LINDO API 17.0 - Release Notes

Changes from version 16.0

1. LP (Linear Program) Solvers

- Enhancements to the simplex and barrier solvers make a typical linear program solve about 2.4 times faster than in version 16.0, and cut the total time over our standard LP benchmark set by 28%.
- A new symmetry-based reduction folds a linear program into a smaller equivalent problem, solves that, and expands the solution back. Either a basic or an interior-point solution can be returned, and alternative optimal bases of the reduced problem can be reported.

2. MIP (Mixed Integer Program) Solvers

Measured over our standard benchmark suite of mixed integer models:

- 9% more models were solved to proven optimality than in the previous version, and 9% more were given a feasible solution.
- Where both versions prove optimality, version 17.0 gets there 17% faster.
- A new family of cuts is generated at the root node, enabled by default.
- Two new heuristics look for integer feasible solutions early: one solves small 0-1 subproblems with a dedicated pseudo-Boolean engine, the other explores the integer points nearest the relaxed solution.
- A solution supplied by the caller through a callback now reaches the solver's internal model, so it can be adopted as a starting incumbent.

3. NLP (Nonlinear Program) Solvers

Measured over our standard benchmark suite of non-convex continuous nonlinear models:

- Models abandoned as locally infeasible fell by 80%, from 15 to 3.
- Models solved to local optimality rose by 5%, and 4% more models finished with a usable solution.
- This costs some extra time: a typical model now takes under three seconds where version 16.0 took under two. That additional effort is what recovers the models the previous version gave up on.
- The multistart search gained several strategies, which together account for most of the recovery above.
- The automatic second pass and tighter global bounds are now part of the default configuration. Models that already find a feasible point on the first pass are unaffected.
- Three new functions are available to model builders: binary entropy, a logit term, and the internal rate of return of a series of cash flows.

4. Global Solvers

Measured over our standard benchmark suite of mixed integer nonlinear models:

- 27% fewer models finished with no verdict and 3% more finished with a usable solution. The number proven globally optimal rose by 2%, and the number proven infeasible doubled.
- Total solution time over the suite fell by 2%, because far fewer models now run to the time limit. The broader search does cost extra time on models the previous version already proved optimal.
- Solving the same suite with the multistart search instead: at default settings 23% fewer models were abandoned as locally infeasible, 18% fewer finished with no verdict and 6% more reached local optimality, with the average remaining gap down from 2.6% to 1.4% and the total time down 7%. With the multistart search enabled, 16% fewer models were abandoned as locally infeasible and 3% more reached local optimality.

5. Platforms

- Windows 32-bit and 64-bit on x86, Linux 64-bit on x64 and on ARM, macOS 64-bit on x86 and on Apple silicon.

Note on upgrading

- The reserved macro `LS_SOLVER_PREP_RES256` is now named `LS_SOLVER_PREP_FLIPGE`. Its value is unchanged.