NEXCHROM · TECHNICAL BACKGROUND

Calculation and uncertainty

How NexChrom calibrates, selects and accepts a regression model, calculates the composition and evaluates the measurement uncertainty, from certificate import to the analysis certificate, and the gas properties according to ISO 6976:2016.

Qualimetric · Nijmegen, The Netherlands · info@qualimetric.nl · nexchrom.nl · 3 November 2026

CONTENTS

1. The calculation chain at a glance
2. Calibration input
3. Regression methods
4. Order selection and acceptance
5. Calibration models and traceability
6. Sample results
7. Measurement uncertainty
8. Analysis certificates
9. Physical properties according to ISO 6976:2016
10. Standards and references
11. Further information

NexChrom takes a gas chromatography laboratory from the peak tables in its CDS files to a reviewed analysis certificate: it assembles calibration points from certified reference gases and replicate injections, fits an analysis function per component, calculates the mole fractions of a sample, propagates their measurement uncertainty and reports the normalised composition with expanded uncertainties. From that composition it also calculates the physical properties of the gas according to ISO 6976:2016, with their uncertainties. This document describes how each of those steps is calculated: the input data, the regression methods, the selection and acceptance of a polynomial order, the calibrated range, the calculation of sample results, the treatment of measurement uncertainty from the certificate of a reference gas to the certificate of an analysis, and the physical properties. It is written for analysts, quality managers and assessors in gas laboratories working to ISO/IEC 17025 who need to understand, and be able to verify, what the software computes. The formulas and thresholds in this document correspond to the behaviour of the software.

The calculation chain at a glance

Every reported mole fraction passes through the same chain of steps. Each step uses only stored, traceable quantities from the step before it.

Step	What NexChrom does	Main output
Calibration input	Joins the certified value of each component in the selected reference gases with the mean peak area of the replicate injections of those gases	Calibration points (x, y) with standard uncertainties u(x) and u(y)
Regression	Fits a polynomial analysis function per component, by default with errors-in-variables regression in conformity with ISO 6143	Coefficients, covariance matrix, fit statistics (goodness of fit Γ, Birge ratio)
Order selection and acceptance	Fits orders 1 to 3 together, proposes the lowest order that passes the fit test of ISO 6143 and flags fits that fail it	One selected fit per component, saved in a calibration model

Step	What NexChrom does	Main output
Sample results	Evaluates the analysis function at the mean peak area of the sample; derives indirect components via relative response factors; adds fixed components	Mole fractions, with a check against the calibrated range
Uncertainty	Propagates the uncertainty of the calibration and the repeatability of the sample according to the GUM	Expanded uncertainty U = k·u per component
Certificate	Normalises the composition, propagates the uncertainty through the normalisation and freezes the result as a PDF	Analysis certificate with normalised composition and expanded uncertainties
Physical properties	Calculates the properties of ISO 6976:2016 from the normalised composition, including the correlations that the normalisation creates	Calorific values, density, relative density, Wobbe indices, molar mass and compression factor, each with U = k·u; optionally a page on the certificate

Two conventions run through the whole chain:

- **Intermediate calculations use standard uncertainties.** Certified uncertainties are converted to standard uncertainties before they enter the regression; only the final results are expanded.
- **One coverage factor.** The coverage factor k is a single laboratory-wide setting (Settings → System → coverage factor k, default 2). It governs the conversion of imported certificate uncertainties, the expansion of results and certificates, and the intercept test described below. The k with which the results of a run are expanded is stored with the run, so that its certificate keeps that k when the setting changes later.

Each component in an analysis method has one of three roles. A **direct** component is calculated from its own calibration function. An **indirect** component is calculated from a direct reference component and a relative response factor (RRF). A **fixed** component carries a value entered in the method, for example a component that is not measured but must be part of the normalised composition.

Calibration input

Calibration points

A calibration point combines one reference gas with its measured response. For every reference gas selected for a calibration, NexChrom joins:

- **y, the certified value**: the mole fraction of the component on the current certificate of the reference gas (in mol/mol unless the certificate states another unit);
- **u(y), its standard uncertainty**: the expanded uncertainty on the certificate divided by the coverage factor k, u(y) = U(y)/k;
- **x, the mean peak area**: the mean of the peak areas of the component over the replicate injections (CDS files) of that reference gas;
- **u(x), its standard uncertainty**: the experimental standard deviation of the mean, u(x) = s/√n, where s is the standard deviation of the n replicate peak areas; for a single injection see below.

$$\bar{x} = (1/n) \cdot \Sigma_j A_j$$
$$u(\bar{x}) = s / \sqrt{n}$$
$$u(y) = U(y) / k$$

The coverage factor applied to each certificate is recorded with the calibration, so that every saved fit shows on which scale its input uncertainties were taken. A component/reference-gas pair without a certified value or

without a peak area is left out.

Two cases have no uncertainty of their own, and the laboratory decides how they are treated (Settings → System → Calibration/Regression).

A level with a single injection has no spread of its own. By default it gets the pooled relative spread of the injections of the same component at its other levels, times its peak area:

$$RSD = \sqrt{\Sigma_i (n_i - 1) \cdot (s_i / \bar{x}_i)^2 / \Sigma_i (n_i - 1)}$$
$$u(\bar{x}) = RSD \cdot \bar{x}$$

ISO 6143:2001 allows this: under statistical control the standard uncertainty may be calculated from the known method standard deviation. The laboratory can set a fixed value instead, relative or absolute. Where no value is known, because no other level has more injections, the level counts as practically exact on the x-axis and the page warns.

A certified value without an uncertainty should not occur on a calibration certificate. For the sporadic case the laboratory sets a fixed standard uncertainty, relative or absolute, by default 1 % of the certified value, and the page names the component and the reference gas.

Both rules are recorded with the fit, so that a saved fit keeps the rule it was fitted with.

Calibration input is assembled for every component that has a certified value in the selected reference gases, whatever its role in the analysis method. The regression itself is method-independent: all components are fitted alike, and the method decides at the moment of calculation which fits are used. The fits of components that a method treats as indirect remain available in the calibration model, for example to compare the curve route with the RRF route. A failing fit of one component never blocks saving the fits of the others.

Analysis sessions

The replicate injections of a reference gas are grouped into analysis sessions by the calendar date of the analysis. For each reference gas, the most recent session is selected by default, so that a recalibration uses fresh data and drift is not averaged away with older injections. A laboratory that calibrates components at different frequencies, for example monthly gases next to yearly gases for minor components, can deliberately select an older session for one reference gas. Selecting several sessions of one gas pools them into one mean; NexChrom shows this and warns about it.

Detectors and duplicate peaks

When a component is detected on more than one detector or signal, NexChrom keeps one response per analysis, component and detector, and uses the response of the detector chosen in the instrument's detector template. Peak areas from different detectors are never summed. The replicates behind one calibration point are expected to come from the same detector; a point whose replicates mix detectors is marked in the calibration input and listed in a warning, both at import and at calibration.

Reviewing and excluding replicates

For every calibration point the analyst can inspect the individual replicates: the peak area of each injection with its deviation from the mean and its z-score, a control chart with the mean and the ±1s, ±2s and ±3s limits of the included replicates, and a Grubbs outlier test (two-sided, $\alpha = 0.05$, applied iteratively). The Grubbs test is advisory.

A single replicate can be excluded from the mean of a calibration point without changing the analysis itself, which may serve other calibrations. An exclusion is stored with the calibration model, with its reason (manual or Grubbs outlier), recorded in the audit log, reversible, and carried over when the model is saved as a new version.

The regression cannot be run while an exclusion has not yet been saved, so that a fit always corresponds to a stored set of points.

Axis convention: the analysis function

NexChrom fits the **analysis function** directly:

```
y = f(x)       x = mean peak area,  y = mole fraction
```

The concentration is the dependent variable and the peak area the independent variable. A sample result is obtained by evaluating the stored polynomial at the measured peak area, without an inversion step. This is the orientation ISO 6143 recommends: the standard prefers the direct determination of the analysis function, the content as a function of the response, over fitting the calibration function (response as a function of content) and inverting it, and its worked examples fit the analysis function as well. NexChrom uses this orientation also because it avoids a numerical inversion at every sample result and lets the uncertainty of a result follow directly from the coefficient covariance. With errors-in-variables regression, both variables are weighted by their own uncertainties in either orientation.

Regression methods

The polynomial model

All methods fit a polynomial in ascending powers of the peak area:

$$y = b_0 + b_1 \cdot x + b_2 \cdot x^2 + \ldots + b_m \cdot x^m$$

m = polynomial order (1 to 3), $p = m + 1$ coefficients

$\nu = n - p$ degrees of freedom (n calibration points)

The design matrix is the Vandermonde matrix V with rows $(1, x_i, x_i^2, \ldots, x_i^m)$. A fit requires at least one degree of freedom, $\nu \geq 1$; calibrations with one or two points are treated separately (see the chapter on order selection). Four regression methods are available. Errors-in-variables regression is pre-selected; the others remain available as references.

EIV: errors-in-variables regression in conformity with ISO 6143

The errors-in-variables (EIV) regression accounts for the uncertainty in both coordinates of every calibration point: the certified concentration and the mean peak area. It is a Deming-type estimator that adjusts both coordinates of every point, each in proportion to its own uncertainty, and finds the polynomial that requires the smallest total weighted adjustment. NexChrom recommends it for traceable calibration whenever replicate injections provide a real estimate of u(x).

Like every other fit, an EIV fit is tested on every calibration point: the goodness-of-fit criterion Γ of ISO 6143, with the fixed limit of 2 of the standard. The test is described under *Order selection and acceptance*.

The estimator minimises, over the adjusted values $\hat{x}_i$, $\hat{y}_i$ and the coefficients b, the Lagrangian

$$L = \tfrac{1}{2} \cdot \Sigma_i \left[(x_i - \hat{x}_i)^2 / u^2(x_i) + (y_i - \hat{y}_i)^2 / u^2(y_i) \right] - \Sigma_i \lambda_i \cdot \left(\hat{y}_i - F(\hat{x}_i, b) \right)$$

where F is the polynomial and the Lagrange multipliers λ_i enforce that every adjusted point lies exactly on the curve. The minimum is found with a Gauss-Newton iteration. Each point carries the scalar weight

$$W_i = 1 / \left(u^2(x_i) \cdot (\partial F/\partial x)^2 + u^2(y_i) \right)$$

which combines both uncertainties, with the x-uncertainty transferred to the y-axis through the local slope of the curve. The coefficient update solves $(J^T \cdot W \cdot J) \cdot \Delta b = J^T \cdot W \cdot c$, with $J = \partial F / \partial b$ the Jacobian of the polynomial with respect to its coefficients and c the linearised constraint residual.

Properties of the implementation:

- **Starting values** from a weighted forward fit, with weights $1/(1/u^2(x) + 1/u^2(y))$, and an unweighted inverse fit to seed the adjusted x values.
- **Convergence** when the relative norm of the coefficient update satisfies $\|\Delta b\| < 10^{-10} \cdot (1 + \|b\|)$, with at most 500 iterations.
- **Divergence guard**: when a step increases the value of the objective, the step is rolled back.
- **Convergence is reported, not assumed.** A fit that stops on the divergence guard or on the iteration limit is marked as not converged, with the reason, in the results overview and in the component detail, and is listed in a warning. The automatic order proposal skips such a fit when a converged order is available.
- **Covariance of the coefficients**: the inverse of the normal matrix at the solution, $\Sigma_b = (J^T \cdot W \cdot J)^{-1}$, computed from the stated uncertainties and not rescaled by the observed scatter. If the matrix is singular, a pseudo-inverse is used and a warning is given.
- **Residual sum over both axes**, the basis of the Birge ratio:

$$TSSD = \Sigma_i \, (x_i - \hat{x}_i)^2 / u^2(x_i) + \Sigma_i \, (y_i - \hat{y}_i)^2 / u^2(y_i)$$

- **Per point, the two signed weighted deviations** $(x_i - \hat{x}_i)/u(x_i)$ and $(y_i - \hat{y}_i)/u(y_i)$ from the adjusted point, the largest absolute value of which is Γ; and the probability of the TSSD under the χ^2 distribution with ν degrees of freedom, reported as information.

The EIV routine has been validated separately with internal-consistency checks, Monte Carlo simulation and a comparison with the worked examples of ISO 6143:2001 Annex B; the results are described in the companion document *Validation of the EIV regression*.

OLS: ordinary least squares

OLS minimises the unweighted sum of squared y-residuals:

$$\text{minimise} \;\; \Sigma_i \, (y_i - \hat{y}_i)^2$$
$$\Sigma_b = s^2 \cdot (V^T V)^{-1}, \quad s^2 = \Sigma_i \, (y_i - \hat{y}_i)^2 / \nu$$

All points carry equal weight, and the coefficient covariance is scaled by the a posteriori residual variance. OLS is kept deliberately as an unweighted reference. On data that span several decades of concentration, its Birge ratio, which is always weighted with u(y) (see below), shows clearly that an unweighted fit neglects the low-concentration points.

WOLS: weighted least squares

WOLS weights each point by the inverse variance of its certified value:

$$\text{minimise} \;\; \Sigma_i \, w_i \cdot (y_i - \hat{y}_i)^2, \quad w_i = 1 / u^2(y_i)$$
$$\Sigma_b = (V^T \cdot W \cdot V)^{-1}$$

The covariance is the a priori covariance, consistent with weights that are true inverse variances. The uncertainty of the peak areas does not enter.

GLM: robust reweighting

The method listed as GLM is a robust iteratively reweighted least-squares fit, an M-estimator with a Huber-type weight function. Starting from unit weights, it solves $(V^T \cdot W \cdot V) \cdot b = V^T \cdot W \cdot y$, computes the residuals r_i and reweights each point with a tuning constant equal to twice the median absolute residual:

$$
\begin{aligned}
c &= 2 \cdot \text{median } |r_i| \\
w_i &= 1 \qquad\qquad \text{if } |r_i| \le c \\
w_i &= c / |r_i| \qquad \text{if } |r_i| > c \\
\Sigma_b &= s^2 \cdot (V^T \cdot W \cdot V)^{-1}, \quad s^2 = \Sigma_i \, w_i \cdot r_i^2 / v
\end{aligned}
$$

The iteration stops when the largest change in a coefficient is below 10^{-8}, with at most 25 iterations. Points with a residual larger than twice the median absolute residual are down-weighted. Certificate uncertainties do not enter the GLM weights.

Order selection and acceptance

Candidate orders

For a component with n calibration points, NexChrom fits every polynomial order from 1 up to min(3, n − 2) in one run, so that each candidate has at least one degree of freedom. Three calibration points therefore allow a straight line, four points a quadratic and five or more points orders up to 3. The analyst sees all candidates side by side and selects one order per component; the proposed order is marked as such. From three points on, NexChrom also computes the straight line through the origin, as the variant *1 (origin)*: selectable per component, tested like every other fit, and never proposed.

The fit statistics: Γ and the Birge ratio

Two statistics describe how well a fit explains its calibration points within their stated uncertainties. The first is the test of every fit; the second is information.

Γ, the goodness-of-fit criterion of ISO 6143, examines every calibration point separately:

$$
\Gamma = \max_i \left(\frac{|x_i - \hat{x}_i|}{u(x_i)} \text{ and } \frac{|y_i - \hat{y}_i|}{u(y_i)} \right)
$$

with $(\hat{x}_i, \hat{y}_i)$ the adjusted point of the errors-in-variables fit. For OLS, WOLS and GLM the peak area is taken as exact ($\hat{x}_i = x_i$), so their Γ is the largest $|y_i - \hat{y}_i| / u(y_i)$, with the uncertainties of the certified values only. In plain words: how far the worst point lies from the curve, in standard uncertainties. A calibration function is admissible when $\Gamma \le 2$, which means that the curve passes through every calibration rectangle $x \pm 2u(x)$, $y \pm 2u(y)$; of several admissible functions the standard prefers the one with the fewest coefficients, unless the fits differ considerably. The weighted deviations per point and Γ were validated against the worked examples of ISO 6143:2001 Annex B (see *Validation of the EIV regression*). The software calls this statistic **GOF**.

The **Birge ratio** (BR), the root of the reduced χ^2, describes the fit as a whole:

$$
\begin{aligned}
BR &= \sqrt{\left(TSSD / v \right)} \\
TSSD &= \Sigma_i \left(\frac{y_i - \hat{y}_i}{u(y_i)} \right)^2 && \text{OLS, WOLS, GLM} \\
TSSD &= \Sigma_i \left(\frac{x_i - \hat{x}_i}{u(x_i)} \right)^2 + \Sigma_i \left(\frac{y_i - \hat{y}_i}{u(y_i)} \right)^2 && \text{EIV}
\end{aligned}
$$

For OLS, WOLS and GLM the TSSD is weighted with the standard uncertainties of the certified values whenever these are available, regardless of the internal weights of the method. For EIV the sum runs over both axes, with the adjusted values $\hat{x}_i$ and $\hat{y}_i$ of the estimator. The Birge ratio and Γ of OLS, WOLS and GLM therefore count the

certificate uncertainties only, not the spread of the peak areas, and are stricter (higher) than those of EIV on the same data; the software says so where these methods are offered.

When the stated standard uncertainties describe the real scatter and the model has the right shape, TSSD follows a χ^2 distribution with ν degrees of freedom and BR^2 has an expectation of 1. A Birge ratio clearly above 1 means that the residuals are larger than the uncertainties predict: either the uncertainties are understated, or the polynomial does not describe the response. In plain words: how many times larger the scatter of the points around the curve is than their stated uncertainties lead you to expect; 1 is as expected, 2 is twice as large.

ISO 6143:2001 uses this sum as an overall measure and names the goodness-of-fit test on Γ as the definitive acceptance test. NexChrom follows the standard: the Birge ratio decides nothing. The component detail shows it with a **BR test**, for information: TSSD against the 95 % point of χ^2 at ν degrees of freedom, so that the limit on the Birge ratio depends on ν (1.96, 1.73, 1.54, 1.45 and 1.35 at ν = 1, 2, 4, 6 and 10). A failed BR test does indicate that something may be wrong.

The two statistics do not map one to one: Γ takes the larger of the two weighted deviations of the worst point, the Birge ratio averages the squares of all of them. A fit can pass Γ with a Birge ratio above 2, and fail it with a Birge ratio below 2. The coefficient of determination R^2 (on unweighted y-residuals) is reported alongside, but it is not used for acceptance and is not comparable with either statistic.

Acceptance: the fit test

Every fit, whatever its regression method, is tested on **$\Gamma \leq 2$**, the goodness-of-fit test of ISO 6143:2001. The limit is fixed by the standard and is not a setting.

The results show the outcome as a **fit test**, *Pass* or *Fail*, next to the GOF value; the limit stands in the column header. A fit that fails is shown in red and listed in a warning banner, and the table of calibration points marks the level and axis whose weighted deviation exceeds 2. A fit that cannot be tested, such as an exact line through two points (ν = 0) or the fall-back through the origin on one or two points, shows no GOF and no fit test and is marked as untested instead.

The acceptance check is **advisory**: it flags and warns, but it does not block saving a model or calculating results with it. Acceptance of a calibration is a review decision, made by people and evidenced by the flag, and every calibration model is released through a second-person approval.

The Auto order rule

The proposed order, shown as *Optimal* in the results, is the lowest order that passes the fit test, raised to the next order when that order passes as well and the two fits differ considerably:

```
m = the lowest order with Γ(m) ≤ 2
while Γ(m + 1) ≤ 2 and TSSD(m) − TSSD(m + 1) > 3.84:
    m = m + 1
propose m
if no order passes:  propose nothing
```

This is the rule of the standard: of the admissible functions, take the one with the fewest coefficients, unless the fits differ considerably. The standard gives no number for *differ considerably*. NexChrom reads it as: the extra coefficient lowers TSSD by more than 3.84, the 95 % point of χ^2 with one degree of freedom. The lowest adequate order guards against overfitting: a higher-order polynomial always lowers the residuals, but it also adds a coefficient with its own uncertainty and can behave poorly between and beyond the calibration points. The same limit governs both the proposal and the fit test, so the two cannot diverge. The analyst can override the proposal per component in *Used*; the proposal stays visible as the optimal order. An order whose iteration did not converge is skipped by the proposal when a converged order exists.

Optimal also carries the verdict of the intercept test of the proposed fit (see *Intercept significance*): the plain order, for example *2*, or *2 ($b_0 \approx 0$)* when the data are consistent with a zero intercept. This is information: *Used* starts at the fit with intercept either way.

When no order passes, nothing is proposed and *Optimal* reads "—". *Used* then starts at the model set in Settings → System → Calibration/Regression, by default the 1st order with intercept, or at the lowest converged order where that model is not available. The model keeps its own fit test, normally *Fail*, and the results page says so. It is a starting value: the analyst can change it.

Regression detail

The results table of a regression shows, per component and order, the GOF with its limit in the header, the fit test and the convergence. The table of saved fits and the fit overview on the Review page show the GOF and the fit test as well. For every fit, both a new regression and a saved fit, NexChrom shows the full statistics:

- the coefficients b_i with their standard uncertainty $u(b_i)$, relative uncertainty and the ratio $|b_i|/u(b_i)$;
- the covariance matrix and the correlation matrix of the coefficients;
- the fit statistics: number of points, ν, TSSD, the Birge ratio with its BR test, Γ, R^2 and the calibrated range, with the GOF and its limit at the top;
- the calibration points with their residuals, their contribution to χ^2, their weighted deviations, and the GOF after a re-fit without that level (*Without this level*);
- the calibration curve and the residual plots.

The contribution of each point to χ^2 is

$$\chi^2_i = \left((y_i - \hat{y}_i) / u(y_i) \right)^2 \qquad \text{OLS, WOLS, GLM}$$

$$\chi^2_i = \left((y_i - \hat{y}_i) / u(y_i) \right)^2 + \left((x_i - \hat{x}_i) / u(x_i) \right)^2 \qquad \text{EIV}$$

For EIV these are the residuals of the adjusted points, the quantities the estimator actually minimises, not the vertical distance to the curve. The sum of the contributions is the TSSD. The weighted deviations of an EIV fit are the two signed terms of its contribution; OLS, WOLS and GLM have one, $(y_i - \hat{y}_i)/u(y_i)$. A deviation beyond 2 is marked in red, because that is the level and axis on which the fit fails Γ. The χ^2 contribution alone does not point reliably to the level that disturbs a fit: a level at the edge of the range pulls the curve towards itself and keeps a small contribution.

Without this level. For every calibration level the table of points therefore also shows Γ of a re-fit with the same method and order without that level. It is computed when the detail is shown and never stored, and it is left empty when the re-fit would have fewer than one degree of freedom. When the fit fails the test and would pass it without one level, the row of the level whose removal helps most is marked in orange, and a caption names it with both values and the other levels without which the fit would also pass. Orange is a pointer to investigate that level, for example the certificate, the cylinder, the integration or the injections; it is not advice to drop the level. When another order passes with all levels, the caption says so: the cause is then more likely the order than a single level. When no single level makes the fit pass, the caption says that as well: several levels deviate, or the order may be too low.

The plots. In the calibration curve and in both residual plots the calibration points are drawn in dark blue; the orange level is drawn in orange and slightly larger, and a point whose weighted deviation exceeds 2 is drawn in red. The hover text of every point gives its χ^2 contribution and the GOF without that level.

Residual diagnostics

A fit that passes its fit test can still show a systematic pattern, such as a bow in the residuals of a straight line. NexChrom therefore reports three advisory diagnostics, which are not required by ISO 6143 or ISO 6974:

- a **runs test** (Wald–Wolfowitz) on the signs of the residuals ordered by peak area, with an exact one-sided p-value and the number of runs;
- the **standardised residuals** $z_i = \pm\sqrt{\chi^2_i}$ (with the sign of the residual), with their mean and the largest $|z_i|$;
- a **Shapiro–Wilk test** for normality of the standardised residuals.

With the six or seven points of a typical calibration these tests have limited power, so the number of runs and the largest $|z_i|$ are shown next to the verdict, to be read together with the residual plots.

Intercept significance

For every fit with an intercept, NexChrom tests whether the intercept is consistent with zero:

```
|b₀| > k · u(b₀)    →  significant intercept
|b₀| ≤ k · u(b₀)    →  consistent with a zero intercept
```

with k the laboratory coverage factor. From three calibration points upwards this indicator is advisory: the intercept is always estimated and never forced to zero, because with enough points a free intercept is estimable and forcing it would replace information by an untested assumption.

Calibrations with one or two points

Some components are present in only one or two reference gases. By default such components are skipped, with a warning; the rest of the calibration proceeds. Behind an explicit opt-in, NexChrom can fit them as follows.

One point. A straight line through the origin is the only possible line, resting on the assumption of a response proportional to concentration:

```
b₁ = y₁ / x₁
u²(b₁) = b₁² · [ (u(y₁)/y₁)² + (u(x₁)/x₁)² ]
```

The calibrated range is then 0 to y_1.

Two points. NexChrom fits both the exact straight line with intercept through both points, with the coefficient covariance obtained by analytical propagation of u(y) and u(x) of both points, and the weighted line through the origin:

```
b₁ = Σᵢ wᵢ·xᵢ·yᵢ / Σᵢ wᵢ·xᵢ²,     wᵢ = 1 / u²(yᵢ)
u²(b₁) = 1 / Σᵢ wᵢ·xᵢ²            ν = 1
```

Both variants are always computed and selectable per component. The selection starts at the line with intercept; the intercept test above shows its verdict in *Optimal* (*1* or *1 ($b_0 \approx 0$)*) without switching the fit.

With one or two points the adequacy of the model (linearity, zero intercept) cannot be tested, and the Birge ratio of an exact line is undefined ($\nu = 0$). Such fits therefore show no GOF and no fit test; they are marked as untested in the calibration results, in a warning, and on the calculation page next to every component that uses them.

Calibration models and traceability

What is stored per component

When a calibration is saved, NexChrom stores one fit per component in a calibration model, at the selected order:

- the coefficients $b_0 \ldots b_m$ in ascending powers of the peak area, and their covariance matrix;
- TSSD, ν, Γ, the Birge ratio and R^2;

- the fit type (polynomial with intercept, or line through the origin) and the regression method;
- the time of the fit and the date of the newest analysis behind it;
- the input set: per reference gas the certificate, the coverage factor applied, the selected session(s) with the number of files, the analyses used, the replicates excluded from the mean, and the rule applied to a level with a single injection or a certified value without an uncertainty;
- the calibrated range, in concentration (the certified values of the lowest and highest reference gas) and in peak area (their mean peak areas).

Because the input set is recorded, a saved fit can be rebuilt from the original data at any time. When a saved fit is opened, NexChrom reconstructs its points from the recorded certificates and analyses and shows the same detail as a new regression; a mismatch with the stored range is reported. For EIV, the fit is recomputed and must reproduce the stored coefficients, as a reproducibility check; the recomputation also supplies the weighted deviations per point and the χ^2 probability.

Models are preserved, not overwritten

Saving into a calibration model that already holds saved fits, or that has been used for results, does not silently overwrite it. The default is to save a new model; the previous model is archived with all its fits, so that every earlier result keeps the calibration it was calculated with. Overwriting is possible only as an explicit choice. Every saved change, the regression input of every save and every activation or archiving of a model are recorded in the audit log.

The same principle applies one level up. Component roles and RRF rules are read from the analysis method at the time of calculation; changing them in a method that has already been used requires an explicit confirmation, and the recommended route is to duplicate the method as a new version.

Age of the calibration

When results are calculated, NexChrom shows the age of the oldest calibration data behind the selected model and warns per component when it exceeds a configurable maximum age (default 365 days). The warning is advisory, so that a deliberate reuse of an older session for a slowly varying component remains possible but is always visible.

Sample results

A sample is measured in one or more replicate injections. Per component, NexChrom uses the mean peak area $\bar{A}$ over the included replicates of the sample and its standard uncertainty $u(\bar{A}) = s/\sqrt{n}$. From here on, x denotes the calculated mole fraction of a component in the sample, and f the analysis function of the calibration.

Direct components

A direct component is calculated by evaluating the stored analysis function at the mean peak area of the sample:

$$x = f(\bar{A}) = b_0 + b_1 \cdot \bar{A} + b_2 \cdot \bar{A}^2 + \ldots + b_m \cdot \bar{A}^m$$

There is no inversion step: the calibration fitted concentration as a function of peak area, and the polynomial is evaluated forward.

Indirect components via relative response factors

An indirect component is calculated from a direct reference component through a relative response factor:

$$x_{ind} = \bar{A}_{ind} \cdot h / RRF, \qquad h = x_{ref} / \bar{A}_{ref}$$

where Ā_ind and Ā_ref are the mean peak areas of the indirect and the reference component, x_ref is the calculated mole fraction of the reference component and h its response factor in the sample. The reference must itself be a direct component that can be calculated; when the reference, a peak area or the RRF is missing, NexChrom reports an error and gives no result for that component.

The RRF is either entered in the method or calculated from the carbon numbers of the two components:

```
RRF = n_C(ind) / n_C(ref)
```

The calculated RRF rests on the effective-carbon-number principle for a flame ionisation detector: a response proportional to the number of carbon atoms. It is intended for hydrocarbons measured on an FID.

Fixed components

A fixed component has its value, and optionally its expanded uncertainty, entered in the analysis method:

```
x = fixed value,     U = expanded uncertainty as entered
```

A fixed component does not need a peak in the sample: it is always reported, takes part in the normalisation and contributes to the normalised uncertainties. The method cannot be saved with a fixed component that has no value, and a validation check before the calculation reports any fixed component without one.

Calibrated range and extrapolation

Every fit carries its calibrated range in concentration: the certified values of the lowest and highest reference gas behind it (from 0 for a line through the origin). At calculation, the mole fraction of each direct component is classified against that range:

Status	Meaning
within	the result lies inside the calibrated range
below	the result lies below the lowest reference gas
above	the result lies above the highest reference gas
ref-below, ref-above	indirect component whose reference component lies outside its range

An indirect component inherits the status of its reference component, because its value rests on the reference's response factor h.

The policy is **flag and annotate, never block**. The result is still calculated, but the calibrated range and the status are shown next to it (in red when outside), an extrapolation warning lists every affected component with its value and range, the status is stored with the result, and the review page shows the range next to each result and repeats the warning before approval. A polynomial is valid within its calibrated range only, and outside that range even a well-fitting polynomial of order 2 or 3 can be seriously wrong; the flag makes every such deviation visible on the result itself and leaves the decision with the reviewer.

The range is stated in concentration because that is how the analyst reads and reports results and how an accreditation scope is stated. For the monotonic analysis functions that calibration produces, a range check in concentration and in peak area coincide. The range in peak area is stored as well, as a diagnostic.

Normalised composition

For a certificate, NexChrom reports the normalised composition, in which the mole fractions of all components, including fixed components, add up to one:

$$S = \Sigma_j x_j$$
$$x'_i = x_i / S$$

The unnormalised mole fractions remain stored with the results and are listed on the certificate as well. The uncertainty of the normalised composition is described in the next chapter.

Measurement uncertainty

Principles

NexChrom evaluates measurement uncertainty according to the GUM (JCGM 100):

- Every input enters as a standard uncertainty. Expanded uncertainties from certificates are divided by their coverage factor before use; the relative RRF uncertainties of ISO 6974-2 are standard uncertainties and are used as such.
- Contributions are combined by first-order propagation, with the correlations that follow from the calculation itself (the covariance of the calibration coefficients, the shared reference of an indirect component, the common divisor of the normalisation).
- Results are reported as expanded uncertainties U = k·u, with the laboratory coverage factor k (default 2, approximately 95 % coverage for a normal distribution).
- A missing or unusable input yields an empty uncertainty, never an understated one. Where a component uncertainty is missing in a normalised composition, it counts as zero and the result is flagged as incomplete.

The independence assumptions are explicit: the measurement of the sample is treated as independent of the calibration data, and the components are treated as independent of one another apart from the correlations listed above.

Certificate import

When the certificate of a reference gas is imported, NexChrom classifies its uncertainty statement:

- an explicit coverage factor k = 2, "expanded uncertainty" or "standard uncertainty multiplied by …": the stated values are expanded uncertainties and are stored as such;
- an explicit coverage factor k = 1, or "standard uncertainty" without a multiplication: the stated values are standard uncertainties and are converted to expanded uncertainties with the laboratory coverage factor;
- a statement that cannot be classified: the user is asked to verify the uncertainty manually.

The stored certificate uncertainty is therefore always an expanded uncertainty on the laboratory coverage factor. Imported certificates are checked by a person before approval.

Uncertainties entering the calibration

The regression works with standard uncertainties on both axes:

Quantity	Source	Scale
u(y), certified value	Stored expanded uncertainty of the reference gas divided by k; without one, the fixed value of the laboratory	Standard
u(x), mean peak area	s/√n of the included replicate injections; with a single injection, the rule under *Calibration points*	Standard

The coverage factor applied is recorded with every calibration point. Weights, TSSD, the Birge ratio, Γ, the EIV weights, the coefficient covariance, the error bars in the plots and the intercept test all operate on these standard

uncertainties. A single injection gives no estimate of the repeatability of its own, and the rule described under *Calibration points* then applies. Replicate injections of every reference gas are therefore recommended, in particular for EIV, whose benefit rests on a real estimate of u(x).

The calibration contributes to the uncertainty of a result through the covariance matrix of the coefficients, Σ_b. For EIV and WOLS this matrix follows from the stated uncertainties; the Birge ratio shows whether those uncertainties are consistent with the observed scatter, and Γ whether every calibration point lies within its uncertainty of the curve.

Uncertainty of direct components

For a direct component x = f(Ā), the combined standard uncertainty contains the uncertainty of the calibration function at the measured peak area and the repeatability of the sample:

$$u^2(x) = g^T \cdot \Sigma_b \cdot g + (f'(\bar{A}))^2 \cdot u^2(\bar{A})$$
$$g = (1, \bar{A}, \bar{A}^2, \ldots, \bar{A}^m)^T$$
$$f'(\bar{A}) = b_1 + 2 \cdot b_2 \cdot \bar{A} + \ldots + m \cdot b_m \cdot \bar{A}^{m-1}$$
$$u(\bar{A}) = s / \sqrt{n}$$
$$U(x) = k \cdot u(x)$$

The first term is the uncertainty of the calibration function, including all correlations between its coefficients; the second is the repeatability of the replicate injections of the sample, transferred through the slope of the analysis function.

A sample with a single injection has no spread of its own either. By default u(Ā) is then the pooled relative spread of the injections behind the calibration model, times the peak area, or a fixed value set by the laboratory (Settings → System → Uncertainty). The calculation page names the components concerned, and says so when no value is known.

Uncertainty of indirect components

For an indirect component x_ind = Ā_ind · h / RRF, the relative standard uncertainty combines three independent inputs: the mean peak area of the indirect component, the RRF, and the response factor h of the reference component:

$$(u(x_{ind}) / x_{ind})^2 = (u(\bar{A}_{ind}) / \bar{A}_{ind})^2 + (u(RRF) / RRF)^2 + (u(h) / h)^2$$

The response factor h = f(Ā_ref) / Ā_ref contains the peak area of the reference component in both numerator and denominator. Its uncertainty takes that correlation into account exactly:

$$u^2(h) = (\partial h / \partial \bar{A}_{ref})^2 \cdot u^2(\bar{A}_{ref}) + g^T \cdot \Sigma_b \cdot g / \bar{A}_{ref}^2$$
$$\partial h / \partial \bar{A}_{ref} = (\bar{A}_{ref} \cdot f'(\bar{A}_{ref}) - f(\bar{A}_{ref})) / \bar{A}_{ref}^2$$

with g and Σ_b taken from the calibration function of the reference component at Ā_ref.

Uncertainty of the relative response factor

ISO 6974-2 (Annex B) gives default relative standard uncertainties for relative response factors, depending on the detector, and allows a thoroughly validated alternative. NexChrom resolves the RRF uncertainty of each indirect component in four steps, the first applicable step being used:

Step	Source of the RRF uncertainty	Value
1	A validated uncertainty entered on the RRF rule in the method, with its source documented	The entered standard uncertainty

Step	Source of the RRF uncertainty	Value
2	The detector type of the component's included responses	FID 2 %, TCD 10 % (relative, standard)
3	For exports without detector information: the type of the instrument's detector, when exactly one detector is registered for the instrument	FID 2 %, TCD 10 % (relative, standard)
4	No detector type known	FID default, 2 % (relative, standard), with a warning

The default values for FID and TCD can be configured in Settings → System → RRF defaults. Steps 3 and 4 are reported in the warnings of the calculation, distinguishing a response without detector information from a detector that is linked but has no type. The RRF uncertainties are deliberately not expanded: like all intermediate quantities they are standard uncertainties, and only the final result is expanded.

Uncertainty of fixed components

The uncertainty of a fixed component is the expanded uncertainty entered with its value in the method, on the laboratory coverage factor. When none is entered, the component is reported without an uncertainty, with a warning, and contributes zero to the normalised uncertainties.

Uncertainty of the normalised composition

The normalised mole fractions $x'_i = x_i/S$ share the common divisor $S = \Sigma_j x_j$, which correlates them. NexChrom propagates the uncertainties through the normalisation according to the GUM, including that common-divisor covariance, with the component uncertainties as independent inputs:

$$\partial x'_i/\partial x_j = (\delta_{ij} \cdot S - x_i) / S^2$$
$$u^2(x'_i) = \Sigma_j (\partial x'_i/\partial x_j)^2 \cdot u^2(x_j)$$
$$= [(1 - 2 \cdot x'_i) \cdot u^2(x_i) + x'^2_i \cdot \Sigma_j u^2(x_j)] / S^2$$

Because this expression is linear in the uncertainties, expanded input uncertainties give expanded output uncertainties on the same coverage factor.

The three propagation paths, for direct components, indirect components and the normalised composition, have been checked against Monte Carlo simulation, with deviations below 0.2 %.

Summary of uncertainty scales

Quantity	Content	Scale
Certificate of a reference gas	Uncertainty of the certified value	Expanded, laboratory k (converted at import if the certificate states k = 1)
Regression input u(y)	Certificate uncertainty divided by k	Standard
Regression input u(x)	s/√n of the replicate peak areas	Standard
Coefficient covariance Σ_b	Per regression method, as described above	Standard
RRF uncertainty	Validated value, or ISO 6974-2 default per detector type	Standard (relative for the defaults)
Fixed component	Uncertainty as entered in the method	Expanded, laboratory k
Result of a direct or indirect component	Propagated combined uncertainty	Expanded, U = k·u

Quantity	Content	Scale
Analysis certificate	Uncertainty of the normalised composition	Expanded, laboratory k

Analysis certificates

An analysis certificate is generated from a stored calculation run. The certificate reports the **normalised composition** as certified values, in mol/mol, with expanded uncertainties on the laboratory coverage factor, propagated through the normalisation as described above. The unnormalised mole fractions are listed on the certificate as well, and the certificate states the regression method of the calibration as it is named in the software, for example EIV (ISO 6143).

What the certificate is built from is fixed when the run is stored for review on the Execute page: the coverage factor with which the uncertainties of the run were expanded, the rounding of the reported values and, when chosen, the page with the ISO 6976 properties and its reference conditions. The certificate is built from exactly these values, whatever the settings are later.

The certificate is issued on the Review page by the reviewer, together with the approval of the run, or later for a run that has already been approved; it is approved from the moment it is issued. NexChrom then writes it as a PDF and records its SHA-256 hash. The stored PDF is the certificate: it is frozen when it is issued and does not change when settings or data change afterwards. A preview on screen is generated live and may therefore differ from a certificate issued earlier. Each calculation run has at most one active certificate; a certificate that is withdrawn is issued again under a new number, and a stored file is never overwritten.

Because the certified values and uncertainties are stored on the same scale as those of an imported certificate (expanded, laboratory coverage factor), an analysis certificate from NexChrom can itself serve as the certificate of a calibration gas in NexChrom, without a scale conversion.

Optionally, a certificate carries an extra last page with the physical properties of the certified composition according to ISO 6976:2016, as described in the next chapter.

Physical properties according to ISO 6976:2016

From the normalised composition of a run, NexChrom calculates the physical properties of the gas according to ISO 6976:2016, each with its uncertainty by the analytical method of Annex B of that standard.

Properties

Property	Basis	Gas	Unit
Molar mass			kg/kmol
Compression factor	Metering conditions	Real	
Gross and net calorific value	Molar	Ideal = real	kJ/mol
Gross and net calorific value	Mass	Ideal = real	MJ/kg
Gross and net calorific value	Volume	Ideal and real	MJ/m^3
Density	Volume	Ideal and real	kg/m^3
Relative density		Ideal and real	
Gross and net Wobbe index	Volume	Ideal and real	MJ/m^3

On the molar and mass basis, ISO 6976 takes the real-gas calorific value equal to the ideal-gas value (clause 5); these are therefore listed once.

The composition and its correlations

The input is the composition of the run as it appears on the analysis certificate: the mole fractions x_i of all components, direct, indirect and fixed, normalised as $y_i = x_i / S$, with the standard uncertainties $u(x_i) = U(x_i) / k$. Because the normalised mole fractions add up to exactly one, they are correlated: what one gains, the others lose. NexChrom propagates the uncertainties through the normalisation, as for the certificate, to the full covariance matrix of the normalised composition:

$$J_{ij} = \partial y_i / \partial x_j = (\delta_{ij} \cdot S - x_i) / S^2$$
$$Cov(y) = J \cdot diag(u^2(x_1), \ldots, u^2(x_n)) \cdot J^T$$

Its diagonal is the uncertainty of the normalised composition on the certificate, and every row sums to zero. The correlation matrix derived from it enters the calculation as the normalisation matrix of ISO 6976:2016, 11.3.1 b). Leaving the correlations out would misstate the uncertainty of the properties. In Example 3 of Annex D of the standard, for instance, the standard uncertainty of the gross calorific value on a volume basis at 15 °C / 15 °C is 0.0163 MJ/m³ with the normalisation matrix and 0.0269 MJ/m³ when the mole fractions are treated as independent.

Components

Each component is linked to its entry in Table 1 of ISO 6976:2016 by its number in that table, and its name must be the Table 1 name of that number. A component without a number, or with a name that does not match its number, stops the calculation and is named; a wrong number can therefore never silently select another substance. A calculated mole fraction below zero also makes the ISO 6976 calculation impossible: the Execute screen reports it, and a certificate with the ISO 6976 page cannot be issued for that run. When the standard does not apply to a composition, NexChrom calculates no properties and states the reason.

Reference conditions

The reference conditions are chosen per run, within the conditions for which ISO 6976:2016 tabulates its data:

Condition	Choices	Default
Combustion temperature	0, 15, 20 or 25 °C, or 60 °F	25 °C
Metering temperature	0, 15 or 20 °C, or 60 °F	0 °C
Metering pressure	90 to 110 kPa	101.325 kPa

The defaults are a laboratory setting (Settings → System → ISO 6976 reference conditions). These are reference conditions, not measured ones: the ambient temperature and pressure that NexChrom can show from an optional Aura2 sensor are for display only and never enter a calculation. 60 °F is converted as exactly 15 5/9 °C (288.705 56 K), as ISO 6976:2016 prescribes in 3.12, Note 3. At a metering temperature of 60 °F, software that converts it as 15.55 °C gives every value on a volume basis $1.92 \cdot 10^{-5}$ (relative) higher; this explains small differences in a comparison.

Uncertainty and rounding

Every property, including the molar mass and the compression factor, is reported with its expanded uncertainty $U = k \cdot u$ on the laboratory coverage factor, where u is the standard uncertainty by the analytical method of Annex B. Annex B combines the uncertainties and correlations of the mole fractions with the uncertainties of the tabulated component data of the standard.

The results are rounded as ISO 6976:2016, 11.5.2 prescribes: the expanded uncertainty to two significant figures, the value to the same decimal place. The reference conditions, the basis of the correlations and the coverage factor are stated with the results.

On screen and on the certificate

On the Execute screen the properties are shown for information, for every run and every type of gas, including calibration and QC gases; the values shown there are not stored.

On an analysis certificate the properties appear only when this is chosen as the run is stored for review, with the reference conditions; the choice is stored with the run. The certificate then gets an extra last page, *Physical properties – ISO 6976:2016*, to which the first page refers. When the certificate is issued, the properties are calculated from the stored results of the run with its coverage factor, and the page is frozen with the PDF; the reference conditions and the unrounded values with their uncertainties are recorded with the certificate. If the properties cannot be calculated, the certificate is not issued. Whether a certificate carries the page cannot be changed once the run has been stored.

Validation

The calculation has been validated against ISO 6976:2016 itself:

- Every worked example in Annex D of the standard, at 15 °C / 15 °C, 60 °F / 60 °F and 25 °C / 0 °C, with and without the normalisation matrix, is reproduced within one unit of the last published digit, for the values and for the published uncertainties, without exception.
- The uncertainties for which Annex D gives no worked example (molar mass, compression factor and the ideal-gas properties) agree with the formulae of Annex B.
- The covariance matrix of the normalised composition is symmetric and positive semi-definite, every row sums to zero, and its diagonal equals the uncertainty of the normalised composition on the certificate.
- A component without a Table 1 number, or with a name that does not match its number, stops the calculation.

Standards and references

Standard or source	Role in the calculations of NexChrom
ISO 6143	Errors-in-variables regression; the goodness-of-fit test on the largest weighted deviation of a calibration point ($\Gamma \leq 2$) and the choice of the simplest admissible function; the worked examples of Annex B, against which the regression routine is validated
ISO/TS 28037	The χ^2 probability of the weighted residual sum, reported as information with an EIV fit
ISO 6974 (series), ISO 6974-2	Determination of natural gas composition by gas chromatography; default RRF uncertainties per detector type (Annex B)
ISO 6976:2016	Calorific values, density, relative density and Wobbe indices from the composition, with their uncertainties (Annex B); worked examples in Annex D
JCGM 100:2008 (GUM)	Evaluation and expression of uncertainty; standard and expanded uncertainty
ISO/IEC 17025	The framework of the laboratories NexChrom is designed for

- Deming, W.E. (1964). *Statistical Adjustment of Data*. Wiley, New York.
- Bremser, W. and Hässelbarth, W. (1997). Controlling uncertainty in calibration. *Analytica Chimica Acta* 348, 61–69.

- ISO 6143:2001, *Gas analysis — Comparison methods for determining and checking the composition of calibration gas mixtures*.
- ISO/TS 28037, *Determination and use of straight-line calibration functions*.
- ISO 6974-2, *Natural gas — Determination of composition and associated uncertainty by gas chromatography — Part 2: Uncertainty calculations*.
- ISO 6976:2016, *Natural gas — Calculation of calorific values, density, relative density and Wobbe indices from composition*.
- JCGM 100:2008, *Evaluation of measurement data — Guide to the expression of uncertainty in measurement* (GUM).

Further information

This document is a summary of the calculation specification of NexChrom. The full specification and the validation documentation, including the numerical validation of the EIV regression, of the uncertainty propagation and of the ISO 6976 calculation, are available to customers with the software, for use in their own software validation and audits.

NexChrom is in development; the release is expected in early 2027. This document describes the software as it is being developed; details may change before release. For questions about the calculations, or to register your interest: info@qualimetric.nl · nexchrom.nl