

flucoma-core Release Notes

Changes in main since production

Comparison: production (da9e6ac, last real commit 31 Aug 2025, v1.0.9) against main (0626d7a, 11 Apr 2026). I counted 34 non-merge commits, 55 files changed, about +1,850 / -390 lines. production has nothing that main lacks except merge commits.

Highlights

- **Random seeding across the stochastic algorithms.** Results can now be reproduced. A new seed parameter defaults to -1, which means a random seed as before.
- **Better K-Means and SKMeans initialisation.** There are three init methods to choose from.
- **Bug fixes** for KNNRegressor loading and MLP training.
- **Dependency cleanup:** `tl:optional` is gone, and a build break from `fmt`'s renamed default branch is fixed.

New features

Settable random seed

- Added a seed parameter to these clients:
 - KMeans
 - SKMeans
 - MLPRegressor
 - MLPClassifier
 - UMAP
 - NMF (BufNMF)
 - NMFMatch
- NMFCross and the NNDSVD seeding code were also changed to take a seed.
- RTPGHI and Griffin-Lim phase initialisation now use the seed too.
- The mechanism is a new `EigenRandom` helper (`algorithms/util/EigenRandom.hpp`). It generates random Eigen containers from the C++ `<random>` library instead of `Eigen::Random()`.
- MLP and `SimpleDataSampler` shuffling is now seedable as well. Per the diff, a positive seed gives repeatable batches, and anything else is non-deterministic.
- New repeatability tests cover each of these algorithms.

K-Means and SKMeans initialisation

- A new `initMethod` enum parameter takes one of three values:
 - Random Assignment (the default)
 - Random Points (Forgy initialisation)
 - Sampling (an MCMC approximation of k-means++ called AFK-MC², using squared Euclidean distance for K-Means and squared cosine distance for SKMeans)
- Training now calls `computeMeans` after assignment, so the update order inside the loop is different.
- SKMeans avoids some matrix copies by using `Eigen::Ref` and `noalias()`.

Build

- `CMAKE_MSVC_RUNTIME_LIBRARY` is now a cache variable, so wrapper projects can override the /MT default (#340, from Eduard Müller / taktik).

Resources

- A rays test dataset was added to `Resources/Data` (four evenly spread clusters centred on zero).

Fixes

- **KNNRegressor:** JSON loading validated the wrong key (`labels` as strings). It now checks `target` as numeric data (#324, fixed in #325).
- **MLP:** if training fails, the model is restored to its exact pre-call state. Before, it was cleared.
- **K-Means:** the "empty cluster" warning used to print to `std::cout` from inside the algorithm. It is now a proper warning result from the client: "There were empty clusters; perhaps `numClusters` is too high."
- **CI:** the temporary `LoudnessClient` change made for a Windows/Max CI investigation was reverted. There is no net difference in the code.

Internal and breaking changes

These affect people who use the C++ headers directly, not Max, Pd or SuperCollider users.

- **Changed signatures** (the new arguments are mostly defaulted or trailing):
 - `KMeans::train(..., InitMethod, seed)`
 - `NMF::processFrame(..., randomSeed, Allocator&)`
 - the `RTPGHI` constructor now takes `seed` before `Allocator&`
 - `MLP::train`
 - `GriffinLim`
 - `SGD`
 - `UMAP::train` and `transform`
- **Removed:**
 - `data/FluidFile.hpp`, which had no remaining includes in `include/`
 - the `tl_optional` dependency (`Optional<T>` is now `std::optional<T>`)
- **fmt:** `FetchContent` now tracks the main branch instead of master, which no longer exists upstream (#342, Nicolas Bouillot).
- **Tests:** new or expanded tests for K-Means, SKMeans, MLP, NMF, NMFCross, NNDSVD, UMAP, RTPGHI, GriffinLim, SGD, EigenRandom and DataSampler. The test CMake was updated, including registering the previously unregistered KMeans test.
- **Style:** clang-format cleanups in the touched files.

Things to flag when releasing

- Results for a given seed on a given platform should be reproducible. They will differ from older versions' random output, because the random number sources changed.
- Existing patches or scripts that hardcode client parameter counts or positions may need updating, since `seed` (and `initMethod` for K-Means and SKMeans) were added to the end of the parameter lists.

- The NMFCross, NNDSVD, NMF Morph and NMFFilter changes were summarised from the diffs without reading every line.

Contributors: Owen Green, Pierre Alexandre Tremblay, Nicolas Bouillot, Eduard Müller.