.1145/192115.192155.
- Cheng RCH (1978). "Generating Beta Variates with Nonintegral Shape Parameters." *Communications of the ACM*, **21**(4), 317–322. doi:10.1145/359460.359482.
- Cran GW, Martin KJ, Thomas GE (1977). "Remark AS R19 and Algorithm AS 109: A Remark on Algorithms AS 63: The Incomplete Beta Integral, AS 64: Inverse of the Incomplete Beta Function Ratio." *Applied Statistics*, **26**(1), 111–114. doi:10.2307/2346887.
- Didonato AR, Morris Jr. AH (1992). "Algorithm 708: Significant Digit Computation of the Incomplete Beta Function Ratios." *ACM Transactions on Mathematical Software*, **18**(3), 360–373. doi:10.1145/131766.131776.
- Frick H (1990). "Algorithm AS R84: A Remark on Algorithm AS 226: Computing Non-Central Beta Probabilities." *Applied Statistics*, **39**(2), 311–312. doi:10.2307/2347780.
- Johnson NL, Kotz S, Balakrishnan N (1995). *Continuous Univariate Distributions*, volume 2, 2nd edition. Wiley, New York. ISBN 978-0-471-58494-0.
- Lam ML (1995). "Remark AS R95: A Remark on Algorithm AS 226: Computing Non-Central Beta Probabilities." *Applied Statistics*, **44**(4), 551–552. doi:10.2307/2986147.
- Lenth RV (1987). "Algorithm AS 226: Computing Noncentral Beta Probabilities." *Applied Statistics*, **36**(2), 241–244. doi:10.2307/2347558.

**Examples**

```

n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dbeta_openc1(rep(0.6, n), shape1 = 2.5, shape2 = 4, fallback = FALSE, verbose = TRUE)
  دنبeta_openc1(rep(0.6, n), shape1 = 2.5, shape2 = 4, ncp = 0.8, fallback = FALSE, verbose = TRUE)
  pbeta_openc1(q = 0.6, shape1 = 2.5, shape2 = 4, ncp = 0, fallback = FALSE, verbose = TRUE)
  rbeta_openc1(n, shape1 = 2.5, shape2 = 4, fallback = FALSE, verbose = TRUE)
} else {
  stats::dbeta(rep(0.6, n), shape1 = 2.5, shape2 = 4)
  stats::pbeta(0.6, shape1 = 2.5, shape2 = 4, ncp = 0)
  stats::rbeta(n, shape1 = 2.5, shape2 = 4)
}

```

---

dbinom\_raw\_openc1

*The Binomial Distribution (OpenCL)*


---

**Description**

OpenCL-backed density, distribution, quantile, and random generation wrappers for the binomial distribution.

**Usage**

```

dbinom_raw_openc1(
  x,
  size,
  prob,
  qprob = NULL,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

dbinom_openc1(
  x,
  size,
  prob,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

pbinom_openc1(
  q,
  size,

```

```

    prob,
    lower.tail = TRUE,
    log.p = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
)

qbinom_openc1(
  p,
  size,
  prob,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

rbinom_openc1(n, size, prob, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile.
size	Number of trials (must be $\geq 0$ ).
prob	Probability of success in $[0, 1]$ .
qprob	Complementary probability. If NULL, uses $1 - \text{prob}$ .
log	log density switch for dbinom* wrappers (stats semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for $p/q$ wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. Density dbinom_* wrappers default FALSE; pbinom_openc1 defaults TRUE temporarily ('inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'); quantile and random wrappers default FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for pbinom_openc1; recycled like stats::pbinom.
lower.tail, log.p	Tail/log-p inputs (stats meanings).
p	Probabilities for qbinom_openc1; ordering matches stats::qbinom.
n	Number of observations (non-negative integer scalar). Used only by rbinom_openc1.

### Value

Numeric vector result from the corresponding binomial-family operation.

## References

Kachitvichyanukul V, Schmeiser BW (1988). “Binomial Random Variate Generation.” *Communications of the ACM*, **31**(2), 216–222. doi:10.1145/42372.42381.

Loader C (2000). “Fast and Accurate Computation of Binomial Probabilities.” Bell Laboratories. <https://www.r-project.org/doc/reports/CLoader-dbinom-2002.pdf>.

## Examples

```
n <- 5L
if (!mathopencl_has_opencl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dbinom_raw_opencl(rep(6, n), size = 10, prob = 0.3, fallback = FALSE, verbose = TRUE)
  dbinom_opencl(rep(6, n), size = 10, prob = 0.3, fallback = FALSE, verbose = TRUE)
  pbinom_opencl(q = 6, size = 10, prob = 0.3, fallback = FALSE, verbose = TRUE)
  qbinom_opencl(rep(0.8, n), size = 10, prob = 0.3, fallback = FALSE, verbose = TRUE)
  rbinom_opencl(n, size = 10, prob = 0.3, fallback = FALSE, verbose = TRUE)
} else {
  stats::dbinom(rep(6, n), size = 10, prob = 0.3)
  stats::dbinom(rep(6, n), size = 10, prob = 0.3)
  stats::pbinom(6, size = 10, prob = 0.3)
  stats::qbinom(rep(0.8, n), size = 10, prob = 0.3)
  stats::rbinom(n, size = 10, prob = 0.3)
}
```

---

dcauchy\_opencl

*The Cauchy Distribution (OpenCL)*


---

## Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the Cauchy distribution.

## Usage

```
dcauchy_opencl(
  x,
  location = 0,
  scale = 1,
  log = FALSE,
  opencl_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
pcauchy_opencl(
  q,
  location = 0,
  scale = 1,
```

```

    lower.tail = TRUE,
    log.p = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

qcauchy_openc1(
  p,
  location = 0,
  scale = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

rcauchy_openc1(n, location = 0, scale = 1, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile.
location	Location parameter.
scale	Scale parameter (must be > 0).
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Quantiles ( <code>pcauchy_openc1</code> ; aligns with <code>stats::pcauchy</code> recycling).
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for <code>qcauchy_openc1</code> (like <code>stats::qcauchy</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rcauchy_openc1</code> .

### Value

Numeric vector result from the corresponding Cauchy-family operation.

## References

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Johnson NL, Kotz S, Balakrishnan N (1994). *Continuous Univariate Distributions*, volume 1, 2nd edition. Wiley, New York. ISBN 978-0-471-58495-7.

## Examples

```
n <- 5L
if (!mathopengl_has_opengl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dcauchy_opengl(rep(0.2, n), location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
  pcauchy_opengl(q = 0.2, location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
  qcauchy_opengl(rep(0.8, n), location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
  rcauchy_opengl(n, location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::dcauchy(rep(0.2, n), location = 0, scale = 1)
  stats::pcauchy(0.2, location = 0, scale = 1)
  stats::qcauchy(rep(0.8, n), location = 0, scale = 1)
  stats::rcauchy(n, location = 0, scale = 1)
}
```

---

dchisq\_opengl

*The Chi-squared Distribution (OpenCL)*


---

## Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the chi-squared distribution, including non-central paths via ncp.

## Usage

```
dchisq_opengl(
  x,
  df,
  ncp = 0,
  log = FALSE,
  opengl_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
pchisq_opengl(
  q,
  df,
  ncp = 0,
  lower.tail = TRUE,
  log.p = FALSE,
```

```

    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

qchisq_openc1(
  p,
  df,
  ncp = 0,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

rchisq_openc1(n, df, ncp = 0, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric vector of quantiles for dchisq_openc1.
df	Degrees of freedom (must be > 0).
ncp	Non-centrality parameter (must be >= 0).
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. dchisq_openc1 defaults FALSE; distribution/quantile and rchisq_openc1 remain TRUE temporarily ('inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'); pass explicit fallback to override.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for pchisq_openc1; recycled like stats::pchisq.
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for qchisq_openc1 (like stats::qchisq).
n	Number of observations (non-negative integer scalar). Used only by rchisq_openc1; dchisq_openc1 takes vector x first (like stats::dchisq).

### Value

Numeric vector of length n.

## References

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Ding CG (1992). “Algorithm AS 275: Computing the Non-Central Chi-Squared Distribution Function.” *Applied Statistics*, **41**(2), 478–482. doi:10.2307/2347584.

Johnson NL, Kotz S, Balakrishnan N (1994). *Continuous Univariate Distributions*, volume 1, 2nd edition. Wiley, New York. ISBN 978-0-471-58495-7.

Johnson NL, Kotz S, Balakrishnan N (1995). *Continuous Univariate Distributions*, volume 2, 2nd edition. Wiley, New York. ISBN 978-0-471-58494-0.

## Examples

```
n <- 5L
if (!mathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dchisq_openc1(rep(4.5, n), df = 6, ncp = 0, fallback = FALSE, verbose = TRUE)
  pchisq_openc1(q = 4.5, df = 6, ncp = 0, fallback = FALSE, verbose = TRUE)
  qchisq_openc1(rep(0.8, n), df = 6, ncp = 0, fallback = FALSE, verbose = TRUE)
  rchisq_openc1(n, df = 6, ncp = 0, fallback = FALSE, verbose = TRUE)
} else {
  stats::dchisq(rep(4.5, n), df = 6, ncp = 0)
  stats::pchisq(4.5, df = 6, ncp = 0)
  stats::qchisq(rep(0.8, n), df = 6, ncp = 0)
  stats::rchisq(n, df = 6, ncp = 0)
}
```

---

dexp\_openc1

*The Exponential Distribution (OpenCL)*


---

## Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the exponential distribution.

## Usage

```
dexp_openc1(
  x,
  rate = 1,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

pexp_openc1(
  q,
  rate = 1,
```

```

    lower.tail = TRUE,
    log.p = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

qexp_openc1(
  p,
  rate = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

rexp_openc1(n, rate = 1, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile.
rate	Rate parameter (must be > 0).
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for <code>pexp_openc1</code> ; recycled like <code>stats::pexp</code> .
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for <code>qexp_openc1</code> (like <code>stats::qexp</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rexp_openc1</code> .

### Value

Numeric vector result from the corresponding exponential-family operation.

### References

- Ahrens JH, Dieter U (1972). “Computer Methods for Sampling from the Exponential and Normal Distributions.” *Communications of the ACM*, **15**(10), 873–882. doi:10.1145/355604.361593.
- Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Johnson NL, Kotz S, Balakrishnan N (1994). *Continuous Univariate Distributions*, volume 1, 2nd edition. Wiley, New York. ISBN 978-0-471-58495-7.

### Examples

```
n <- 5L
if (!mathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dexp_openc1(rep(1.2, n), rate = 1, fallback = FALSE, verbose = TRUE)
  pexp_openc1(q = 1.2, rate = 1, fallback = FALSE, verbose = TRUE)
  qexp_openc1(rep(0.8, n), rate = 1, fallback = FALSE, verbose = TRUE)
  rexp_openc1(n, rate = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::dexp(rep(1.2, n), rate = 1)
  stats::pexp(1.2, rate = 1)
  stats::qexp(rep(0.8, n), rate = 1)
  stats::rexp(n, rate = 1)
}
```

---

df\_openc1

*The F Distribution (OpenCL)*


---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the F distribution.

### Usage

```
df_openc1(
  x,
  df1,
  df2,
  ncp = 0,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
pf_openc1(
  q,
  df1,
  df2,
  ncp = 0,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
)
```

```

    verbose = FALSE
  )

  qf_openc1(
    p,
    df1,
    df2,
    ncp = 0,
    lower.tail = TRUE,
    log.p = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  rf_openc1(n, df1, df2, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric vector of quantiles ( <code>df_openc1</code> ).
df1	Numerator degrees of freedom (must be > 0).
df2	Denominator degrees of freedom (must be > 0).
ncp	Non-centrality parameter (must be >= 0). Used by <code>df_openc1()</code> , <code>pf_openc1()</code> , and <code>qf_openc1()</code> .
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. <code>df_openc1</code> defaults FALSE; <code>pf_openc1</code> , <code>qf_openc1</code> , and <code>rf_openc1</code> default TRUE temporarily ('inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md').
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles ( <code>pf_openc1</code> ); recycled like <code>stats::pf</code> .
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for <code>qf_openc1</code> (like <code>stats::qf</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rf_openc1</code> ; <code>df_openc1</code> takes vector <i>x</i> first (like <code>stats::df</code> ).

### Value

For `df_openc1`, `qf_openc1`, `rf_openc1`: numeric vector result. For `pf_openc1`: numeric vector of recycled length (see `stats::pf`).

### Known OpenCL limitations

qf\_openc1() can fail on some GPU/driver combinations with CL\_OUT\_OF\_RESOURCES. This has been observed in both central and non-central settings, with non-central paths typically more fragile.

### References

Abramowitz M, Stegun IA (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York. ISBN 0486612724.

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Johnson NL, Kotz S, Balakrishnan N (1995). *Continuous Univariate Distributions*, volume 2, 2nd edition. Wiley, New York. ISBN 978-0-471-58494-0.

### Examples

```
n <- 5L
if (!mathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  df_openc1(rep(2, n), df1 = 5, df2 = 9, ncp = 0, fallback = FALSE, verbose = TRUE)
  df_openc1(rep(2, n), df1 = 5, df2 = 9, ncp = 1.1, fallback = FALSE, verbose = TRUE)
  pf_openc1(q = 2, df1 = 5, df2 = 9, ncp = 0, fallback = FALSE, verbose = TRUE)
  pf_openc1(q = 2, df1 = 5, df2 = 9, ncp = 1.1, fallback = FALSE, verbose = TRUE)
  rf_openc1(n, df1 = 5, df2 = 9, fallback = FALSE, verbose = TRUE)
  ## qf_openc1: disabled - see inst/OPENCL_KERNEL_KNOWN_FAILURES.md
} else {
  stats::df(rep(2, n), df1 = 5, df2 = 9, ncp = 0)
  stats::df(rep(2, n), df1 = 5, df2 = 9, ncp = 1.1)
  stats::pf(2, df1 = 5, df2 = 9, ncp = 0)
  stats::pf(2, df1 = 5, df2 = 9, ncp = 1.1)
  stats::rf(n, df1 = 5, df2 = 9)
}
```

---

dgamma\_openc1

*The Gamma Distribution (OpenCL)*


---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the gamma distribution. These mirror the base stats gamma family while adding OpenCL dispatch and optional CPU fallback behavior.

### Usage

```
dgamma_openc1(
  x,
  shape,
  scale = 1,
  log = FALSE,
```

```

    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  pgamma_openc1(
    q,
    shape,
    rate = 1,
    scale = 1/rate,
    lower.tail = TRUE,
    log.p = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  rgamma_openc1(n, shape, scale = 1, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric vector of quantiles for <code>dgamma_openc1</code> .
shape	Shape parameter (must be > 0).
scale	Scale parameter (must be > 0). For <code>pgamma_openc1</code> , combined with <code>rate</code> like <code>stats::pgamma</code> .
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <code>pgamma_openc1</code> ; parallel dispatch reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. <code>dgamma_openc1</code> defaults FALSE; distribution/quantile wrappers follow 'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'. Pass explicit fallback where needed.
verbose	Logical; print informational fallback messages.
q	Numeric vector of quantiles for <code>pgamma_openc1</code> (same role as <code>stats::pgamma</code> ).
rate	Optional rate for <code>pgamma_openc1</code> ; see <code>pgamma</code> .
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
n	Number of observations (non-negative integer scalar). Used only by <code>rgamma_openc1</code> ; <code>dgamma_openc1</code> takes vector <code>x</code> first (like <code>stats::dgamma</code> ).

### Details

`pgamma_openc1` follows `pgamma` argument names and `rate/scale` handling (including the error when both are supplied inconsistently). Recycling of `q`, `shape`, and `scale` follows `stats::pgamma`. Vector `lower.tail` and `log.p` are recycled row-wise with those arguments (like `pnorm_openc1`);

a single-vector `stats::pgamma()` call does not apply tail flags element-wise. There is no leading `n` argument for `pgamma_openc1`. On the GPU path each recycled row runs `pgamma_kernel` once with `n_out = 1`. Missing or non-finite values after recycling, or non-positive shape/scale, use row-wise `stats::pgamma`.

## Value

Numeric vector result from the corresponding gamma-family operation.

## Known OpenCL limitations

Compilation of `qgamma_kernel` can fail (ptxas: unresolved `stirlerr_cycle_free`), so no `qgamma` wrapper is exported; use `stats::qgamma()`. See `'inst/OPENCL_KERNEL_KNOWN_FAILURES.md'`.

## References

- Abramowitz M, Stegun IA (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York. ISBN 0486612724.
- Ahrens JH, Dieter U (1974). “Computer Methods for Sampling from Gamma, Beta, Poisson and Binomial Distributions.” *Computing*, **12**(3), 223–246. doi:10.1007/bf02293108.
- Ahrens JH, Dieter U (1982). “Generating Gamma Variates by a Modified Rejection Technique.” *Communications of the ACM*, **25**(1), 47–54. doi:10.1145/358315.358390.
- Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.
- Best DJ, Roberts DE (1975). “Algorithm AS 89: The Upper Tail Probabilities of Spearman’s Rho.” *Applied Statistics*, **24**(3), 377–379. doi:10.2307/2347111.
- NIST (2024). “NIST Digital Library of Mathematical Functions.” <https://dlmf.nist.gov/>. Section 8.2.
- Shea BL (1988). “Algorithm AS 239: Chi-Squared and Incomplete Gamma Integral.” *Applied Statistics*, **37**(3), 466–473. doi:10.2307/2347328.

## Examples

```
n <- 5L
if (!mathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dgamma_openc1(rep(1.2, n), shape = 2, scale = 1, fallback = FALSE, verbose = TRUE)
  pgamma_openc1(q = 1.2, shape = 2, scale = 1, fallback = FALSE, verbose = TRUE)
  rgamma_openc1(n, shape = 2, scale = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::dgamma(rep(1.2, n), shape = 2, scale = 1)
  stats::pgamma(1.2, shape = 2, scale = 1)
  stats::rgamma(n, shape = 2, scale = 1)
}
```

---

dgeom\_opengl                      *The Geometric Distribution (OpenCL)*

---

## Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the geometric distribution.

## Usage

```
dgeom_opengl(  
  x,  
  prob,  
  log = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
pgeom_opengl(  
  q,  
  prob,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
qgeom_opengl(  
  p,  
  prob,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
rgeom_opengl(n, prob, fallback = FALSE, verbose = FALSE)
```

## Arguments

x	Numeric scalar quantile.
prob	Probability of success in $[0, 1]$ .
log	log flag for densities (stats <i>d</i> -family semantics).

openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for $p/q$ wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. <code>dgeom_openc1</code> defaults FALSE. <code>pgeom_openc1</code> , <code>qgeom_openc1</code> , and <code>rgeom_openc1</code> default FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for <code>pgeom_openc1</code> ; recycled like <code>stats::pgeom</code> .
lower.tail, log.p	Tail/log- $p$ inputs (stats meanings).
p	Numeric vector of probabilities for <code>qgeom_openc1</code> (like <code>stats::qgeom</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rgeom_openc1</code> .

**Value**

Numeric vector result from the corresponding geometric-family operation.

**References**

Devroye L (1986). *Non-Uniform Random Variate Generation*. Springer, New York.

**Examples**

```
n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dgeom_openc1(rep(4, n), prob = 0.3, fallback = FALSE, verbose = TRUE)
  pgeom_openc1(q = 4, prob = 0.3, fallback = FALSE, verbose = TRUE)
  qgeom_openc1(rep(0.8, n), prob = 0.3, fallback = FALSE, verbose = TRUE)
  rgeom_openc1(n, prob = 0.3, fallback = FALSE, verbose = TRUE)
} else {
  stats::dgeom(rep(4, n), prob = 0.3)
  stats::pgeom(4, prob = 0.3)
  stats::qgeom(rep(0.8, n), prob = 0.3)
  stats::rgeom(n, prob = 0.3)
}
```

**Description**

OpenCL-backed density, distribution, quantile, and random generation wrappers for the hypergeometric distribution.

**Usage**

```
dhyper_openc1(  
  x,  
  m,  
  n_black,  
  k,  
  log = FALSE,  
  openc1_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
phyper_openc1(  
  q,  
  m,  
  n_black,  
  k,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  openc1_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
qhyper_openc1(  
  p,  
  m,  
  n_black,  
  k,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  openc1_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
rhyper_openc1(n, m, n_black, k, fallback = FALSE, verbose = FALSE)
```

**Arguments**

x	Numeric scalar quantile.
m	Number of white balls in the urn (must be $\geq 0$ ).
n_black	Number of black balls in the urn (must be $\geq 0$ ).
k	Number of draws (must be $\geq 0$ ).
log	log flag for densities (stats <i>d</i> -family semantics).

openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for $p/q$ wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. <code>dhyper_openc1</code> defaults FALSE; <code>phyper_openc1</code> and <code>rhyper_openc1</code> default TRUE temporarily ( <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'</code> ); <code>qhyper_openc1</code> defaults FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for <code>phyper_openc1</code> ; recycled like <code>stats::phyper</code> .
lower.tail, log.p	Tail/log- $p$ inputs (stats meanings).
p	Numeric vector of probabilities for <code>qhyper_openc1</code> (like <code>stats::qhyper</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rhyper_openc1</code> .

**Value**

Numeric vector result from the corresponding hypergeometric-family operation.

**References**

Johnson NL, Kotz S, Kemp AW (1993). *Univariate Discrete Distributions*, 2nd edition. Wiley, New York. ISBN 978-0471548973.

Kachitvichyanukul V, Schmeiser BW (1985). "Computer Generation of Hypergeometric Random Variates." *Journal of Statistical Computation and Simulation*, **22**(2), 127–145. doi:10.1080/00949658508810839.

**Examples**

```
n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dhyper_openc1(rep(3, n), m = 10, n_black = 12, k = 8, fallback = FALSE, verbose = TRUE)
  phyper_openc1(q = 3, m = 10, n_black = 12, k = 8, fallback = FALSE, verbose = TRUE)
  qhyper_openc1(rep(0.8, n), m = 10, n_black = 12, k = 8, fallback = FALSE, verbose = TRUE)
  rhyper_openc1(n, m = 10, n_black = 12, k = 8, fallback = FALSE, verbose = TRUE)
} else {
  stats::dhyper(rep(3, n), m = 10, n = 12, k = 8)
  stats::phyper(3, m = 10, n = 12, k = 8)
  stats::qhyper(rep(0.8, n), m = 10, n = 12, k = 8)
  stats::rhyper(n, m = 10, n = 12, k = 8)
}
```

---

dlnorm\_openc1                      *The Lognormal Distribution (OpenCL)*

---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the lognormal distribution.

### Usage

```
dlnorm_openc1(  
  x,  
  meanlog = 0,  
  sdlog = 1,  
  log = FALSE,  
  openc1_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
plnorm_openc1(  
  q,  
  meanlog = 0,  
  sdlog = 1,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  openc1_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
qlnorm_openc1(  
  p,  
  meanlog = 0,  
  sdlog = 1,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  openc1_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
rlnorm_openc1(n, meanlog = 0, sdlog = 1, fallback = FALSE, verbose = FALSE)
```

### Arguments

x                      Numeric scalar quantile (must be  $\geq 0$ ).

meanlog	Mean of the distribution on the log scale.
sdlog	Standard deviation on the log scale (must be > 0).
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for <code>plnorm_openc1</code> ; recycled like <code>stats::plnorm</code> .
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for <code>qlnorm_openc1</code> (like <code>stats::qlnorm</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rlnorm_openc1</code> .

### Value

Numeric vector of length `n`.

### References

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Johnson NL, Kotz S, Balakrishnan N (1994). *Continuous Univariate Distributions*, volume 1, 2nd edition. Wiley, New York. ISBN 978-0-471-58495-7.

### Examples

```
n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dlnorm_openc1(rep(1.2, n), meanlog = 0.1, sdlog = 0.8, fallback = FALSE, verbose = TRUE)
  plnorm_openc1(q = 1.2, meanlog = 0.1, sdlog = 0.8, fallback = FALSE, verbose = TRUE)
  qlnorm_openc1(rep(0.8, n), meanlog = 0.1, sdlog = 0.8, fallback = FALSE, verbose = TRUE)
  rlnorm_openc1(n, meanlog = 0.1, sdlog = 0.8, fallback = FALSE, verbose = TRUE)
} else {
  stats::dlnorm(rep(1.2, n), meanlog = 0.1, sdlog = 0.8)
  stats::plnorm(1.2, meanlog = 0.1, sdlog = 0.8)
  stats::qlnorm(rep(0.8, n), meanlog = 0.1, sdlog = 0.8)
  stats::rlnorm(n, meanlog = 0.1, sdlog = 0.8)
}
```

---

dlogis_opengl	<i>The Logistic Distribution (OpenCL)</i>
---------------	---

---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the logistic distribution.

### Usage

```
dlogis_opengl(  
  x,  
  location = 0,  
  scale = 1,  
  log = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
plogis_opengl(  
  q,  
  location = 0,  
  scale = 1,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
qlogis_opengl(  
  p,  
  location = 0,  
  scale = 1,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
rlogis_opengl(n, location = 0, scale = 1, fallback = FALSE, verbose = FALSE)
```

### Arguments

x                    Numeric scalar quantile.

location	Location parameter.
scale	Scale parameter (must be > 0).
log	log density switch for dlogis_openc1 (stats semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for plogis_openc1 and qlogis_openc1; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. See <code>'inst/OPENCL_KERNEL_KNOWN_FAILURES.md'</code> .
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for plogis_openc1; recycled like <code>stats::plogis.lower.tail, log.p</code>
	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for qlogis_openc1 (like <code>stats::qlogis</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rlogis_openc1</code> .

### Value

Numeric vector result from the corresponding logistic-family operation.

### Known OpenCL limitations

Some platforms fail to link `qlogis_kernel` (ptxas unresolved `Rf_qlogis`). Runnable examples omit GPU `qlogis_openc1` until resolved. See `'inst/OPENCL_KERNEL_KNOWN_FAILURES.md'`.

### References

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Johnson NL, Kotz S, Balakrishnan N (1995). *Continuous Univariate Distributions*, volume 2, 2nd edition. Wiley, New York. ISBN 978-0-471-58494-0.

### Examples

```
n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dlogis_openc1(rep(0.2, n), location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
  plogis_openc1(q = 0.2, location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
  ## qlogis_openc1: disabled - see inst/OPENCL_KERNEL_KNOWN_FAILURES.md
  rlogis_openc1(n, location = 0, scale = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::dlogis(rep(0.2, n), location = 0, scale = 1)
  stats::plogis(0.2, location = 0, scale = 1)
  stats::rlogis(n, location = 0, scale = 1)
}
```

---

dnbinom\_opengl      *The Negative Binomial Distribution (OpenCL)*

---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the negative binomial distribution, with variants parameterized by prob and by mu.

### Usage

```
dnbinom_opengl(  
  x,  
  size,  
  prob,  
  log = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
pnbinom_opengl(  
  q,  
  size,  
  prob,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
qnbinom_opengl(  
  p,  
  size,  
  prob,  
  lower.tail = TRUE,  
  log.p = FALSE,  
  opengl_parallel = NA,  
  fallback = FALSE,  
  verbose = FALSE  
)
```

```
rnbinom_opengl(n, size, prob, fallback = FALSE, verbose = FALSE)
```

```
dnbinom_mu_opengl(  
  x,  
  size,
```

```

    mu,
    log = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
)

pnbinom_mu_openc1(
  q,
  size,
  mu,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

qnbinom_mu_openc1(
  p,
  size,
  mu,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

rnbinom_mu_openc1(n, size, mu, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile (must be $\geq 0$ ).
size	Dispersion/size parameter (must be $\geq 0$ ).
prob	Probability of success in $[0, 1]$ .
log	log density switch for density wrappers (stats semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for $p/q$ wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. Density wrappers ( <code>dnbinom_openc1</code> , <code>dnbinom_mu_openc1</code> ) default FALSE; <code>pnbinom_openc1/pnbinom_mu_openc1</code> default TRUE temporarily ( <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'</code> ); quantiles and random wrappers default FALSE.
verbose	Logical; print fallback/error diagnostics.

**q** Quantiles for pnbinom\* wrappers (stats::pnbinom recycling).  
**lower.tail, log.p** Tail/log-p inputs (stats meanings).  
**p** Probabilities for qnbinom\* wrappers (stats::qnbinom semantics).  
**n** Observations scalar; used by negative-binomial r\* wrappers only.  
**mu** Mean parameter (must be  $\geq 0$ ).

**Value**

Numeric vector result from the corresponding negative-binomial operation.

**References**

Devroye L (1986). *Non-Uniform Random Variate Generation*. Springer, New York.

**Examples**

```

n <- 5L
if (!nmathopencl_has_opencl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dnbinom_opencl(rep(4, n), size = 7, prob = 0.4, fallback = FALSE, verbose = TRUE)
  pnbinom_opencl(q = 4, size = 7, prob = 0.4, fallback = FALSE, verbose = TRUE)
  qnbinom_opencl(rep(0.8, n), size = 7, prob = 0.4, fallback = FALSE, verbose = TRUE)
  rnbinom_opencl(n, size = 7, prob = 0.4, fallback = FALSE, verbose = TRUE)

  dnbinom_mu_opencl(rep(4, n), size = 7, mu = 5, fallback = FALSE, verbose = TRUE)
  pnbinom_mu_opencl(q = 4, size = 7, mu = 5, fallback = FALSE, verbose = TRUE)
  qnbinom_mu_opencl(rep(0.8, n), size = 7, mu = 5, fallback = FALSE, verbose = TRUE)
  rnbinom_mu_opencl(n, size = 7, mu = 5, fallback = FALSE, verbose = TRUE)
} else {
  stats::dnbinom(rep(4, n), size = 7, prob = 0.4)
  stats::pnbinom(4, size = 7, prob = 0.4)
  stats::qnbinom(rep(0.8, n), size = 7, prob = 0.4)
  stats::rnbinom(n, size = 7, prob = 0.4)

  stats::dnbinom(rep(4, n), size = 7, mu = 5)
  stats::pnbinom(4, size = 7, mu = 5)
  stats::qnbinom(rep(0.8, n), size = 7, mu = 5)
  stats::rnbinom(n, size = 7, mu = 5)
}

```

---

dnorm\_opencl

*The Normal Distribution (OpenCL)*


---

**Description**

OpenCL-backed density, distribution, quantile, and random generation wrappers for the normal distribution. These mirror the base stats normal family while adding OpenCL dispatch and optional CPU fallback behavior.

**Usage**

```
dnorm_openc1(
  x,
  mean = 0,
  sd = 1,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
pnorm_openc1(
  q,
  mean = 0,
  sd = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
qnorm_openc1(
  p,
  mean = 0,
  sd = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
rnorm_openc1(n, mean = 0, sd = 1, fallback = FALSE, verbose = FALSE)
```

**Arguments**

x	Numeric vector of quantiles.
mean	Location parameter (rnorm: scalar). For pnorm_openc1 and qnorm_openc1, recycled like the corresponding stats function.
sd	Scale parameter (rnorm: scalar, sd >= 0). For pnorm_openc1 and qnorm_openc1, recycled like the corresponding stats function. The GPU path runs only when every recycled value is strictly positive (sd > 0); sd == 0 is evaluated with stats::pnorm.
log	log flag for dnorm_openc1 (stats semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for p/q wrappers on this page; parallel kernels reserved.

fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print informational fallback messages.
q	Numeric vector of quantiles for <code>pnorm_openc1</code> (same role as <code>stats::pnorm</code> ).
lower.tail, log.p	Recycling for <code>pnorm_openc1</code> ; see Details for contrasts with some vector <code>stats::pnorm</code> calls.
p	Numeric vector of probabilities for <code>qnorm_openc1</code> (like <code>stats::qnorm</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rnorm_openc1</code> ; not an argument to <code>pnorm_openc1</code> / <code>qnorm_openc1</code> .

### Details

`pnorm_openc1` matches `pnorm` on the first five statistical arguments (`q`, `mean`, `sd`, `lower.tail`, `log.p`) and the usual defaults. It also accepts `openc1_parallel`, `fallback`, and `verbose`. Only `rnorm_openc1` uses a leading `n`.

Recycling follows `pnorm` once arguments are aligned. On the GPU path, each recycled row calls the scalar `pnorm_kernel` once (`len` submissions).

Vector `q`, `mean`, `sd`, `lower.tail`, and `log.p` yield one OpenCL evaluation per output index.

Length-zero `q` returns `numeric(0)` as in `pnorm`.

Missing or non-finite values (after recycling), or any `sd == 0`, use CPU `pnorm`. Zero-length recycling errors and negative `sd` still error.

### Value

Numeric vector result from the corresponding normal-family operation.

### References

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

Cody WJ (1993). "Algorithm 715: SPECFUN – A Portable FORTRAN Package of Special Function Routines and Test Drivers." *ACM Transactions on Mathematical Software*, **19**(1), 22–30. doi:10.1145/151271.151273.

Johnson NL, Kotz S, Balakrishnan N (1994). *Continuous Univariate Distributions*, volume 1, 2nd edition. Wiley, New York. ISBN 978-0-471-58495-7.

Wichura MJ (1988). "Algorithm AS 241: The Percentage Points of the Normal Distribution." *Applied Statistics*, **37**(3), 477–484. doi:10.2307/2347330.

### Examples

```
n <- 5L
x <- c(-1, 0, 1)
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dnorm_openc1(x, mean = 0, sd = 1, fallback = FALSE, verbose = TRUE)
  pnorm_openc1(q = 0.2, mean = 0, sd = 1, fallback = FALSE, verbose = TRUE)
}
```

```

qnorm_opengl(rep(0.8, n), mean = 0, sd = 1, fallback = FALSE, verbose = TRUE)
rnorm_opengl(n, mean = 0, sd = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::dnorm(x, mean = 0, sd = 1)
  stats::pnorm(0.2, mean = 0, sd = 1)
  stats::qnorm(rep(0.8, n), mean = 0, sd = 1)
  stats::rnorm(n, mean = 0, sd = 1)
}

```

---

dpois\_raw\_opengl

*The Poisson Distribution (OpenCL)*


---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the Poisson distribution.

### Usage

```

dpois_raw_opengl(
  x,
  lambda,
  log = FALSE,
  opengl_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

dpois_opengl(
  x,
  lambda,
  log = FALSE,
  opengl_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

ppois_opengl(
  q,
  lambda,
  lower.tail = TRUE,
  log.p = FALSE,
  opengl_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

qpois_opengl(

```

```

    p,
    lambda,
    lower.tail = TRUE,
    log.p = FALSE,
    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  rpois_openc1(n, lambda, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile (must be $\geq 0$ ).
lambda	Mean/rate parameter (must be $\geq 0$ ).
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. Density <code>dpois_*</code> wrappers default FALSE; <code>ppois_openc1</code> defaults TRUE temporarily ( <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'</code> ); <code>qpois_openc1</code> and <code>rpois_openc1</code> default FALSE.
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles for <code>ppois_openc1</code> ; recycled like <code>stats::ppois</code> .
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for <code>qpois_openc1</code> (like <code>stats::qpois</code> ).
n	Number of observations (non-negative integer scalar). Used only by <code>rpois_openc1</code> .

### Value

Numeric vector result from the corresponding Poisson-family operation.

### References

Ahrens JH, Dieter U (1982). "Computer Generation of Poisson Deviates from Modified Normal Distributions." *ACM Transactions on Mathematical Software*, **8**(2), 163–179. doi:10.1145/355993.355997.

### Examples

```

n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dpois_raw_openc1(rep(4, n), lambda = 4, fallback = FALSE, verbose = TRUE)
  dpois_openc1(rep(4, n), lambda = 4, fallback = FALSE, verbose = TRUE)
}

```

```

ppois_openc1(q = 4, lambda = 4, fallback = FALSE, verbose = TRUE)
qpois_openc1(rep(0.8, n), lambda = 4, fallback = FALSE, verbose = TRUE)
rpois_openc1(n, lambda = 4, fallback = FALSE, verbose = TRUE)
} else {
  stats::dpois(rep(4, n), lambda = 4)
  stats::dpois(rep(4, n), lambda = 4)
  stats::ppois(4, lambda = 4)
  stats::qpois(rep(0.8, n), lambda = 4)
  stats::rpois(n, lambda = 4)
}

```

---

dt\_openc1

*The Student t Distribution (OpenCL)*


---

### Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the Student t distribution.

### Usage

```

dt_openc1(
  x,
  df,
  ncp = 0,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

pt_openc1(
  q,
  df,
  ncp = 0,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

```

```

qt_openc1(
  p,
  df,
  ncp = 0,
  lower.tail = TRUE,
  log.p = FALSE,

```

```

    openc1_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  rt_openc1(n, df, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric vector of quantiles (dt_openc1).
df	Degrees of freedom (must be > 0).
ncp	Non-centrality parameter.
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. dt_openc1 defaults FALSE; pt_openc1, qt_openc1, and rt_openc1 default TRUE temporarily ('inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md').
verbose	Logical; print fallback/error diagnostics.
q	Numeric vector of quantiles (pt_openc1); recycled like stats::pt.
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for qt_openc1 (like stats::qt).
n	Number of observations (non-negative integer scalar). Used only by rt_openc1; dt_openc1 takes vector x first (like stats::dt).

### Value

For dt\_openc1, qt\_openc1, rt\_openc1: numeric vector result. For pt\_openc1: numeric vector of recycled length (see stats::pt).

### Known OpenCL limitations

Some platforms fail to link qnt\_kernel (ptxas unresolved Rf\_qnt). Runnable examples omit GPU qt\_openc1 until resolved. See 'inst/OPENCL\_KERNEL\_KNOWN\_FAILURES.md'.

### References

- Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.
- Hill GW (1970). "ACM Algorithm 395: Student's t-Distribution." *Communications of the ACM*, **13**(10), 617–619. doi:10.1145/355598.355599.
- Hill GW (1981). "Remark on "Algorithm 396: Student's t-Quantiles"." *ACM Transactions on Mathematical Software*, **7**(2), 250–251. doi:10.1145/355945.355956.

Johnson NL, Kotz S, Balakrishnan N (1995). *Continuous Univariate Distributions*, volume 2, 2nd edition. Wiley, New York. ISBN 978-0-471-58494-0.

Lenth RV (1989). "Algorithm AS 243: Cumulative Distribution Function of the Non-Central t Distribution." *Applied Statistics*, **38**(1), 185–189. doi:[10.2307/2347693](https://doi.org/10.2307/2347693).

## Examples

```
n <- 5L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dt_openc1(rep(1.5, n), df = 8, ncp = 0, fallback = FALSE, verbose = TRUE)
  dt_openc1(rep(1.5, n), df = 8, ncp = 1.2, fallback = FALSE, verbose = TRUE)
  pt_openc1(q = 1.5, df = 8, ncp = 0, fallback = FALSE, verbose = TRUE)
  pt_openc1(q = 1.5, df = 8, ncp = 1.2, fallback = FALSE, verbose = TRUE)
  ## qt_openc1: disabled - see inst/OPENCL_KERNEL_KNOWN_FAILURES.md
  rt_openc1(n, df = 8, fallback = FALSE, verbose = TRUE)
} else {
  stats::dt(rep(1.5, n), df = 8, ncp = 0)
  stats::dt(rep(1.5, n), df = 8, ncp = 1.2)
  stats::pt(1.5, df = 8, ncp = 0)
  stats::pt(1.5, df = 8, ncp = 1.2)
  stats::rt(n, df = 8)
}
```

---

dunif\_openc1

*The Uniform Distribution (OpenCL)*

---

## Description

OpenCL-backed density, distribution, quantile, and random generation wrappers for the uniform distribution. These mirror the base stats uniform family while adding OpenCL dispatch and optional CPU fallback behavior.

## Usage

```
dunif_openc1(
  x,
  min = 0,
  max = 1,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
punif_openc1(
  q,
  min = 0,
  max = 1,
```

```

    lower.tail = TRUE,
    log.p = FALSE,
    opengl_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  qunif_opengl(
    p,
    min = 0,
    max = 1,
    lower.tail = TRUE,
    log.p = FALSE,
    opengl_parallel = NA,
    fallback = FALSE,
    verbose = FALSE
  )

  runif_opengl(n, min = 0, max = 1, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile for <code>dunif_opengl</code> .
min	Lower limit of the distribution.
max	Upper limit of the distribution.
log	log flag for densities (stats <i>d</i> -family semantics).
opengl_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopengl_has_opengl()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. See ‘inst/OPENCL_KERNEL_KNOWN_FAILURES.md’.
verbose	Logical; print informational fallback messages.
q	Numeric quantiles ( <code>punif_opengl</code> ; like <code>stats::punif</code> ).
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Probabilities for <code>qunif_opengl</code> (stats: <code>qunif</code> semantics).
n	Draw count ( <i>r*</i> ); density wrappers still lead with <i>x</i> .

### Value

Numeric vector result from the corresponding uniform-family operation.

### Known OpenCL limitations

Some platforms fail to link `qunif_kernel` (ptxas unresolved `Rf_qunif`). Runnable examples omit GPU `qunif_opengl` until resolved. See ‘inst/OPENCL\_KERNEL\_KNOWN\_FAILURES.md’.

**References**

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

**Examples**

```
n <- 5L
if (!mathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dunif_openc1(rep(0.4, n), min = 0, max = 1, fallback = FALSE, verbose = TRUE)
  punif_openc1(q = 0.4, min = 0, max = 1, fallback = FALSE, verbose = TRUE)
  ## qunif_openc1: disabled - see inst/OPENCL_KERNEL_KNOWN_FAILURES.md
  runif_openc1(n, min = 0, max = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::dunif(rep(0.4, n), min = 0, max = 1)
  stats::punif(0.4, min = 0, max = 1)
  stats::runif(n, min = 0, max = 1)
}
```

---

dweibull\_openc1

*The Weibull Distribution (OpenCL)*


---

**Description**

OpenCL-backed density, distribution, quantile, and random generation wrappers for the Weibull distribution.

**Usage**

```
dweibull_openc1(
  x,
  shape,
  scale = 1,
  log = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

pweibull_openc1(
  q,
  shape,
  scale = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```

qweibull_openc1(
  p,
  shape,
  scale = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)

rweibull_openc1(n, shape, scale = 1, fallback = FALSE, verbose = FALSE)

```

### Arguments

x	Numeric scalar quantile (must be $\geq 0$ ).
shape	Shape parameter (must be $> 0$ ).
scale	Scale parameter (must be $> 0$ ).
log	log flag for densities (stats <i>d</i> -family semantics).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for <i>p/q</i> wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. Defaults TRUE only for <code>qweibull_openc1</code> temporarily ( <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL.density/survival/rand</code> stay FALSE.
verbose	Logical; print fallback/error diagnostics.
q	<i>p</i> *-wrapper quantiles (stats::pweibull semantics).
lower.tail, log.p	Tail/log- <i>p</i> inputs (stats meanings).
p	Numeric vector of probabilities for <code>qweibull_openc1</code> (like stats::qweibull).
n	Number of observations (non-negative integer scalar). Used only by <code>rweibull_openc1</code> .

### Value

Numeric vector result from the corresponding Weibull-family operation.

### References

Johnson NL, Kotz S, Balakrishnan N (1994). *Continuous Univariate Distributions*, volume 1, 2nd edition. Wiley, New York. ISBN 978-0-471-58495-7.

**Examples**

```
n <- 5L
if (!nmathopencl_has_opencl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  dweibull_opencl(rep(1.2, n), shape = 2, scale = 1.5, fallback = FALSE, verbose = TRUE)
  pweibull_opencl(q = 1.2, shape = 2, scale = 1.5, fallback = FALSE, verbose = TRUE)
  qweibull_opencl(rep(0.8, n), shape = 2, scale = 1.5, fallback = FALSE, verbose = TRUE)
  rweibull_opencl(n, shape = 2, scale = 1.5, fallback = FALSE, verbose = TRUE)
} else {
  stats::dweibull(rep(1.2, n), shape = 2, scale = 1.5)
  stats::pweibull(1.2, shape = 2, scale = 1.5)
  stats::qweibull(rep(0.8, n), shape = 2, scale = 1.5)
  stats::rweibull(n, shape = 2, scale = 1.5)
}
```

---

extract\_library\_subset

*Extract a Minimal Library Subset for a Set of Kernels*

---

**Description**

Given one or more kernel .cl files and a library directory, determines the minimal set of library files required (union of all kernels' dependency annotations) and copies them — in dependency order — to a destination directory.

**Usage**

```
extract_library_subset(
  kernel_paths,
  library_dir,
  dest_dir,
  depends_tag = "all_depends",
  index = NULL,
  overwrite = FALSE
)
```

**Arguments**

kernel_paths	Character vector of paths to kernel .cl files. Each file is scanned for the annotation tag given by depends_tag.
library_dir	Path to the source library directory.
dest_dir	Path where files would be written. Must already exist for any copying to occur. If absent, a warning is issued, no directories are created, and nothing is copied; the returned data.frame still describes the planned subset (sources, intended destinations, copied = FALSE). In addition to the .cl files, kernel_dependency_index.rds is copied here when copying runs so the extracted subset can be used with a pre-loaded index.

depends_tag	Name of the annotation tag listing library file stems. Defaults to "all_depends". Pass "all_depends_nmath" for kernels that annotate their nmath dependencies with that tag.
index	Pre-loaded dependency index. If NULL, read from <code>file.path(library_dir, "kernel_dependency_index.rds")</code> with a <code>message()</code> nudging toward the recommended pattern.
overwrite	Logical; if FALSE (default) existing files in <code>dest_dir</code> are not overwritten — the copy is skipped and <code>copied = FALSE</code> in the returned data frame.

## Details

Use this to populate a project-local copy containing only the library files actually needed by your kernels. The result can then be committed alongside your kernel files, removing a runtime dependency on the full library.

Bundled libraries such as `inst/cl/nmath` ship `kernel_dependency_index.rds` next to their `.cl` shards. Pass `index =` when you pass a pre-loaded index object to avoid redundant reads; regenerate the files with [write\\_kernel\\_dependency\\_index](#), for example after porting via `nmathtools/port_inst_cl_nmath_from_src` in the `openclport` package source tree.

```
lib_dir <- system.file("cl/nmath", package = "opencltools")
kernel_paths <- system.file(
  c("cl/src/dnorm_kernel.cl", "cl/src/pnorm_kernel.cl"),
  package = "opencltools"
)
dest_dir <- tempfile("ex_subset"); dir.create(dest_dir)
## on.exit(unlink(dest_dir, recursive = TRUE), add = TRUE)
result <- extract_library_subset(
  kernel_paths, lib_dir, dest_dir,
  depends_tag = "all_depends_nmath")
```

`extract_library_subset()` evaluates `inst/extdata/opencl_known_failures.json` against the union of `depends_tag` annotations and launcher paths.

## Value

A `nmathopencl_lib_extract_df` subclass of `data.frame` with one row per library shard (`.cl` files in dependency order, followed by companion index files when planned or copied) and columns:

`stem` Stem name (filename without `.cl`), or index filenames.

`source` Full path to the source file under `library_dir`.

`dest` Intended destination path under `dest_dir`.

`copied` TRUE if copied; otherwise FALSE including when `dest_dir` is missing (`copied = FALSE` for every row), a source path is missing, or an existing destination was skipped.

**See Also**[printing methods](#)[load\\_library\\_for\\_kernel](#)[write\\_kernel\\_dependency\\_index](#)Other OpenCL kernel library subsets: [load\\_library\\_for\\_kernel\(\)](#), [write\\_kernel\\_dependency\\_index\(\)](#)**Examples**

```
##### Start of extract_library_subset example #####

## Small library (fast; runs on CRAN check)
lib_small <- system.file("cl/nmath_small", package = "opencltools")
kpath <- system.file("cl/src/dnorm_kernel.cl", package = "opencltools")
idx_small <- write_kernel_dependency_index(library_dir = lib_small, write = FALSE)
dest_small <- file.path(tempdir(), "opencltools_extract_small")
if (dir.exists(dest_small)) unlink(dest_small, recursive = TRUE)
dir.create(dest_small, recursive = TRUE)
on.exit(unlink(dest_small, recursive = TRUE), add = TRUE)
df_small <- extract_library_subset(
  kpath, lib_small, dest_small,
  depends_tag = "all_depends_nmath",
  index = idx_small
)
sum(df_small$copied)

## Full nmath (slow)
lib_dir <- system.file("cl/nmath", package = "opencltools")
kernel_paths <- system.file(
  c("cl/src/dnorm_kernel.cl", "cl/src/pnorm_kernel.cl"),
  package = "opencltools"
)
idx <- write_kernel_dependency_index(library_dir = lib_dir, write = FALSE)
dest_dir <- file.path(tempdir(), "opencltools_extract_example")
if (dir.exists(dest_dir)) unlink(dest_dir, recursive = TRUE)
dir.create(dest_dir, recursive = TRUE)
on.exit(unlink(dest_dir, recursive = TRUE), add = TRUE)
df <- extract_library_subset(
  kernel_paths, lib_dir, dest_dir,
  depends_tag = "all_depends_nmath",
  index = idx
)
print(df)
sum(df$copied)

#####
## End of extract_library_subset example
#####
```

---

Ex\_EnvelopeEval      *Evaluate Negative Log-Likelihood and Gradients*

---

### Description

These functions implement the grid evaluation logic used in envelope construction for rejection sampling. They make use of the theory described in (Nygren and Nygren 2006) and the general implementation outlined in (Nygren 2025).

### Usage

```
Ex_EnvelopeEval(G4, y, x, mu, P, alpha, wt,
               family, link,
               use_openc1 = FALSE, verbose = FALSE)
```

### Arguments

G4	Numeric matrix of parameter values (parameters * grid points).
y	Numeric response vector.
x	Numeric design matrix.
mu	Numeric matrix of offsets or prior means.
P	Numeric matrix representing the portion of the prior precision shifted into the likelihood.
alpha	Numeric offset vector of length m
wt	Numeric vector of weights.
family	Character string; model family (e.g. "gaussian").
link	Character string; link function (e.g. "identity").
use_openc1	Logical; if TRUE, attempt OpenCL acceleration.
verbose	Logical; if TRUE, print diagnostic output.

### Details

Compact overview.

Derivations:\cr vignettes/equations ((Nygren and Nygren 2006)).

- Dispatcher fans out to CPU and OpenCL kernel runners.
- NegLL vectors plus cbars matrices feed envelopes in rNormal\_reg internals.
- Acceptance inequalities mirror the vignette “Simulation execution” chapter.

**Value**

**Ex\_EnvelopeEval** NegLL: negatives logLik; cbars: tangent gradients.

**f2\_f3\_non\_openc1** qf/grad from CPU kernels.

**f2\_f3\_openc1** qf/grad from OpenCL.

**run\_openc1\_pilot** Pilot runtime estimate (seconds).

**References**

Nygren K (2025). “Chapter A05: Simulation Methods – Likelihood Subgradient Densities.” Vignette in the glmbayes R package. R vignette name: Chapter-A05.

Nygren K~N, Nygren L~M (2006). “Likelihood Subgradient Densities.” *Journal of the American Statistical Association*, **101**(475), 1144–1156. doi:10.1198/016214506000000357.

**See Also**

[Ex\\_EnvelopeSize](#) (Nygren 2025) (Nygren 2025) (Nygren 2025) (Nygren 2025)

**Examples**

```
##### Start of Ex_EnvelopeEval example #####

# This example demonstrates Ex_EnvelopeEval in isolation. Ex_EnvelopeEval evaluates
# the negative log-likelihood and gradients at a grid of parameter values.
# It is called internally by EnvelopeBuild. Here we build the same inputs
# (grid G4, standardized model) using Ex_EnvelopeSize and expand.grid, then
# call Ex_EnvelopeEval directly. The setup mirrors Ex_EnvelopeBuild through
# the standardization step.

data(menarche, package = "MASS")
Age2 <- menarche$Age - 13

x <- matrix(as.numeric(1.0), nrow = length(Age2), ncol = 2)
x[, 2] <- Age2

y <- menarche$Menarche / menarche$Total
wt <- menarche$Total

mu <- matrix(as.numeric(0.0), nrow = 2, ncol = 1)
mu[2, 1] <- (log(0.9 / 0.1) - log(0.5 / 0.5)) / 3

V1 <- 1 * diag(as.numeric(2.0))
V1[1, 1] <- ((log(0.9 / 0.1) - log(0.5 / 0.5)) / 2)^2
V1[2, 2] <- (3 * mu[2, 1] / 2)^2

famfunc <- Ex_glmbfamfunc(binomial(logit))
f2 <- famfunc$f2
f3 <- famfunc$f3

dispersion2 <- as.numeric(1.0)
```

```

start <- mu
offset2 <- rep(as.numeric(0.0), length(y))
P <- solve(V1)
n <- 1000

wt2 <- wt / dispersion2
alpha <- x %*% as.vector(mu) + offset2
mu2 <- 0 * as.vector(mu)
P2 <- P
x2 <- x

parin <- start - mu
opt_out <- optim(parin, f2, f3,
  y = as.vector(y), x = as.matrix(x), mu = as.vector(mu2),
  P = as.matrix(P), alpha = as.vector(alpha), wt = as.vector(wt2),
  method = "BFGS", hessian = TRUE
)

bstar <- opt_out$par
A1 <- opt_out$hessian

Standard_Mod <- Ex_glm_Standardize_Model(
  y = as.vector(y), x = as.matrix(x), P = as.matrix(P),
  bstar = as.matrix(bstar, ncol = 1), A1 = as.matrix(A1)
)

bstar2 <- Standard_Mod$bstar2
A <- Standard_Mod$A
x2 <- Standard_Mod$x2
mu2 <- Standard_Mod$mu2
P2 <- Standard_Mod$P2

#####
# Build grid G4 via Ex_EnvelopeSize and expand.grid (as EnvelopeBuild does)
#####
a <- diag(A)
omega <- (sqrt(2) - exp(-1.20491 - 0.7321 * sqrt(0.5 + a))) / sqrt(1 + a)
b2 <- as.vector(bstar2)
G1 <- rbind(b2 - omega, b2, b2 + omega)

size_info <- Ex_EnvelopeSize(a, G1, Gridtype = 3L, n = n)
G2 <- size_info$G2

G3 <- as.matrix(do.call(expand.grid, G2))
G4 <- t(G3)

#####
# Ex_EnvelopeEval: negative log-likelihood and gradients at grid points
#####
eval_out <- Ex_EnvelopeEval(
  G4 = G4,
  y = y,
  x = as.matrix(x2),

```

```

mu = as.matrix(mu2, ncol = 1),
P = as.matrix(P2),
alpha = as.vector(alpha),
wt = as.vector(wt2),
family = "binomial",
link = "logit",
use_openc1 = FALSE,
verbose = FALSE
)

eval_out$NegLL
eval_out$cbars

#####
# End of Ex_EnvelopeEval example
#####

```

---

Ex\_EnvelopeSize      *Envelope Sizing and Optimization*

---

### Description

Ex\_EnvelopeSize() is the high-level entry point that constructs per-dimension grids and expected draw counts, while Ex\_EnvelopeOpt() performs the adaptive optimization used when Gridtype = 2.

### Usage

```
Ex_EnvelopeSize(a, G1, Gridtype = 2L, n = 1000L, n_envopt = -1,
               use_openc1 = FALSE, verbose = FALSE)
```

### Arguments

a	Numeric vector of diagonal precisions for the log-likelihood (posterior precision is $1 + a_i$ ).
G1	Numeric matrix of candidate grid points (3 * 11).
Gridtype	Integer code controlling grid sizing logic: <ul style="list-style-type: none"> <li>• 1 = static threshold test</li> <li>• 2 = adaptive optimization via Ex_EnvelopeOpt()</li> <li>• 3 = always three-point grid</li> <li>• 4 = always single-point grid</li> </ul>
n	Integer; number of posterior draws to generate (used for grid sizing).
n_envopt	Integer; effective sample size passed to Ex_EnvelopeOpt. Defaults to -1, which means "use n".
use_openc1	Logical; if TRUE, attempt GPU acceleration.
verbose	Logical; if TRUE, print progress messages.

## Details

These functions implement the grid sizing logic used in envelope construction for rejection sampling. They make use of the theory described in (Nygren and Nygren 2006) and the general implementation outlined in (Nygren 2025).

Ex\_EnvelopeSize() returns the constructed grid (G2), index vectors (GIndex1), expected draw count (E\_draws), and the per-dimension grid index.

Ex\_EnvelopeOpt() implements the adaptive optimization used in Gridtype = 2, ranking dimensions by posterior variance and promoting them to three-point tangents when the tradeoff is favorable.

## Value

Ex\_EnvelopeSize() A list with components G2, GIndex1, E\_draws, and gridindex.

Ex\_EnvelopeOpt() An integer vector of length  $l1$  with entries 1 (single-point) or 3 (three-point).

## Gridtype Logic and Candidates per Draw

The envelope sizing logic follows the analysis of (Nygren and Nygren 2006).

**Gridtype 1: Static Threshold** For each dimension  $i$ , if  $\sqrt{1 + a_i} \leq 2/\sqrt{\pi} \approx 1.128379$ , then a single tangent at the posterior mode suffices. Expected candidates per draw in that dimension:  $\sqrt{1 + a_i}$ . Otherwise, a symmetric three-point envelope is used at  $(\theta_i^* - \omega_i, \theta_i^*, \theta_i^* + \omega_i)$ , with expected candidates per draw bounded above by  $2/\sqrt{\pi}$ .

**Gridtype 2: Adaptive Optimization** Each dimension is assigned either a single-point or three-point envelope by minimizing

$$T_{\text{total}}(g_i) = T_{\text{build}}(g_i) + T_{\text{sample}}(n, \text{acc}_i(g_i)).$$

The optimizer balances build cost (grows with number of tangents) against sampling cost (decreases as acceptance improves). Expected candidates per draw:  $\prod_j \text{scaleest}_{i,j}$ , where each factor is either  $\sqrt{1 + a_j}$  (single-point) or  $2/\sqrt{\pi}$  (three-point), depending on the optimization outcome.

**Gridtype 3: Always Three-Point** Every dimension uses a symmetric three-point envelope. Expected candidates per draw:

$$\left(\frac{2}{\sqrt{\pi}}\right)^k$$

for  $k$  dimensions, as shown in Theorem 3 of (Nygren and Nygren 2006).

**Gridtype 4: Always Single-Point** Every dimension uses a single tangent at the posterior mode. Expected candidates per draw:

$$\prod_{i=1}^k \sqrt{1 + a_i}$$

(Example 1 in (Nygren and Nygren 2006)).

## References

Nygren K (2025). “Chapter A05: Simulation Methods – Likelihood Subgradient Densities.” Vignette in the glmbayes R package. R vignette name: Chapter-A05.

Nygren K~N, Nygren L~M (2006). “Likelihood Subgradient Densities.” *Journal of the American Statistical Association*, **101**(475), 1144–1156. doi:10.1198/016214506000000357.

## See Also

[Ex\\_EnvelopeEval](#) for evaluating these grids. (Nygren 2025) (Nygren 2025)

---

Ex_glmbfamfunc	<i>Return family functions used during simulation and post processing</i>
----------------	---

---

## Description

This function takes as input a [family](#) object and returns a set of functions that are used during simulation and summarization of models using the simulation functions in this package.

## Usage

```
Ex_glmbfamfunc(family)

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

## Arguments

family	an object of class <a href="#">family</a>
x	an object of class "Ex_glmbfamfunc" for which a printed output is desired.
...	additional optional arguments

## Details

Compiled paths dominate;  
retain Ex\_glmbfamfunc closures for scripted workflows.

## Value

A list (class "Ex\_glmbfamfunc") whose first four components are **always** present for every supported family and link. The names f1–f4 are stable: they mean the same roles across families (only the internal formulas change).

f1 Neg log-likelihood in coefficients b (usual data args).

f2 Neg log-posterior: likelihood plus Normal( $\mu$ , P) quadratic penalty.

f3 Gradient of f2 w.r.t.  $\mathbf{b}$  (argument pattern mirrors f2).

f4 Deviance gap vs saturation; honors dispersion;  
quasi / DIC helper.

f7 Weighted curvature / Hessian proxy at b.

Slots f5 and f6 are **not** returned: they were reserved for alternate or C++-aligned likelihood/posterior routines and remain commented out in the implementation (only f1, f2, f3, f4, and f7 are assigned in the returned list).

---

Ex\_glmb\_Standardize\_Model

*Standardize A Non-Gaussian Model*

---

### Description

Standardizes a Non-Gaussian Model prior to Envelope Creation

### Usage

Ex\_glmb\_Standardize\_Model(y, x, P, bstar, A1)

### Arguments

y	a vector of observations of length m
x	a design matrix of dimension m*p
P	Positive-definite prior precision ( <i>m times m</i> ).
bstar	a matrix containing the posterior mode from an optimization step
A1	a matrix containing the posterior precision matrix at the posterior mode

### Details

Starts from the posterior mode quantities and proceeds in three transformations.  
 Step 1: posterior-precision eigendecomp.  
 Interim model gets identity posterior precision.  
 Step 2: isolate diagonal epsilon; remainder behaves like likelihood information.  
 Step 3: another eigendecomp diagonalizes data precision A.  
 The prior tied to epsilon becomes identity.  
 (Nygren and Nygren 2006)

### Value

A list with the following components

bstar2	Standardized Posterior Mode
A	Standardized Data Precision Matrix

x2	Standardized Design Matrix
mu2	Standardized Prior Mean vector
P2	Standardized Precision Matrix Added to log-likelihood
L2Inv	A matrix used when undoing the first step in standardization described below
L3Inv	A matrix used when undoing the second step in standardization described below

## References

Nygren K~N, Nygren L~M (2006). “Likelihood Subgradient Densities.” *Journal of the American Statistical Association*, **101**(475), 1144–1156. doi:[10.1198/016214506000000357](https://doi.org/10.1198/016214506000000357).

---

gammafn_openc1	<i>Special Functions (OpenCL)</i>
----------------	-----------------------------------

---

## Description

OpenCL-backed wrappers for selected special functions from R Mathlib.

## Usage

```
gammafn_openc1(x, fallback = FALSE, verbose = FALSE)
lgammafn_openc1(x, fallback = FALSE, verbose = FALSE)
digamma_openc1(x, fallback = FALSE, verbose = FALSE)
trigamma_openc1(x, fallback = FALSE, verbose = FALSE)
tetragamma_openc1(x, fallback = FALSE, verbose = FALSE)
pentagramma_openc1(x, fallback = FALSE, verbose = FALSE)
psigamma_openc1(x, deriv, fallback = FALSE, verbose = FALSE)
beta_openc1(a, b, fallback = FALSE, verbose = FALSE)
lbeta_openc1(a, b, fallback = FALSE, verbose = FALSE)
choose_openc1(n, k, fallback = FALSE, verbose = FALSE)
lchoose_openc1(n, k, fallback = FALSE, verbose = FALSE)
```

**Arguments**

x	Numeric vector (and additional vectors where listed); arguments are recycled to a common length like the corresponding base functions.
fallback	When TRUE while <code>nmathopengl_has_opengl()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.
deriv	Derivative order for <code>psigamma_opengl</code> (recycled with x).
a, b	Parameters for <code>beta_opengl</code> / <code>lbeta_opengl</code> (recycled together).
n, k	Arguments for <code>choose_opengl</code> / <code>lchoose_opengl</code> , like <code>base::choose(n, k)</code> (recycled together).

**Value**

Numeric vector of recycled common length.

**References**

Abramowitz M, Stegun IA (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York. ISBN 0486612724.

Amos DE (1983). "Algorithm 610: A Portable FORTRAN Subroutine for Derivatives of the Psi Function." *ACM Transactions on Mathematical Software*, **9**(4), 494–502. doi:10.1145/356056.356065.

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

**Examples**

```
if (!nmathopengl_has_opengl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  gammafn_opengl(x = 2.5, fallback = FALSE, verbose = TRUE)
  lgammafn_opengl(x = 2.5, fallback = FALSE, verbose = TRUE)
  digamma_opengl(x = 2.5, fallback = FALSE, verbose = TRUE)
  trigamma_opengl(x = 2.5, fallback = FALSE, verbose = TRUE)
  tetragamma_opengl(x = 2.5, fallback = FALSE, verbose = TRUE)
  pentagramma_opengl(x = 2.5, fallback = FALSE, verbose = TRUE)
  psigamma_opengl(x = 2.5, deriv = 1, fallback = FALSE, verbose = TRUE)

  beta_opengl(a = 2.5, b = 3.0, fallback = FALSE, verbose = TRUE)
  lbeta_opengl(a = 2.5, b = 3.0, fallback = FALSE, verbose = TRUE)
  choose_opengl(n = 10, k = 4, fallback = FALSE, verbose = TRUE)
  lchoose_opengl(n = 10, k = 4, fallback = FALSE, verbose = TRUE)
} else {
  base::gamma(2.5)
  base::lgamma(2.5)
  base::digamma(2.5)
  base::trigamma(2.5)
  base::psigamma(2.5, deriv = 2L)
  base::psigamma(2.5, deriv = 3L)
  base::psigamma(2.5, deriv = 1L)
}
```

```

base::beta(2.5, 3.0)
base::lbeta(2.5, 3.0)
base::choose(10, 4)
base::lchoose(10, 4)
}

```

---

glmbyesEnvelopeExample

*Envelope Evaluation Utilities*

---

### Description

Core utilities for envelope-based posterior simulation using GPU-accelerated OpenCL kernels. These functions support the construction and evaluation of log-likelihood envelopes used in rejection sampling, and serve as an example of how downstream packages can build custom OpenCL kernels on top of the ported nmath routines in **nmathopenc1**.

### References

There are no references for Rd macro \insertAllCites on this help page.

---

imax2\_openc1

*Math Support Functions (OpenCL)*

---

### Description

OpenCL-backed wrappers for miscellaneous scalar support functions.

### Usage

```

imax2_openc1(x, y, fallback = FALSE, verbose = FALSE)
imin2_openc1(x, y, fallback = FALSE, verbose = FALSE)
fmax2_openc1(x, y, fallback = FALSE, verbose = FALSE)
fmin2_openc1(x, y, fallback = FALSE, verbose = FALSE)
sign_openc1(x, fallback = FALSE, verbose = FALSE)
fprec_openc1(x, digits, fallback = FALSE, verbose = FALSE)
fround_openc1(x, digits, fallback = FALSE, verbose = FALSE)
fsign_openc1(x, y, fallback = FALSE, verbose = FALSE)
ftrunc_openc1(x, fallback = FALSE, verbose = FALSE)

```

**Arguments**

x, y	Numeric vectors recycled together (where both appear).
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.
digits	Numeric vector recycled with x for precision/rounding helpers.

**Value**

Numeric vector of recycled common length.

**References**

Abramowitz M, Stegun IA (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York. ISBN 0486612724.

Becker RA, Chambers JM, Wilks AR (1988). *The New S Language*. Chapman and Hall/CRC, London. ISBN 053409192X.

**Examples**

```
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  imax2_openc1(x = 7, y = 3, fallback = FALSE, verbose = TRUE)
  imin2_openc1(x = 7, y = 3, fallback = FALSE, verbose = TRUE)
  fmax2_openc1(x = 7.2, y = 3.1, fallback = FALSE, verbose = TRUE)
  fmin2_openc1(x = 7.2, y = 3.1, fallback = FALSE, verbose = TRUE)
  sign_openc1(x = -2.5, fallback = FALSE, verbose = TRUE)
  fprec_openc1(x = 123.456, digits = 4, fallback = FALSE, verbose = TRUE)
  fround_openc1(x = 123.456, digits = 2, fallback = FALSE, verbose = TRUE)
  fsign_openc1(x = -2.5, y = 4.0, fallback = FALSE, verbose = TRUE)
  ftrunc_openc1(x = 123.456, fallback = FALSE, verbose = TRUE)
} else {
  as.double(pmax(7L, 3L))
  as.double(pmin(7L, 3L))
  pmax(7.2, 3.1)
  pmin(7.2, 3.1)
  base::sign(-2.5)
  signif(123.456, digits = 4)
  base::round(123.456, digits = 2)
  base::sign(-2.5) * abs(4.0)
  base::trunc(123.456)
}
```

---

 load\_library\_for\_kernel

*Load a Minimal OpenCL Library Subset for a Single Kernel*


---

## Description

Given a single kernel .cl file and a library directory (with an associated dependency index), reads the annotation tag that lists needed library files and returns their source code concatenated in the correct dependency order.

## Usage

```
load_library_for_kernel(
  kernel_path,
  library_dir,
  depends_tag = "all_depends",
  index = NULL
)
```

## Arguments

kernel_path	Path to a single .cl kernel file. The file is scanned for the annotation tag given by depends_tag.
library_dir	Path to the library directory containing the .cl source files (e.g. system.file("cl/nmath", package = "opencltools")).
depends_tag	Name of the annotation tag in the kernel file that lists the required library file stems. Defaults to "all_depends". For kernels annotated with @all_depends_nmath, pass depends_tag = "all_depends_nmath".
index	Optional RDS list; NULL triggers lazy reads.

## Details

For repeated calls in R code, loading kernel\_dependency\_index.rds once and passing index = avoids redundant disk reads. The bundled cl/nmath directory ships kernel\_dependency\_index.rds beside the .cl files; use [write\\_kernel\\_dependency\\_index](#) to regenerate it after porting (for example via nmathtools/port\_inst\_cl\_nmath\_from\_src.R in the openclport package source tree).

```
lib_dir <- system.file("cl/nmath", package = "opencltools")
kpath <- system.file("cl/src/dnorm_kernel.cl", package = "opencltools")
src <- load_library_for_kernel(
  kpath, lib_dir,
  depends_tag = "all_depends_nmath")
```

Subsets come only from each launcher file's transitive @all\_depends\_nmath-style annotations (use [attach\\_cross\\_library\\_tags](#) where dependency graphs span libraries). Deliberately loading every .cl under cl/nmath remains available via [load\\_kernel\\_library](#)(..., "nmath") where appropriate.

When `inst/extdata/opencv_known_failures.json` matches the launcher path or the declared / loaded stems, `warning(...)` points to fragile ports.

### Value

A character vector subclass `nmathopencv_concatenated_lib` holding concatenated sources (often length 1; blank annotations yield length-zero concatenation). Attachments describe requested and loaded library stems, paths, and byte size; see [print methods](#).

### See Also

[extract\\_library\\_subset](#)

[printing methods](#)

[write\\_kernel\\_dependency\\_index](#)

Other OpenCL kernel library subsets: [extract\\_library\\_subset\(\)](#), [write\\_kernel\\_dependency\\_index\(\)](#)

### Examples

```
##### Start of load_library_for_kernel example #####

## Small library (fast; runs on CRAN check)
lib_small <- system.file("cl/nmath_small", package = "opencvtools")
kpath <- system.file("cl/src/dnorm_kernel.cl", package = "opencvtools")
idx_small <- write_kernel_dependency_index(library_dir = lib_small, write = FALSE)
src_small <- load_library_for_kernel(
  kpath, lib_small,
  depends_tag = "all_depends_nmath",
  index = idx_small
)
nzchar(src_small)

## Full nmath (slow; rebuilds index over all shards)
lib_dir <- system.file("cl/nmath", package = "opencvtools")
idx <- write_kernel_dependency_index(library_dir = lib_dir, write = FALSE)
src <- load_library_for_kernel(
  kpath, lib_dir,
  depends_tag = "all_depends_nmath",
  index = idx
)
print(src)
nzchar(src)

#####
## End of load_library_for_kernel example
#####
```

---

nmathopencl\_has\_opencl

*GPU and OpenCL diagnostics for nmathopencl*


---

## Description

Compile-time and device-selection probes for **nmathopencl**.

Low-level workstation probes (GPU vendor detection, driver and ICD checks, PATH validation, `verify_opencl_runtime()`, PATH helpers, combined diagnostic reports, etc.) live in **opencltools**; call them as `opencltools::detect_environment_and_gpus()`, `opencltools::diagnose_glbayes()`, and related topics documented under `?opencltools`.

## Usage

```
nmathopencl_has_opencl()
```

```
nmathopencl_opencl_device_info(force = FALSE, details = FALSE)
```

```
nmathopencl_opencl_fp64_available(force = FALSE)
```

```
nmathopencl_opencl_reset_device_selection()
```

## Arguments

<code>force</code>	If TRUE, rerun discovery even when a previous selection is cached.
<code>details</code>	If TRUE, include a candidates list describing every platform/device pair (extension flag and probe result per device).

## Details

GPU acceleration uses OpenCL kernels and `*_opencl` wrappers when `nmathopencl_has_opencl()` is TRUE and a suitable device is available ((Stone et al. 2010)). CPU fallbacks apply for many routines when OpenCL is absent at compile time or runtime.

For a readable host/runtime report, use `opencltools::diagnose_glbayes()`; use `nmathopencl_has_opencl()` for this package's compile-time flag. Setup: `vignette("Chapter-01", package = "nmathopencl")`; packaged GPU API: `vignette("Chapter-12", package = "nmathopencl")`.

## Value

A list with `ok` (logical), `reason` (character), `indices`, `vendor/name` strings, `device_type`, `extension_cl_khr_fp64`, `probe_fp64_ok`, `selection_policy`, and optionally `candidates`.

## Functions

- `nmathopencl_openc1_device_info()`: Cached OpenCL device selection for double-precision (`cl_khr_fp64`) kernels in **nmathopencl**: enumerates platforms and devices, prefers GPU, checks the extension token, then verifies with a tiny `clBuildProgram` probe. Override with environment variables `NMATHOPENCL_PLATFORM_INDEX` and/or `NMATHOPENCL_DEVICE_INDEX` (0-based; device index is within the platform's device list). Use `nmathopencl_openc1_reset_device_selection()` to clear the cache (e.g. after driver changes).
- `nmathopencl_openc1_fp64_available()`: Returns TRUE if a cached OpenCL device passes the `cl_khr_fp64` extension check and build probe used for double kernels in **nmathopencl** (uses this package's compile-time OpenCL build and device cache, not **opencltools**).
- `nmathopencl_openc1_reset_device_selection()`: Clears the process-local OpenCL device selection cache so the next kernel or `nmathopencl_openc1_device_info()` run re-enumerates devices.

## Diagnostics exported from nmathopencl

- `nmathopencl_has_openc1()` — TRUE if this build was compiled with OpenCL support.
- `nmathopencl_openc1_device_info()`, `nmathopencl_openc1_fp64_available()` — cached double-precision device selection for kernels.
- `nmathopencl_openc1_reset_device_selection()` — clear device cache.
- `opencltools::get_openc1_core_count()` — compute units on the opencltools default device.

## Host / runtime checks (opencltools)

- `detect_environment_and_gpus()`
- `detect_compute_runtimes()`
- `verify_openc1_runtime()`
- `check_runtime_env()`
- `opencltools::diagnose_glmbyes()`
- `add_to_path_windows()` and related PATH helpers

## References

Stone JE, Gohara D, Shi G (2010). "OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems." *Computing in Science & Engineering*, 12(3), 66–72. doi:10.1109/MCSE.2010.69.

## See Also

`nmathopencl_has_openc1`, `nmathopencl_openc1_device_info`, **opencltools**.

---

norm_rand_opengl	<i>OpenCL-backed RNG Core linkage checks</i>
------------------	--

---

### Description

Linkage wrappers for core RNG primitives used by translated Mathlib paths.

### Usage

```
norm_rand_opengl(n, fallback = FALSE, verbose = FALSE)
```

```
unif_rand_opengl(n, fallback = FALSE, verbose = FALSE)
```

```
r_unif_index_opengl(n, dn, fallback = FALSE, verbose = FALSE)
```

```
exp_rand_opengl(n, fallback = FALSE, verbose = FALSE)
```

### Arguments

n	Number of observations. Non-negative integer scalar.
fallback	When TRUE while <code>nmathopengl_has_opengl()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.
dn	Positive upper bound used by <code>r_unif_index_opengl</code> .

### Value

Numeric vector of length n.

### Examples

```
n <- 5L
if (!nmathopengl_has_opengl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  norm_rand_opengl(n, fallback = FALSE, verbose = TRUE)
  unif_rand_opengl(n, fallback = FALSE, verbose = TRUE)
  r_unif_index_opengl(n, dn = 10, fallback = FALSE, verbose = TRUE)
  exp_rand_opengl(n, fallback = FALSE, verbose = TRUE)
} else {
  stats::rnorm(n)
  stats::runif(n)
  floor(stats::runif(n, min = 0, max = 10))
  stats::rexp(n)
}
```

---

 port\_to\_opencl\_configure

*Port an existing static src/Makevars to use OpenCL configure scripts*


---

## Description

Migrates a package that already has a static committed `src/Makevars` (and optionally `src/Makevars.win`) to the configure-script pattern required for CRAN-safe OpenCL support.

The function renames the existing `src/Makevars` to `src/Makevars.in` (the maintained source template) and generates `configure` (Linux/macOS) and `configure.win` (Windows) scripts that read `src/Makevars.in` at R CMD INSTALL time, run OpenCL detection, and write the final `src/Makevars` with OpenCL flags merged in – or copied verbatim from `src/Makevars.in` for CPU-only builds.

The generated scripts **always succeed**: if no OpenCL SDK is found the package installs cleanly as CPU-only. This is the property that makes packages safe for CRAN submission on build machines without a GPU SDK.

For packages with no existing `src/Makevars`, use `use_opencl_configure` instead. This function is for *migrating* an existing static `Makevars`.

## Usage

```
port_to_opencl_configure(path = ".", backup = TRUE, overwrite = FALSE)
```

## Arguments

<code>path</code>	Character. Root directory of the target package. Defaults to the current working directory ( <code>"."</code> ).
<code>backup</code>	Logical. If TRUE (default), rename <code>src/Makevars</code> to <code>src/Makevars.in</code> before writing the configure scripts. If FALSE, only the configure scripts are written (the existing <code>src/Makevars</code> is left in place).
<code>overwrite</code>	Logical. If TRUE, overwrite existing configure scripts. Defaults to FALSE.

## Value

Invisibly returns a character vector of the file paths written.

## src/Makevars.in workflow

After porting, **maintain** `src/Makevars.in` for your base build flags (OpenMP, **RcppParallel**, LAPACK, etc.). The configure script reads it at install time and appends (or omits) the OpenCL flags. The generated `src/Makevars` is a build artifact – add it to `.gitignore` and never commit it. To update base flags, edit `src/Makevars.in` and reinstall.

**configure -> USE\_OPENCL -> has\_opencl()**

```

configure / configure.win
  -> reads src/Makevars.in for base flags
  -> detects CL/cl.h + libOpenCL (+ runtime probe on Linux)
  -> writes -DUSE_OPENCL into Makevars (or copies .in verbatim)

#ifdef USE_OPENCL in C++ source
  -> guards all GPU code; package compiles cleanly either way

has_opencl() in R
  -> mirrors the compile-time flag; TRUE only if USE_OPENCL was set

```

**Limitations**

- += (append) assignments in src/Makevars are detected and trigger a warning; review the generated configure carefully.
- If src/Makevars looks like a generated file (contains absolute paths, -lOpenCL, or -DUSE\_OPENCL), the function warns. Run on the static committed file, not a build artifact.
- Packages that already have configure or configure.win are refused unless overwrite = TRUE. Users with existing configure scripts should integrate the OpenCL block manually; see `system.file("configure-templates", "README.md", package = "opencltools")`.

**See Also**

[use\\_opencl\\_configure](#) for packages without an existing src/Makevars. `vignette("Chapter-02", package = "nmathopencl")` in **nmathopencl** for a full guide when building on the ported nmath kernel library.

**Examples**

```

##### Start of port_to_opencl_configure example #####

tmp <- tempfile("port_opencl_pkg")
dir.create(tmp)
dir.create(file.path(tmp, "src"))
on.exit(unlink(tmp, recursive = TRUE), add = TRUE)

writeLines(
  c(
    "PKG_CXXFLAGS = $(SHLIB_OPENMP_CXXFLAGS)",
    "PKG_LIBS = $(LAPACK_LIBS) $(BLAS_LIBS) $(FLIBS)"
  ),
  file.path(tmp, "src", "Makevars")
)

written <- port_to_opencl_configure(path = tmp, backup = TRUE)
written
file.exists(file.path(tmp, "configure"))
file.exists(file.path(tmp, "src", "Makevars.in"))

```

```
## Regenerate configure scripts after editing Makevars.in (same temp tree)
written2 <- port_to_openc1_configure(path = tmp, backup = FALSE, overwrite = TRUE)
length(written2)

#####
## End of port_to_openc1_configure example
#####
```

---

ptukey\_openc1

*The Studentized Range Distribution (OpenCL)*


---

## Description

OpenCL-backed distribution and quantile wrappers for the studentized range (Tukey) distribution.

## Usage

```
ptukey_openc1(
  q,
  nmeans,
  df,
  nranges = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

```
qtukey_openc1(
  p,
  nmeans,
  df,
  nranges = 1,
  lower.tail = TRUE,
  log.p = FALSE,
  openc1_parallel = NA,
  fallback = FALSE,
  verbose = FALSE
)
```

## Arguments

q	Numeric vector of quantiles for ptukey_openc1; recycled like stats::ptukey.
nmeans	Number of means in each range (must be >= 2).
df	Degrees of freedom (must be > 0).
nranges	Number of groups whose maxima/minima define the range (must be >= 1).

lower.tail, log.p	Tail/log- $p$ inputs (stats meanings).
openc1_parallel	Dispatch hint (TRUE, FALSE, NA) for $p/q$ wrappers on this page; parallel kernels reserved.
fallback	When TRUE while <code>nmathopenc1_has_openc1()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL. Defaults TRUE temporarily ( <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK.md'</code> ); pass FALSE to surface failures.
verbose	Logical; print fallback/error diagnostics.
p	Numeric vector of probabilities for <code>qtukey_openc1</code> (like <code>stats::qtukey</code> ).

**Value**

Numeric vector result from `ptukey_openc1` or `qtukey_openc1`.

**References**

Copenhaver MD, Holland BS (1988). "Computation of the Distribution of the Maximum Studentized Range Statistic with Application to Multiple Significance Testing of Simple Effects." *Journal of Statistical Computation and Simulation*, **30**(1), 1–15. doi:10.1080/00949658808811082.

Odeh RE, Evans JO (1974). "Algorithm AS 70: The Percentage Points of the Normal Distribution." *Applied Statistics*, **23**(1), 96–97. doi:10.2307/2347061.

**Examples**

```
n <- 1L
if (!nmathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  ptukey_openc1(q = 3.4, nmeans = 5, df = 10, nranges = 1, fallback = FALSE, verbose = TRUE)
  qtukey_openc1(rep(0.8, n), nmeans = 5, df = 10, nranges = 1, fallback = FALSE, verbose = TRUE)
} else {
  stats::ptukey(3.4, nmeans = 5, df = 10, nranges = 1)
  stats::qtukey(rep(0.8, n), nmeans = 5, df = 10, nranges = 1)
}
```

---

 rmultinom\_openc1

*The Multinomial Distribution (OpenCL linkage subset)*


---

**Description**

OpenCL-backed linkage wrapper for multinomial sampling. The current OpenCL kernel supports a 2-category multinomial parameterization via scalar `prob`, and returns a  $2 \times n$  count matrix.

**Usage**

```
rmultinom_openc1(n, size, prob, fallback = FALSE, verbose = FALSE)
```

**Arguments**

n	Number of observations. Non-negative integer scalar.
size	Number of trials per draw (non-negative integer scalar).
prob	Probability of the first category in $[0, 1]$ .
fallback	When TRUE while <code>nmathopengl_has_opengl()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.

**Value**

Integer matrix with 2 rows and n columns.

**Examples**

```
n <- 5L
if (!nmathopengl_has_opengl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  rmultinom_opengl(n, size = 12L, prob = 0.4, fallback = FALSE, verbose = TRUE)
} else {
  stats::rmultinom(n, size = 12L, prob = c(0.4, 0.6))
}
```

---

r\_check\_user\_interrupt\_opengl

*OpenCL-backed R\_ext runtime utility linkage checks*

---

**Description**

Wrappers for utility hooks used by translated R\_ext-dependent kernels.

**Usage**

```
r_check_user_interrupt_opengl(n, fallback = FALSE, verbose = FALSE)
```

**Arguments**

n	Number of observations. Non-negative integer scalar.
fallback	When TRUE while <code>nmathopengl_has_opengl()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). Defaults to FALSE.
verbose	Logical; print fallback/error diagnostics.

**Value**

Numeric vector of length n.

**Known OpenCL limitations**

The stack-check linkage wrapper can fail in device compilation or runtime due to missing host/runtime stack symbols, so it is not exported.

**Examples**

```
n <- 5L
if (!mathopenc1_has_openc1() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  r_check_user_interrupt_openc1(n, fallback = FALSE, verbose = TRUE)
} else {
  as.numeric(seq_len(n))
}
```

---

r\_pow\_openc1

*OpenCL-backed R Math runtime linkage checks*


---

**Description**

Wrappers for low-level Mathlib runtime helpers used by translated kernels.

**Usage**

```
r_pow_openc1(x, y, fallback = FALSE, verbose = FALSE)
```

```
r_pow_di_openc1(x, n_exp, fallback = FALSE, verbose = FALSE)
```

```
log1pmx_openc1(x, fallback = FALSE, verbose = FALSE)
```

```
log1pexp_openc1(x, fallback = FALSE, verbose = FALSE)
```

```
log1mexp_openc1(x, fallback = FALSE, verbose = FALSE)
```

```
lgamma1p_openc1(x, fallback = FALSE, verbose = FALSE)
```

```
pow1p_openc1(x, y, fallback = FALSE, verbose = FALSE)
```

```
logspace_add_openc1(logx, logy, fallback = FALSE, verbose = FALSE)
```

```
logspace_sub_openc1(logx, logy, fallback = FALSE, verbose = FALSE)
```

```
logspace_sum_openc1(logx, logy, fallback = FALSE, verbose = FALSE)
```

**Arguments**

- x            Numeric vector(s): primary input, recycled together like stats family functions. Length-zero returns `numeric(0)`.
- y            Numeric vector for `r_pow_openc1` and `pow1p_openc1` (recycled against x).

fallback	When TRUE while <code>nmathopengl_has_opengl()</code> reports OpenCL present, recover with CPU if the OpenCL call fails. Ignored when the runtime reports no OpenCL (CPU path is chosen automatically). For <code>log1pmx_opengl</code> , <code>lgamma1p_opengl</code> , <code>pow1p_opengl</code> , and the <code>logspace_*</code> wrappers, defaults to TRUE temporarily while <code>pgamma_utils</code> -stitching kernels are stabilized; see <code>'inst/OPENCL_PGAMMA_UTILS_KERNEL_FALLBACK'</code> .
verbose	Logical; print fallback/error diagnostics.
n_exp	Integer vector for <code>r_pow_di_opengl</code> , recycled against <code>x</code> .
logx, logy	Numeric vectors for log-space combination helpers (recycled together).

### Details

On the GPU path, arguments are recycled to a common length `len` (maximum argument length, R recycling rules). Each output index runs one scalar kernel launch.

### Value

Numeric vector of the recycled common length (see Details).

### Examples

```
if (!nmathopengl_has_opengl() || identical(Sys.getenv("NOT_CRAN"), "true")) {
  r_pow_opengl(x = 1.2, y = 2, fallback = FALSE, verbose = TRUE)
  r_pow_di_opengl(x = 1.2, n_exp = 3L, fallback = FALSE, verbose = TRUE)
  log1pmx_opengl(x = 0.2, fallback = FALSE, verbose = TRUE)
  log1pexp_opengl(x = 0.2, fallback = FALSE, verbose = TRUE)
  log1mexp_opengl(x = 0.5, fallback = FALSE, verbose = TRUE)
  lgamma1p_opengl(x = 0.2, fallback = FALSE, verbose = TRUE)
  pow1p_opengl(x = 0.2, y = 3, fallback = FALSE, verbose = TRUE)
  logspace_add_opengl(logx = -2, logy = -3, fallback = FALSE, verbose = TRUE)
  logspace_sub_opengl(logx = -2, logy = -3, fallback = FALSE, verbose = TRUE)
  logspace_sum_opengl(logx = -2, logy = -3, fallback = FALSE, verbose = TRUE)
  log1pmx_opengl(x = seq(-0.5, 0.5, by = 0.25), fallback = FALSE, verbose = TRUE)
} else {
  (1.2 + seq_len(1L) * 1e-3)^2
  1.2^3
  log1p(0.2) - 0.2
  ifelse(0.2 > 0, 0.2 + log1p(exp(-0.2)), log1p(exp(0.2)))
  ifelse(0.5 <= log(2), log(-expm1(-0.5)), log1p(-exp(-0.5)))
  lgamma(1.2)
  exp(3 * log1p(0.2))
  {
    m <- max(-2, -3)
    m + log1p(exp(min(-2, -3) - m))
  }
  -2 + log1p(-exp(-3 - (-2)))
  {
    m <- max(-2, -3)
    m + log1p(exp(min(-2, -3) - m))
  }
  {
    xv <- seq(-0.5, 0.5, by = 0.25)
  }
```

```

    log1p(xv) - xv
  }
}

```

---

```
stage_kernel_dependency_sort
```

*Stage Kernel Library Dependency Sort Results*

---

## Description

Run the file-level @depends sort without requiring full success. Files that can be sorted are copied in order, and files blocked by unresolved dependencies are copied into a separate folder with CSV reports.

## Usage

```
stage_kernel_dependency_sort(library_dir, output_dir, overwrite = FALSE)
```

## Arguments

library_dir	Directory containing .cl files with @depends tags.
output_dir	Directory where sorted and unresolved files/reports should be written.
overwrite	Logical; remove and recreate output_dir if it exists.

## Value

A named list with:

sorted Data frame of files placed in dependency order (order, pass, file, source\_type, depends).

unresolved Data frame of files blocked by missing dependencies.

sorted\_dir, unresolved\_dir Output directory paths; CSV reports 'sorted\_files.csv' and 'unresolved\_files.csv' are written under output\_dir.

## Examples

```
##### Start of stage_kernel_dependency_sort example #####

lib_dir  <- system.file("cl/ex_glmbyes_nmath", package = "opencltools")
output_dir <- tempfile("stage_sort")
on.exit(unlink(output_dir, recursive = TRUE), add = TRUE)

res <- stage_kernel_dependency_sort(lib_dir, output_dir, overwrite = TRUE)
nrow(res$sorted)
length(list.files(output_dir, recursive = TRUE))

#####
## End of stage_kernel_dependency_sort example
#####
```

---

use\_opengl\_configure *Set up OpenCL configure scripts in a downstream R package*

---

### Description

Copies generic OpenCL configure and configure.win scripts to the root directory of a package. The scripts detect CL/cl.h and libOpenCL at compile time and generate src/Makevars (Linux / macOS) or src/Makevars.win (Windows) with or without -DUSE\_OPENCL, depending on what is found.

The scripts **always succeed**. When no OpenCL SDK is present they produce a CPU-only Makevars with no -lOpenCL. This is the key property that makes packages safe for CRAN submission without requiring a GPU SDK on the build machine.

### Usage

```
use_opengl_configure(path = ".", overwrite = FALSE)
```

### Arguments

path	Character. Root directory of the target package. Defaults to the current working directory (".").
overwrite	Logical. If TRUE, overwrite existing configure scripts. Defaults to FALSE to avoid accidentally replacing a customized script.

### Value

Invisibly returns a character vector of the file paths that were written (empty if all files were skipped).

### Why configure scripts are necessary

A package that references -lOpenCL or CL/cl.h in a static src/Makevars will **fail to compile** on 'CRAN' build machines (which have no GPU SDK installed), and no binary will be produced. The configure scripts here avoid this by probing for the SDK at install time and falling back to a CPU-only build when it is absent. The relationship is:

```
configure / configure.win
  -> detects CL/cl.h + libOpenCL
  -> writes -DUSE_OPENCL into Makevars (or omits it)

#ifdef USE_OPENCL in C++ source
  -> guards all GPU code; package compiles cleanly either way

has_opengl() in R
  -> mirrors the compile-time flag; returns TRUE only if USE_OPENCL was set
```

**See Also**

[port\\_to\\_openc1\\_configure](#) for packages with an existing static src/Makevars. [openc1toolsLdFlags](#) for linking against this package from C++. [vignette\("Chapter-02", package = "nmathopenc1"\)](#) in **nmathopenc1** for a full downstream package guide when using the ported nmath kernel library. Template source: `system.file("configure-templates", package = "openc1tools")`.

**Examples**

```
##### Start of use_openc1_configure example #####

tmp <- tempfile("use_openc1_pkg")
dir.create(tmp)
on.exit(unlink(tmp, recursive = TRUE), add = TRUE)

written <- use_openc1_configure(path = tmp)
written
file.exists(file.path(tmp, "configure"))
file.exists(file.path(tmp, "configure.win"))

## Overwrite path (same temp tree; safe when run.dontest = TRUE)
written2 <- use_openc1_configure(path = tmp, overwrite = TRUE)
length(written2)

#####
## End of use_openc1_configure example
#####
```

---

write\_kernel\_dependency\_index

*Build and save a kernel dependency index*

---

**Description**

Writes two companion index files next to the .cl files in the kernel library directory:

**Usage**

```
write_kernel_dependency_index(
  library_dir = NULL,
  tags = NULL,
  output_path = NULL,
  write = TRUE,
  verbose = FALSE
)
```

**Arguments**

library_dir	Library root with annotated .cl files. When tags is supplied, must agree with tags\$library_dir.
tags	Optional attachment result from <a href="#">attach_kernel_dependency_tags</a> . Must have ok = TRUE when reused to skip resorting.
output_path	Path for the RDS file. Defaults to file.path(<library_dir>, "kernel_dependency_index.rds"). The .tsv file is always written to the same directory with the .tsv extension.
write	If FALSE, builds the index object and returns it without writing either file.
verbose	If TRUE, emits short messages with the output paths.

**Details**

- RDS helper shard for loaders such as [load\\_library\\_for\\_kernel](#).
- Tab-separated (.tsv) stem map for C++ (rows stem<TAB>dependencies, pre-sorted).

Both files encode the same information and are always written together so they remain in sync.

**Value**

Invisibly, the index `list()` (version schema):

- version: integer schema version.
- generated\_at: timestamp from `Sys.time()`.
- library\_dir, library\_name: resolved library path / basename.
- stems\_ordered: stems in global load order.
- load\_order: named integer vector stem -> rank.
- depends: named list stem -> character() (direct @depends).
- all\_depends: named list stem -> character() (transitive dependencies, order consistent with global load order).
- n\_files: file count.

**See Also**

[load\\_library\\_for\\_kernel](#)

[extract\\_library\\_subset](#)

Other OpenCL kernel library subsets: [extract\\_library\\_subset\(\)](#), [load\\_library\\_for\\_kernel\(\)](#)

**Examples**

```
##### Start of write_kernel_dependency_index example #####
```

```
lib_dir <- system.file("cl/ex_glmbyes_nmath", package = "opencltools")
idx <- write_kernel_dependency_index(library_dir = lib_dir, write = FALSE)
names(idx)
```

```
length(idx$stems_ordered)
idx$n_files
```

```
#####
## End of write_kernel_dependency_index example
#####
```

# Index

- \* **diagnostics**
    - nmathopencil\_has\_opencil, 64
  - \* **environment**
    - nmathopencil\_has\_opencil, 64
  - \* **gpu**
    - nmathopencil\_has\_opencil, 64
  - \* **opencil**
    - nmathopencil\_has\_opencil, 64
- add\_to\_path\_windows, 65
- attach\_cross\_library\_tags, 5, 8, 9, 62
- attach\_kernel\_call\_tags, 8
- attach\_kernel\_dependency\_tags, 6, 9, 10, 77
- beta\_opencil (gammafn\_opencil), 58
- check\_runtime\_env, 65
- choose\_opencil (gammafn\_opencil), 58
- dbeta\_opencil, 11
- dbinom\_opencil (dbinom\_raw\_opencil), 14
- dbinom\_raw\_opencil, 14
- dcauchy\_opencil, 16
- dchisq\_opencil, 18
- detect\_compute\_runtimes, 65
- detect\_environment\_and\_gpus, 65
- dexp\_opencil, 20
- df\_opencil, 22
- dgamma\_opencil, 24
- dgeom\_opencil, 27
- dhyper\_opencil, 28
- digamma\_opencil (gammafn\_opencil), 58
- dlnorm\_opencil, 31
- dlogis\_opencil, 33
- dnbeta\_opencil (dbeta\_opencil), 11
- dnbinom\_mu\_opencil (dnbinom\_opencil), 35
- dnbinom\_opencil, 35
- dnorm\_opencil, 37
- dpois\_opencil (dpois\_raw\_opencil), 40
- dpois\_raw\_opencil, 40
- dt\_opencil, 42
- dunif\_opencil, 44
- dweibull\_opencil, 46
- Ex\_EnvelopeEval, 4, 5, 51, 56
- Ex\_EnvelopeSize, 52, 54
- Ex\_glmb\_Standardize\_Model, 57
- Ex\_glmbfamfunc, 56
- exp\_rand\_opencil (norm\_rand\_opencil), 66
- extract\_library\_subset, 48, 63, 77
- family, 56
- fmax2\_opencil (imax2\_opencil), 60
- fmin2\_opencil (imax2\_opencil), 60
- fprec\_opencil (imax2\_opencil), 60
- fround\_opencil (imax2\_opencil), 60
- fsign\_opencil (imax2\_opencil), 60
- ftrunc\_opencil (imax2\_opencil), 60
- gammafn\_opencil, 58
- glmbayesEnvelopeExample, 60
- gpu\_diagnostics
  - (nmathopencil\_has\_opencil), 64
- imax2\_opencil, 60
- imin2\_opencil (imax2\_opencil), 60
- lbeta\_opencil (gammafn\_opencil), 58
- lchoose\_opencil (gammafn\_opencil), 58
- lgamma1p\_opencil (r\_pow\_opencil), 72
- lgammafn\_opencil (gammafn\_opencil), 58
- load\_kernel\_library, 62
- load\_library\_for\_kernel, 6, 50, 62, 77
- log1mexp\_opencil (r\_pow\_opencil), 72
- log1pexp\_opencil (r\_pow\_opencil), 72
- log1pmx\_opencil (r\_pow\_opencil), 72
- logspace\_add\_opencil (r\_pow\_opencil), 72
- logspace\_sub\_opencil (r\_pow\_opencil), 72
- logspace\_sum\_opencil (r\_pow\_opencil), 72

- nmathopenc1 (nmathopenc1-package), 3
- nmathopenc1-package, 3
- nmathopenc1\_has\_openc1, 4, 5, 12, 15, 17, 19, 21, 23, 25, 28, 30, 32, 34, 36, 39, 41, 43, 45, 47, 59, 61, 64, 64, 65, 66, 70, 71, 73
- nmathopenc1\_openc1\_device\_info, 4, 5, 65
- nmathopenc1\_openc1\_device\_info (nmathopenc1\_has\_openc1), 64
- nmathopenc1\_openc1\_fp64\_available, 4, 65
- nmathopenc1\_openc1\_fp64\_available (nmathopenc1\_has\_openc1), 64
- nmathopenc1\_openc1\_reset\_device\_selection, 65
- nmathopenc1\_openc1\_reset\_device\_selection (nmathopenc1\_has\_openc1), 64
- norm\_rand\_openc1, 66
- openc1tools::diagnose\_glmbyes, 4, 64, 65
- openc1tools::get\_openc1\_core\_count, 65
- openc1tools::load\_kernel\_library, 5
- openc1toolsLdFlags, 76
- packageStartupMessage, 4
- pbeta\_openc1 (dbeta\_openc1), 11
- pbinom\_openc1 (dbinom\_raw\_openc1), 14
- pcauchy\_openc1 (dcauchy\_openc1), 16
- pchisq\_openc1 (dchisq\_openc1), 18
- pentagamma\_openc1 (gammafn\_openc1), 58
- pexp\_openc1 (dexp\_openc1), 20
- pf\_openc1 (df\_openc1), 22
- pgamma, 25
- pgamma\_openc1, 25
- pgamma\_openc1 (dgamma\_openc1), 24
- pgeom\_openc1 (dgeom\_openc1), 27
- phyper\_openc1 (dhyper\_openc1), 28
- plnorm\_openc1 (dlnorm\_openc1), 31
- plogis\_openc1 (dlogis\_openc1), 33
- pnbinom\_mu\_openc1 (dnbinom\_openc1), 35
- pnbinom\_openc1 (dnbinom\_openc1), 35
- pnorm, 39
- pnorm\_openc1, 25
- pnorm\_openc1 (dnorm\_openc1), 37
- port\_to\_openc1\_configure, 67, 76
- pow1p\_openc1 (r\_pow\_openc1), 72
- ppois\_openc1 (dpois\_raw\_openc1), 40
- print methods, 63
- print(), 11
- print.Ex\_glmfamfunc (Ex\_glmfamfunc), 56
- printing methods, 50, 63
- psigamma\_openc1 (gammafn\_openc1), 58
- pt\_openc1 (dt\_openc1), 42
- ptukey\_openc1, 69
- punif\_openc1 (dunif\_openc1), 44
- pweibull\_openc1 (dweibull\_openc1), 46
- qbinom\_openc1 (dbinom\_raw\_openc1), 14
- qcauchy\_openc1 (dcauchy\_openc1), 16
- qchisq\_openc1 (dchisq\_openc1), 18
- qexp\_openc1 (dexp\_openc1), 20
- qf\_openc1 (df\_openc1), 22
- qgeom\_openc1 (dgeom\_openc1), 27
- qhyper\_openc1 (dhyper\_openc1), 28
- qlnorm\_openc1 (dlnorm\_openc1), 31
- qlogis\_openc1 (dlogis\_openc1), 33
- qnbinom\_mu\_openc1 (dnbinom\_openc1), 35
- qnbinom\_openc1 (dnbinom\_openc1), 35
- qnorm\_openc1 (dnorm\_openc1), 37
- qpois\_openc1 (dpois\_raw\_openc1), 40
- qt\_openc1 (dt\_openc1), 42
- qtukey\_openc1 (ptukey\_openc1), 69
- qunif\_openc1 (dunif\_openc1), 44
- qweibull\_openc1 (dweibull\_openc1), 46
- r\_check\_user\_interrupt\_openc1, 71
- r\_pow\_di\_openc1 (r\_pow\_openc1), 72
- r\_pow\_openc1, 72
- r\_unif\_index\_openc1 (norm\_rand\_openc1), 66
- rbeta\_openc1 (dbeta\_openc1), 11
- rbinom\_openc1 (dbinom\_raw\_openc1), 14
- rcauchy\_openc1 (dcauchy\_openc1), 16
- rchisq\_openc1 (dchisq\_openc1), 18
- rexp\_openc1 (dexp\_openc1), 20
- rf\_openc1 (df\_openc1), 22
- rgamma\_openc1 (dgamma\_openc1), 24
- rgeom\_openc1 (dgeom\_openc1), 27
- rhyper\_openc1 (dhyper\_openc1), 28
- rlnorm\_openc1 (dlnorm\_openc1), 31
- rlogis\_openc1 (dlogis\_openc1), 33
- rmultinom\_openc1, 70
- rnbinom\_mu\_openc1 (dnbinom\_openc1), 35
- rnbinom\_openc1 (dnbinom\_openc1), 35
- rnorm\_openc1, 39
- rnorm\_openc1 (dnorm\_openc1), 37

rpois\_opengl (dpois\_raw\_opengl), 40  
rt\_opengl (dt\_opengl), 42  
runif\_opengl (dunif\_opengl), 44  
rweibull\_opengl (dweibull\_opengl), 46  
  
sign\_opengl (imax2\_opengl), 60  
stage\_kernel\_dependency\_sort, 74  
Sys.time(), 77  
  
tetragamma\_opengl (gammafn\_opengl), 58  
trigamma\_opengl (gammafn\_opengl), 58  
  
unif\_rand\_opengl (norm\_rand\_opengl), 66  
use\_opengl\_configure, 67, 68, 75  
  
verify\_opengl\_runtime, 65  
  
warning, 63  
write\_kernel\_dependency\_index, 6, 49, 50,  
62, 63, 76