

# BI two dimensional search and Grid mode

---

*LISE\*\* 18.8.82 | 16 October 2026*

Grid (PID) mode supports the search for low-count peaks whose centers form an approximately uniform grid. This guide describes the 2D workflow, recovery criteria, safeguards, output files, and release validation. Recovered low-count peaks are provisional candidates and should be checked against the original events. Version 18.8.75 also provides optional single-event grid annotations, kept separate from accepted peaks. The release sections include the 18.8.76 output columns and rounding rules and the 18.8.80 correction to result plotting. Version 18.8.82 adds Counts from original bins in common grid ellipses, global sigmas weighted by  $\sqrt{\text{Area}}$ , a Bound factor control and an information dialog. Results now use the .csv extension introduced in 18.8.81.

Extends the optional Grid mode introduced in 18.8.70.

## Changes introduced in version 18 8 72

- Added Guess sigma values for the original input within selected limits. It measures resolved individual profiles, uses a median width, and removes the variance introduced by its short smoothing kernel. It does not confuse full-spectrum peak spacing with individual peak width. Suggestions remain editable; insufficient evidence leaves the previous values intact.
- Reduced the search settings panel font by one point (9 to 8 pt).
- Added raw-count recovery, local widths, confidence labels and original-data inspection, as specified below. The background workflow and input cache apply to 2D analysis; 1D keeps its established search workflow.
- Fixed zero-based line-count handling: minimum rows = 1 now allows one row, and the .bir nb column reports cur + 1. One/two-row profiles are not dropped merely because a Y-splitting operation has no interior maximum.
- Bounded smoothing by available profile length; removed per-line temporary smoothing allocations. Added peak-capacity guards and corrected delete[[]].
- Corrected the sigma fallback to use the active axis. Zero sigma-relation now disables the width cut safely; finite numeric input is required and valid sigma-relation values lie between zero and one. Area cuts are inclusive.
- Preserved full-precision input in the new pipeline instead of converting it to signed 16-bit counts. Added input/output path separation checks, atomic individual output commits, and local-footprint Gaussian rasterization.

## Validation of version 18 8 72

The full MSVC application build and standard regression tests passed before the version-only increment from 18.8.71 to 18.8.72. This release commit does not include generated executables. The counter change requires a rebuild to appear in the installed executable.

Synthetic tests verify recovery of exactly two events at an omitted grid site, rejection of an isolated off-grid pair, background screening, minimum-area cuts, cancellation, cropped input, tall matrices and oversized smoothing radii. The sigma test checks a known-width Gaussian and insufficient-structure input. The grid test includes a low-weight on-grid point.

With Z\_A3Q\_MSU.spa, full search limits, sigma X=0.15 / Y=0.13, minimum rows=1, minimum area=0.1, Grid on and tolerance=25%, the regression run returned 158 strong and 20 weak grid-supported candidates, plus 80

empty grid sites. Analysis of the already-loaded matrix took approximately 0.05-0.09 seconds on the development machine; this excludes file loading, rendering and output. These counts are regression observations, not a ground-truth accuracy claim. Interactive GUI automation was not performed. The optional AddressSanitizer executable exited without a usable diagnostic, so that check is inconclusive.

Search uses the established BI line detector. Grid mode then fits strong detected centers to a skewed lattice and examines raw data at unoccupied lattice sites. The original spectrum is never changed. The new pipeline applies to 2D spectra; the 1D search retains its existing workflow.

## Quick workflow

1. Open the original spectrum and choose search limits.
2. Click Guess sigma values, then review the suggested X/Y widths.
3. Enable Grid (PID) mode when peak centers form an approximately uniform grid.
4. Search. Cancel stops processing without replacing previous output files.
5. Open Overlay / diagnostics. Click an ellipse for the reason and measurements. Mouse wheel zooms; dragging pans; Fit restores the full view.

## Confidence and screening

- Strong: passed BI cuts, at least 25 raw counts within a two-sigma aperture, and local background screening tail below 0.001.
- Weak grid-supported: a low-count detection consistent with the grid, or raw data recovered at an unoccupied grid site. These are provisional candidates.
- Uncertain: failed a cut or needs review. In Grid mode these are not exported as accepted peaks. With Grid off, accepted legacy detections can be uncertain and remain exported, to preserve normal BI behavior.
- Missing: an expected grid site with zero counts in its recovery aperture.

Recovery uses the median local 2D widths of up to seven nearby strong grid detections. The aperture has radius two in normalized X/Y sigma coordinates; background is estimated from an annulus between three and four sigma, excluding four-sigma regions around neighboring strong grid anchors. Excluded bins are removed from both counts and sampled area. At least eight usable bins and one quarter of the original annulus must remain. An empty annulus has a 0.25-count floor before scaling by the aperture/annulus bin-area ratio. Recovery requires at least two raw counts, the selected minimum area, a local screening tail  $\leq 0.05$ , and a measured center within the grid tolerance. This is a local Poisson screening score with an estimated background, NOT a calibrated discovery significance or a correction for searching many sites. Use original, unweighted event-count spectra when interpreting these scores.

Recovery area is the measured two-sigma aperture count, or that count minus the estimated background when subtraction is selected. It is not extrapolated to an infinite Gaussian area. Legacy detections retain their BI-derived areas. The diagnostic CSV separates area, raw aperture counts, and local background.

## Limits and safeguards

- At least six strong peaks must support a lattice. If the fit fails, standard BI detections are kept and a status message explains the fallback.
- Only sites within the search limits, one grid cell beyond the anchor bounds, and within 2.5 grid steps of an anchor are inspected. At most 10000 sites.

- Legacy minimum-row, smoothing, and width cuts apply to legacy detection. Grid recovery independently inspects raw counts, so those cuts do not erase an otherwise supported two-event locus. Minimum area still applies.
- Sigma relation must lie in [0,1]; zero disables the global width cut.
- Candidate limits are reported. Large matrices are limited to 16 million bins.

## Files and performance

Current results use the .csv extension. Standard 2D results retain ten columns. Grid mode adds Counts after Area and  $x_l$  and  $y_l$  after  $y_{mean}$ . Counts sums original bins inside common ellipses centered on fitted grid sites, using the global widths weighted by  $\sqrt{\text{Area}}$  and the selected Bound factor. Floating-point fields use scientific notation with three digits after the decimal point. The Name field records confidence. A companion .csv.diagnostics.csv includes every inspected candidate and its reason. The release sections below specify the center fallback, boundaries and weighted input behavior. Older sections retain the historical .bir extension and earlier column layouts.

2D reading, sigma estimation and detection run in a background worker with progress and cancellation. Search and Guess sigma share a full-precision cache keyed by canonical filename, size, modification time, and input format. Output Gaussian rasterization visits only local footprints around each peak. Original-data overlays retain noninteger and high-count bins without the old 16-bit clipping. Legacy standalone spectrum plotting still has its older integer display limitations; the new overlay uses the full-precision cache.

Tests: tests\bi\_recovery\_test.cmd [path-to-Z\_A3Q\_MSU.sp] Includes two-event recovery, off-grid noise, background rejection, minimum-area cuts, cancellation, cropped regions, tall matrices and short-profile smoothing.

## Local widths and neighboring peak background in version 18 8 73

The sigma suggestion now uses individual profile cuts as seeds for local 2D patches. Populated patches are measured once, with iterative marginal moments, an edge-median pedestal, and a Gaussian truncation correction. At least three resolved patches are required to replace the profile-based fallback estimate. Search limits select the region used for the suggestion; a full-spectrum median is not necessarily representative of every region. Sparse loci are not fitted individually: recovery borrows the median of up to seven nearby anchor widths, now remeasured from populated local 2D patches bounded by grid spacing.

Recovery background excludes four-sigma ellipses around neighboring strong grid anchors. Excluded bins are removed from both counts and background area. At least eight usable bins and one quarter of the original annulus must remain; otherwise the candidate stays uncertain. The 0.25-count floor, minimum of two events, minimum area, center tolerance and local screening threshold are retained. Only established neighboring peaks are masked, not arbitrary nonzero bins.

Regression coverage includes tilted Gaussian marginal widths, sparse-width fallback, contaminated background, area-normalized uniform background, and insufficient unmasked area. With Z\_A3Q\_MSU.sp and the reported settings (sigma  $X=0.127948$ ,  $Y=0.0701201$ , minimum rows=1, area=0.1, tolerance=25%), the loci near (-14.0954,25.9702) and (-14.0587,25.0334) are both recovered as provisional weak grid-supported candidates, containing 12 and 2 aperture counts. The real-data test also requires both loci to survive a search using the revised full-spectrum sigma suggestion. No discovery-significance claim is implied.

## Validation of version 18 8 73

- Original screenshot settings: X limits [-19,3], Y limits [6,32], sigma X=0.127948, Y=0.0701201, minimum rows=1, area=0.1, Grid enabled, tolerance=25%, and background subtraction disabled.
- Upper locus: center (-14.0954,25.9702), 12 raw aperture counts, estimated background 0.909091, local screening tail 2.88038e-10. Before the correction the estimated background was 30.1828 and the candidate was rejected. It now passes as weak\_grid, still provisional.
- Lower locus: center (-14.0587,25.0334), 2 raw aperture counts, estimated background 0.181818, local screening tail 0.0146556. Before the correction the estimated background was 4.031 and the candidate was rejected. It now passes as weak\_grid, still provisional.
- Revised full-spectrum sigma suggestion: X=0.121453, Y=0.060830, using 77 populated patches. Both loci also pass a search using these values.
- Regional suggestion for X=[-14.8,-11.5], Y=[24,30]: X=0.118864, Y=0.099593, using 6 populated patches. Local widths need not equal the full-spectrum median; the correction does not simply double every sigma.
- The recovery regression suite and the original profile-sigma tests passed. Uniform background and isolated off-grid counts remain rejected. Too little unmasked background area gives an uncertain result rather than acceptance.
- The complete MSVC application build passed before the version-only change from 18.8.72 to 18.8.73. Generated executables and sample output files are not included in this source commit. Rebuild to display the new counter. Interactive GUI automation was not performed.

## Single event grid candidates in version 18 8 75

The Include single-event grid candidates checkbox defaults to off and is only available for 2D spectra with Grid mode enabled. A reliable grid must already exist. At an unoccupied grid site, exactly one unit-count bin inside the borrowed two-sigma aperture can be annotated if its position meets the grid tolerance and sufficient unmasked surrounding bins remain for a background estimate. Fractional weights that sum to one are not treated as a single event.

These are positional annotations, not accepted peaks or statistical detections. Their widths come from neighboring strong peaks; they are not fitted from one event. The local background and screening tail are still calculated and shown. Minimum-area and accepted-peak screening cuts do not apply to these annotations; this is stated explicitly in their diagnostic reason, not a silent acceptance.

The overlay uses filled purple diamonds, independently controlled by Show single-event candidates. The status line reports their count separately. The .bir.diagnostics.csv export uses status single\_event\_candidate with accepted=0 and recovered=1. They are excluded from the accepted .bir file, accepted-peak totals and fitted output spectrum. Turning the option off restores the usual uncertain/missing classification without changing accepted results.

Regression tests cover opt-in/off behavior, separate CSV status, unchanged accepted results, off-grid rejection, fractional-count rejection, missing-grid fallback, Grid-off behavior and continued two-event recovery. The option does not identify an isotope or establish the existence of a peak from one event. On Z\_A3Q\_MSU.spa with the revised full-spectrum sigma suggestion, the regression run reported 9 single-event annotations and identical accepted-peak flags with the option on and off. These are candidate counts, not validated identifications.

## Validation of version 18 8 75

The recovery regression suite and original profile-sigma tests passed. The single-event tests explicitly check default-off behavior, one on-grid candidate, no change in accepted flags or centers, CSV status `single_event_candidate` with `accepted=0`, high-area-cut independence for annotations, off-grid rejection, fractional counts, Grid disabled, and the absence of a reliable lattice. The real-spectrum comparison uses the revised full-spectrum sigma suggestion and still recovers the previously reported 12-event and 2-event loci.

The complete MSVC feature build passed before the version-only increment to 18.8.75. Generated executables and analysis outputs are excluded from the commit; rebuild to display the new counter. Interactive GUI automation was not used.

## Release history for BI output and plotting

### Version 18.8.76 output format and rounding

Grid mode writes 12 columns in this order: N, Area, Max, xmean, ymean, xl, yl, Sigx, Sigy, nb, Name, Good. With Grid disabled, the existing 10-column 2D schema remains unchanged. The extra columns depend on the selected mode, not on successful lattice fitting. Readers using positional columns must account for the two inserted Grid fields.

xl and yl use nearest-integer rounding of the full-precision mean coordinates, with exact halves rounded away from zero. Examples: -14.1 becomes -14, -14.6 becomes -15, -14.5 becomes -15, and 25.5 becomes 26. These labels are not fitted lattice indices. Rounding precedes text formatting, so a displayed mean of 1.500e+00 can have xl=1 when its full-precision value is 1.49999.

Every floating-point field in .bir uses %.3e scientific notation in standard and Grid mode, including legacy 1D output. Integer columns stay integral. The diagnostic CSV retains its existing precision and status fields. Single-event annotations remain excluded from accepted .bir rows, totals and model spectra.

Regression tests cover exact headers and rows, positive and negative rounding, half ties, rounding before formatting, empty accepted results, and exclusion of single-event candidates. The real Z\_A3Q\_MSU spectrum still recovers both the 12-event and 2-event loci; the single-event option gives 9 annotations without changing accepted results. The output-format MSVC build and recovery suite passed.

### Version 18.8.80 BI result plot correction

The output plot previously read .bir fields by fixed position. With the two Grid columns introduced in 18.8.76, it mistook xl and yl for Sigx and Sigy, and shifted the remaining row fields. This produced oversized ellipses even though the exported mean coordinates and peak widths were correct.

The result plot now locates fields by trimmed, case-insensitive header names. It reads Sigx and Sigy explicitly and ignores the coordinate labels xl and yl. Both old ten-column files and new twelve-column Grid files are supported, including decimal and scientific numeric notation. Required fields must parse successfully and widths must be finite and positive. Incomplete or invalid rows are skipped; early end-of-file no longer reuses stale row data.

The export schema and 18.8.76 rounding rules are unchanged. This fix does not change the detection algorithm, accepted peaks, single-event candidates, diagnostic CSV, or Gaussian spectrum generation. Use the corrected executable when plotting results; the old reader remains incompatible with Grid files.

Tests cover legacy decimal rows, standard and Grid export/read-back, malformed rows, and every accepted peak in the Z\_A3Q\_MSU regression spectrum. Both modes preserve all accepted widths within  $3e$  rounding precision, row counts and acceptance flags. The two sparse loci and nine single-event annotations still pass their existing tests. The full MSVC plot-fix build passed before the version-only increment from 18.8.79 to 18.8.80; rebuild to display the counter.

## Version 18.8.81 CSV result filenames

The default text result filename now ends in .csv for both 1D and 2D searches. The results Browse dialog shows CSV result files (\*.csv), uses the title Save CSV results, and appends .csv when the chosen filename has no extension. The All files filter remains available. Cancel keeps the previous filename. The original comma-separated contents, Grid columns, scientific precision, rounding, and plot reading remain the same. Existing .bir files remain readable. Earlier release sections below use the historical .bir extension. The companion diagnostics filename continues to append .diagnostics.csv to the selected results path, for example Z\_A3Q\_MSU.csv.diagnostics.csv.

## Version 18.8.82 Grid Counts and Bound factor

With Grid mode enabled, the results CSV adds Counts immediately after Area: N, Area, Counts, Max, xmean, ymean, xl, yl, Sigx, Sigy, nb, Name, Good. Standard mode retains its existing columns. Counts is independent of the BI-derived Area and of the local recovery aperture used for peak selection.

After final peak selection, the global widths are calculated over accepted peaks with finite positive Area and widths:  $\sigma_{x\_global} = \text{Sum}(\sigma_{x\_i} * \sqrt{\text{Area}_i}) / \text{Sum}(\sqrt{\text{Area}_i})$   $\sigma_{y\_global} = \text{Sum}(\sigma_{y\_i} * \sqrt{\text{Area}_i}) / \text{Sum}(\sqrt{\text{Area}_i})$ . Both axes use the same contributing peaks and denominator. Rejected peaks, empty sites and single-event annotations do not contribute.

The Grid-only Bound factor control defaults to 2.00, with range 0.01 to 100. The common ellipse semiaxes are Bound factor \*  $\sigma_{x\_global}$  and Bound factor \*  $\sigma_{y\_global}$ . Changing the factor requires Search again. Each accepted peak is associated with its nearest fitted lattice site, using the fitted grid basis; Counts integrates original bins around that site. Measured xmean/ymean and their rounded xl/yl labels are retained as before. If no reliable lattice is available, counting uses measured peak centers and the information dialog explicitly reports this fallback.

Counts sums original bin contents whose bin centers lie inside or on the axis-aligned ellipse. It uses the original loaded matrix, including bins outside the search limits when an ellipse crosses those limits. Ellipses are clipped at matrix edges. No smoothing, Gaussian extrapolation or background subtraction is applied to Counts. For unweighted spectra the sum is an event count; weighted spectra retain their fractional sums. Overlapping ellipses may count the same bin, so adding all Counts is not a unique-event total. Counts uses the established  $3e$  numeric output format.

After a successful Grid search, an information dialog reports both weighted global sigmas, the number of contributing peaks, Bound factor, both ellipse semiaxes and the center convention. The plot header uses these global sigmas. An accepted peak's overlay details report its Counts, actual counting center and ellipse semiaxes. If no valid average exists, the dialog explains this and Counts cells are empty. Single-event annotations remain diagnostic only.

## Validation of version 18.8.82

Focused numerical tests distinguish  $\sqrt{\text{Area}}$  weights from Area weights and cover ellipse boundaries, excluded rectangle corners, fitted-site projection, original bins beyond search limits, fractional input, clipped

spectrum edges, Grid-off behavior, empty averages and cancellation. The independent full-matrix check matches Counts for all 186 accepted sites in Z\_A3Q\_MSU.spa. For the regression settings sigma X=0.15, Y=0.13, area=0.1, minimum rows=1, Grid on, tolerance=25% and Bound factor=2, weighted globals are X=0.148777 and Y=0.0900391. These are regression observations for that input and settings. The previous two-event and twelve-event recovery tests, nine single-event annotations and both CSV plot-reader round trips also pass. The edited dialog sources passed syntax checks. No full application build was performed for this feature; compilation was not repeated for this release at user request.