

CasaXPS Version 2.3.25

Introduction

Developments in software features and methodologies for applying these new and old features are described in the following text. The approach taken within this text is to describe an analysis for a specific experiment, then review CasaXPS features that were used during the preparation of the figures and results featured in the analysis section.

XPS of Sucrose

Experiment

Sucrose crystalline powder was measured by XPS using a Thermo K-Alpha+ at University of Cardiff. Spectra were collected over two binding energy intervals corresponding to C 1s and O 1s photoemission due to monochromatic x-rays with photon energy 1486.68 eV. The measurement as specified was repeated 101 times. Spectra recorded from sucrose (Figure Text1) are presented overlaid in Figure Text2. Charge compensation was used throughout these measurements resulting in apparent binding energy for C 1s and O 1s shown in Figure Text2. The evidence from overlaying C 1s and O 1s spectra in Figure Text2 is charge compensation has maintained a sufficiently stable uniform potential during the course of these measurements. It is also evident from these plots in Figure Text2 that C 1s spectra evolve in shape with elapse time for each measurement cycle. No other modifications to the sample were performed during these measurements. Therefore the variation for C 1s spectra apparent in Figure Text2 is due entirely to the environment within the instrument and the use of X-rays combined with charge compensation.

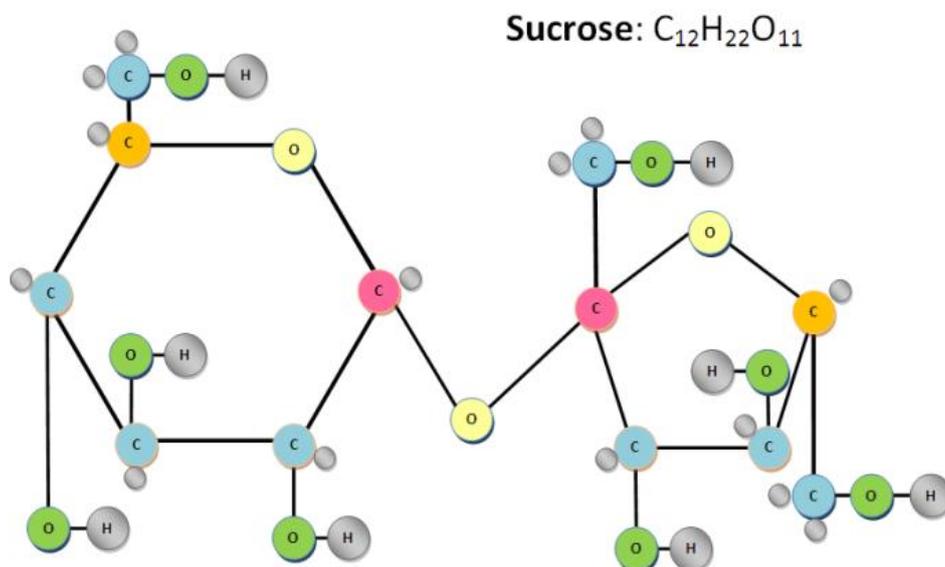


Figure Text1. Expected chemistry for sucrose.

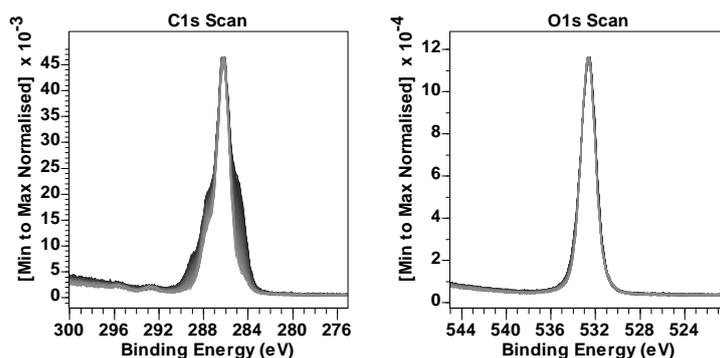


Figure Text2. C 1s and O 1s spectra measured from sucrose. These spectra represent the overlay of 101 C 1s spectra and 101 O 1s spectra measured sequentially.

Data Analysis

An investigation of chemical states for sucrose is typically performed using peak models constructed from a background approximation (to account for inelastic scattered intensity) upon which bell-shaped curves are used to fit shapes in spectra that can be attributed to chemical shifts to photoemission from a given element. It is clear from Figure Text2 that only C 1s offers shapes that might be open to analysis by peak models. O 1s signal from sucrose must include chemical state, but the energy resolution for data in Figure Text2 is insufficient to expect mathematical fitting of two components to a peak model with reasonable estimates for binding energy and relative intensities unless some forms of constraints are used to guide optimisation parameters to meaningful results. By contrast, C 1s spectra include variation in intensity of components that might be constructed making use of the understanding for the sample depicted in Figure Text1, where chemical binding energy offsets offer the prospect of demonstrating that the chemistry shown in Figure Text1 can be shown to be plausible by means of a peak model. However, such a peak model for C 1s is expected to apply only for spectra at the start of the experiment where changes to C 1s features are maintained with a degree of consistency. Further, the peak model must account for carbon chemistry not shown in Figure Text1 which is often assumed to be contamination by adventitious carbon or some form of degradation to the sample incurred before commencement of the current experiment.

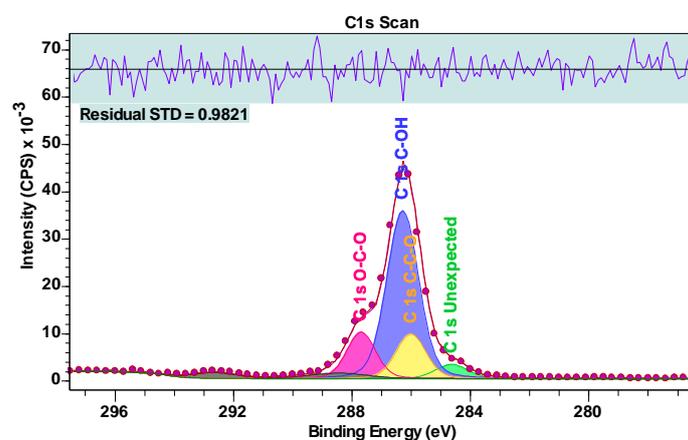


Figure Text3. Peak model for as-received sucrose C 1s spectrum recorded first in the sequence of spectra shown in Figure Text2. Bell-shaped components are designed with line shape LA(50) and area constraints are used to enforce intensity for three different bonds between oxygen and carbon shown in Figure Text1.

The peak model illustrated in Figure Text3 is one interpretation of the chemistry in Figure Text1. However, the chemistry shown in Figure Text1 does not include chemistry that would explain photoemission contributing to a peak with binding energy below the most intense peak in Figure Text3. Therefore, an additional peak is introduced to account for CH type chemistry not anticipated for a sample consisting of sucrose alone. Satellite peaks are labelled that are introduced to account for possible loss peaks due to the ring nature of sucrose illustrated in Figure Text1. Satellite loss peaks are often observed in polymers with ring structures. There are several other educated guesses involved in the design of the peak model in Figure Text3. An example is the position for photoemission associated with C-OH and C-C-O, which is not easily determined by optimisation as several equally plausible solutions can be often by changing the starting positions for these component peaks prior to optimisation. Area constraints are also used to force the expected ratio 8:2:2 for C-OH:C-C-O:O-C-O component areas. If contamination or sample degradation is already present in these data then area constraint assumptions may result in solutions with greater uncertainty than if all signal were derived from pure sucrose. We therefore reserve judgement of the efficacy of the peak model in describing the chemistry of the sample when applied to these as-received data in Figure Text3.

Reservations about the peak model shown in Figure Text3 are now investigated by considering the overall trends in C 1s spectra (Figure Text2) and applications of linear algebra to spectra which allow the derivation of spectral shapes that help explain any deficiencies in the model presented in Figure Text3.

Applying linear algebra to sets of spectra necessitates treating spectra as vectors in an n-dimensional vector space, where n is the number of data channels into which signal is distributed. In making this statement there is an assumption that the number of spectra used is less than the number of data channels defining each spectrum. It is also assumed that spectra are pre-processed to ensure signal within a given data channel in one spectrum has the same physical meaning in terms of binding energy as signal in the corresponding data channel for all spectra. This assumption may be restated in terms of measurement of spectra by assuming charge compensation has achieved an unaltered steady state potential experienced by photoelectrons for all spectra recorded. Failing to obtain adequate charge compensation uniformly for all measured spectra, the assumption requires that each individual measurement is performed with adequate charge compensation that may differ for each measurement, but it is possible to perform charge correction for signal in software to align signal with binding energy for all spectra. As stated above, data shown overlaid in Figure Text2 suggests these spectra are acquired under near identical charge compensation state and therefore analysis of these data will proceed on the assumption that spectra as-measured can be treated as vectors without the need for pre-processing of signal in software. In general such ideal measurements are not to be expected and careful charge correction allowing re-distribution of signal in terms of energy bins is required for analysis based on linear algebra.

The first application of linear algebra is to determine the dimension for a subspace that represents systematic shapes within the set of C 1s spectra in Figure Text2. The method used to make an estimate for the number of vectors required to approximate systematic shapes in C 1s spectra is principal component analysis (PCA). PCA applies a linear least squares criterion to determine an alternative set of vectors to the vectors formed from C 1s spectra in Figure Text2 which represent an orthogonal basis set for the subspace defined by all C 1s spectra in Figure Text2. That is, the as-measured C 1s spectra are assumed to be 101 linearly independent vectors by dint of random noise alone. Given 101 linearly independent vectors, Grams Schmidt ensures it is possible to create an orthogonal basis for the vector space spanned by a non-orthogonal set of vectors. PCA applied to a

set of vectors similarly constructs a mutually orthogonal set of vectors, however in the case of PCA rather than randomly selecting vectors to perform Grams Schmidt, PCA implements a least squares optimisation result to sequentially compute mutually orthogonal vectors in the order of greatest variance. The result of PCA is therefore a set of vectors in which systematic variations in signal intensity are preferentially selected for compared to random variations. Hence PCA applied to C 1s spectra in Figure Text2 identifies chemically influenced variations in signal and collects this information into a set of vectors which can be distinguish from random noise. In terms of XPS data, a subset of vectors chosen from these PCA determined vectors is a optimally selected by choosing the least number of vectors that allow raw spectra to be approximated by PCA vectors in a linear least squares sense with residual plots uniformly distributed about zero and a figure of merit computed from the normalised residual standard deviation close to unity. To illustrate the result of fitting spectra using three PCA determined vectors, Figure Text4 provides four examples of PCA signal enhanced data together with the residual plot and figure of merit for these approximations compared to raw data. While Figure Text4 illustrates typical results using three PCA vectors, the quality of fit is only marginally improved over the use of two PCA vectors (Figure Text5). The implication from these results shown in Figures 4 and 5 is that the set of C 1s spectra can be considered as a two-dimensional subspace representative of systematic shapes.

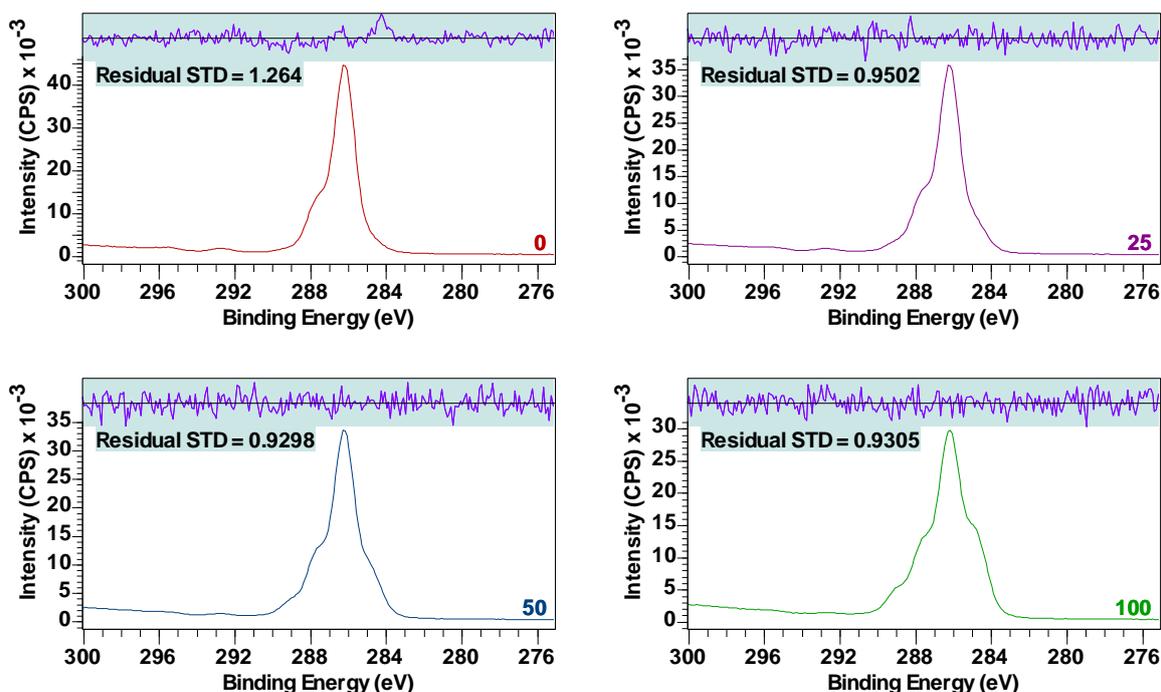


Figure Text4. Four examples of data reproduction based on fitting *three* most significant PCA vectors to spectra acquired at measurement cycles 0, 25, 50 and 100. Only the initial measurement cycle 0 is the residual plot and residual standard deviation greater than unity. In all other measurement cycles the residual plot and residual standard deviation are comparable to results for measurement cycles 25, 50 and 100.

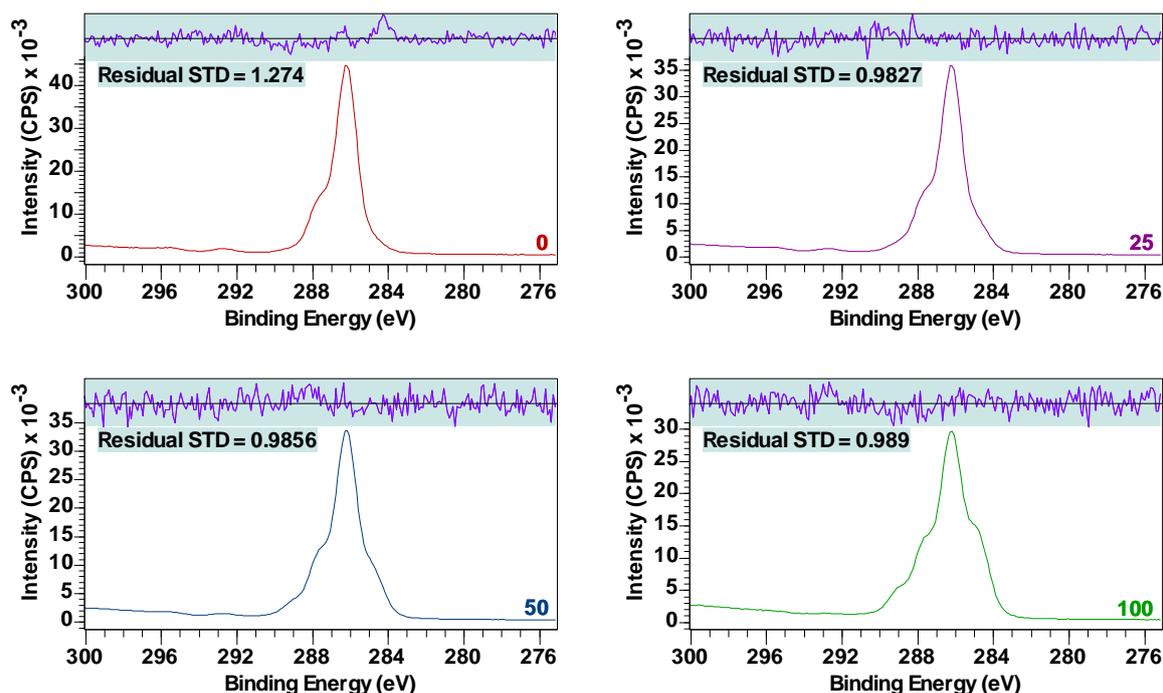


Figure Text5. Four examples of data reproduction based on fitting two most significant PCA vectors to spectra acquired at measurement cycles 0, 25, 50 and 100. Only the initial measurement cycle 0 is the residual plot and residual standard deviation greater than unity. In all other measurement cycles the residual plot and residual standard deviation are comparable to results for measurement cycles 25, 50 and 100.

The second use of linear algebra is to effect noise reduction in C 1s spectra. While two PCA vectors are sufficient to represent systematic shapes in C 1s spectra in Figure Text2, removing a proportion of noise from each spectrum prior to forming vectors by subtraction of vectors represents an advantage to the construction of difference spectra. The advantage is gained when these vectors with coordinate values that can be interpreted as signal are compared to spectral shapes. When PCA vectors are used to approximate spectra, in this example, the results in terms of noise are improved by making use of three PCA vectors to reconstruct the original data. A trial and error approach is used to determine the number of PCA vectors when approximating spectra. The consequence for the analysis that follows of different numbers of PCA vectors is subjective but often the choice for a suitable number of PCA vectors is apparent. An overestimate for the number of PCA vectors results in the reintroduction of noise to the reconstructed data. Since the process by which new vectors are calculated involves scaling and subtracting vectors, the best quality data is required because subtraction of near equal vectors is responsible for loss of significant digits in floating point machine numbers which amplifies noise relative to signal.

When presented with data such as C 1s spectra in Figure Text2, uncovering spectral shapes from these as-measured data involves selecting spectra two at a time v_n and v_m , then performing operations on these spectra that are in essence of the form

$$\text{Equation 1} \quad w_1 = \alpha v_n - (1 - \alpha)v_m, \text{ where } 0 \leq \alpha \leq 1$$

$$w_2 = \beta v_n - (1 - \beta)v_m, \text{ where } 0 \leq \beta \leq 1$$

The method relies on visual inspection (Figure Text6) of vectors w_1 and w_2 to determine spectral shapes that are consistent with expectations for sample chemistry (Figure Text7). In the case of C 1s spectra in Figure Text2, PCA determined that these spectra in the absence of noise all lie within a 2 dimensional subspace. Vectors w_1 and w_2 by their definition must lie within the same 2-dimensional subspace as v_n and v_m . It is therefore possible to calculate vectors $u_i = c_{i1}w_1 + c_{i2}w_2$ that approximate the original set of vectors $\{v_i: i = 0 \text{ to } 100\}$ (Figure Text8). By definition $v_n = u_n$ and $v_m = u_m$, but the success or failure of $v_i \approx u_i$ is the measure by which a solution is accepted or rejected. Figure Text8 illustrates four C 1s spectra sampled from the set of 101 C 1s spectra showing the result of fitting two difference spectra (calculated by these means) to four original spectra.

If a set of spectra equivalent to C 1s spectra in Figure Text2 lie within more than 2 dimensions, then the task is to identify subsets of spectra for which the subspace defined by the subset is of dimension two (as indicated by PCA applied to each subset of spectra). The process applied to C 1s spectra in Figure Text2 can then be applied to each subset to yield the first iteration of vectors that ultimately form the appropriate number of vectors to span the subspace which defines systematic shapes in the original data set. Each iteration involves forming vectors of the form Equation [1] using the set of spectra derived from subsets of spectra. Each iteration step involves testing least squares approximations to the original spectra. The objective is to identify the least number of vectors so determined such that good data reproduction through least square fitting of derived vectors to the original data is achieved.

The ultimate goal for the analysis of C 1s spectra is to identify bell-shaped components that can be assigned chemical state for carbon atoms. Figure Text9 is a set of peak models fitted to C 1s spectra that provide insight into the number of distinct components and possible chemical meaning that can be attributed to these component curves. These peak models are a combination of two peak models previously fitted to difference spectra which are assumed to have meaningful shapes. One distinguishing feature of these peak models in Figure Text7 is the presence of satellite peaks in the one assigned to pure sucrose, while the spectral shape thought to represent a material that is created by the measurement process appears to be without satellites. The implication is that ring structures in sucrose and is in some way damaged by the measurement process. Further, the binding energy for components in the degraded material suggests new bond states between carbon and oxygen are created. Input to the selections process based on observations relating to satellite peaks is an example of information used in identifying spectral shapes within the list of different spectra. Criterion such as the existence of satellite peaks for certain materials is very difficult to include in derivation of component curves purely based on mathematical heuristics. If we accept the outcomes from linear algebraic manipulations of spectra to be plausible, then the analysis of the entire data set therefore suggests the basic peak model in Figure Text3 is less precise than might otherwise be possible. In practise, constructing peak models and application of linear algebra to spectra both provide insight into chemistry that can be feedback into both approaches to allow iteration to a superior outcome for both.

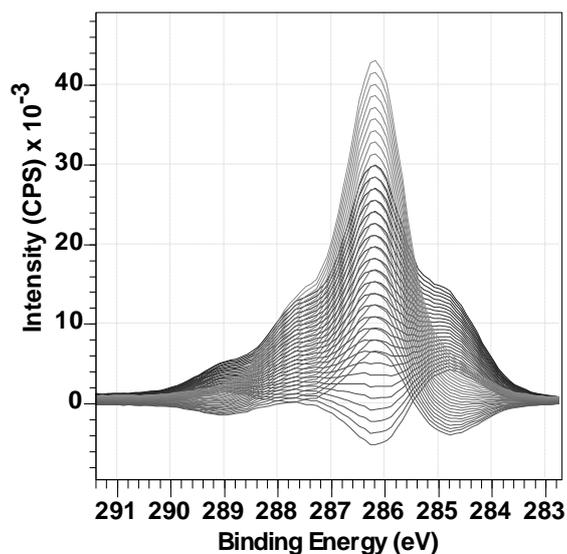


Figure Text6. Subset of difference spectra constructed using Equation 1. Within the full set of difference spectra are two spectral shapes with the potential for physical interpretation as shown in Figure Text7.

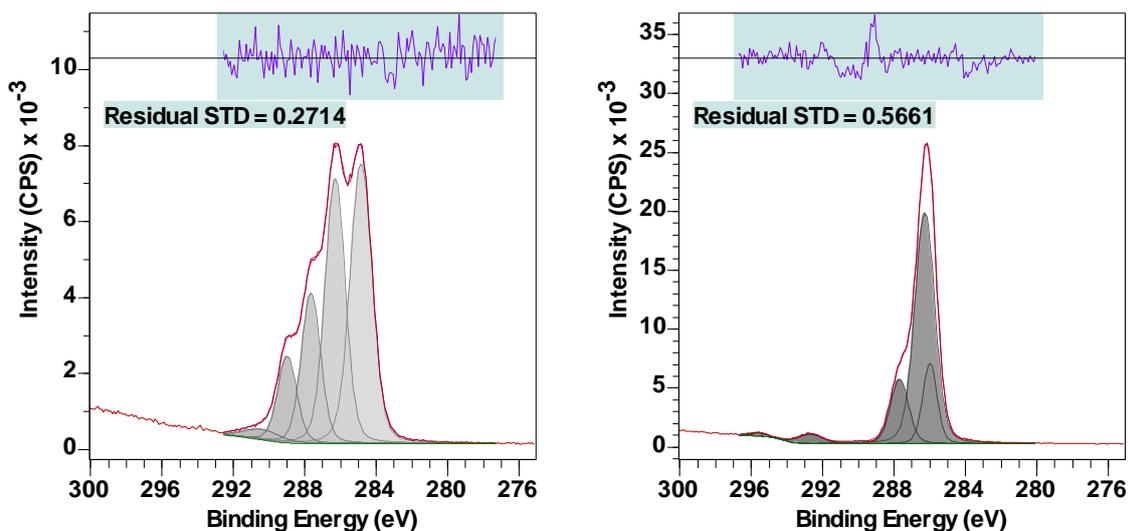


Figure Text7. Difference spectra identified from the full set of difference spectra corresponding to Figure Text6. Fitting bell-shaped curves to these two spectral shapes provides evidence that an initial sucrose sample degrades under the influence of measurement by XPS to produce a complex carbon chemistry (left). It is also possible to identify (right) a peak model with components compatible with the chemistry illustrated in Figure Text1 for sucrose.

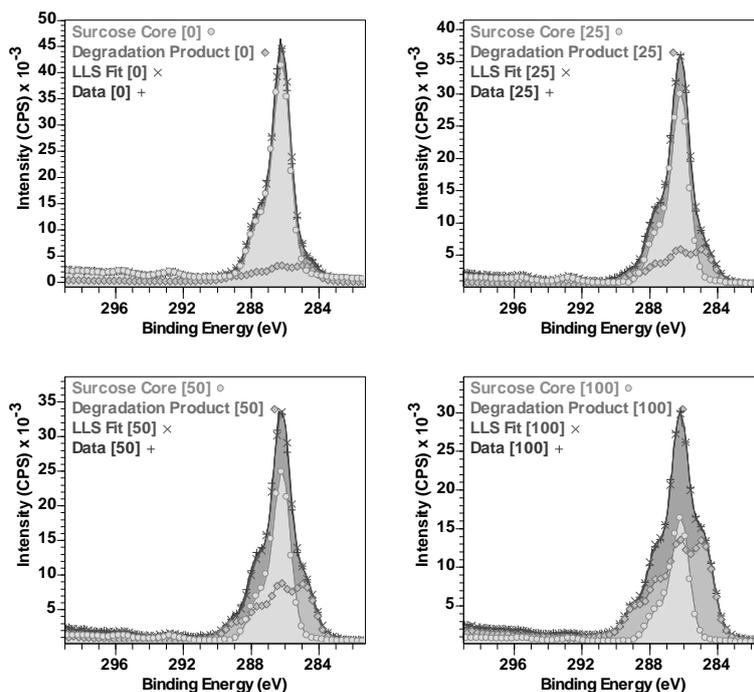


Figure Text8. C 1s Spectra from the data set in Figure Text2 fitted in a linear least squares sense using basis spectral components shown in Figure Text7. The peak models shown in Figure Text7 are scaled using the fits shown in this figure with the resulting peak models shown in Figure Text9.

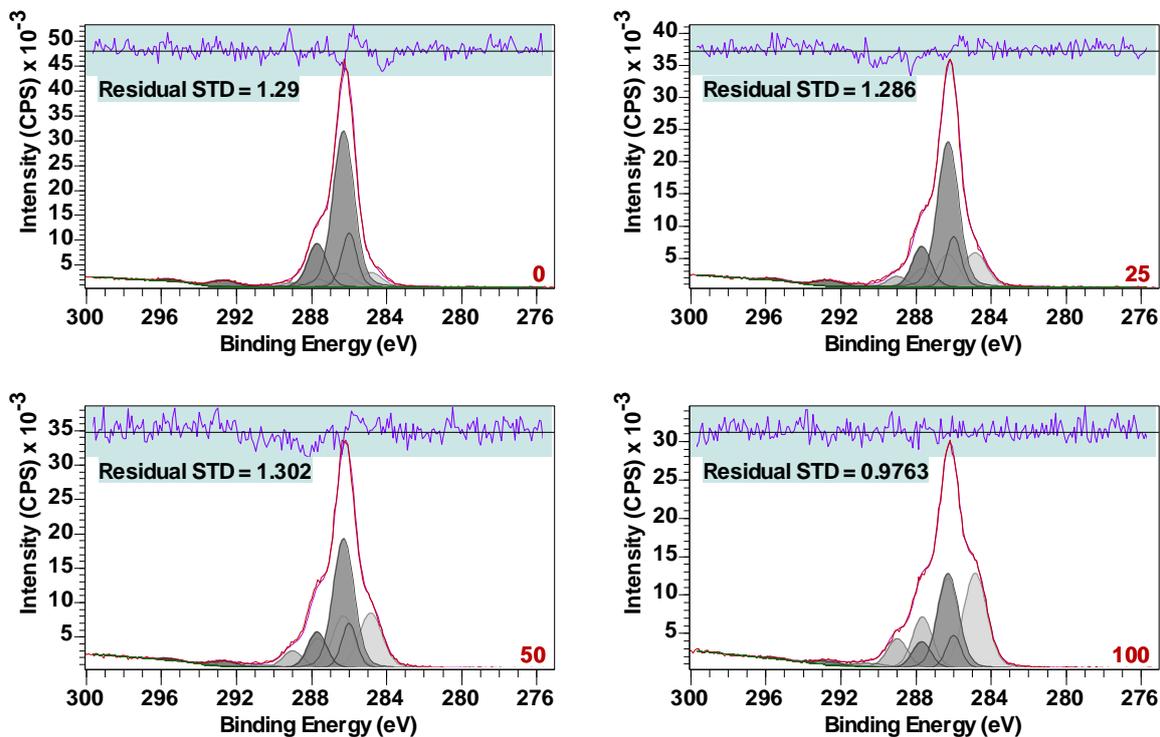


Figure Text9. Spectra shown in Figure Text8 fitted with bell-shaped curves based on peak models shown in Figure Text7. Numerical labels for each plot indicate the measurement cycle with the total experiment from which these C 1s spectra are sampled. Note that even the first C 1s spectrum measured labelled 0 includes a proportion of the degraded sample chemistry dominant in the last C 1s spectrum measured labelled 100.

CasaXPS Review of Sucrose Analysis

Converting and Opening Data Files

There are many file formats for which conversion to VAMAS format (ISO14976) is performed by CasaXPS. Conversion is possible for Ulvac PHI Miltipak binary files, Kratos Vision ASCII KAL files, Scienta-Omicron pxt files and a range of ASCII X/Y pair data formats. In the case of the sucrose data set, these data were acquired using a Thermo instrument and therefore the following explanation is focused on data from the Thermo data system Avantage.

Data collected during the measurement from Sucrose is exported from Thermo Avantage data system as ASCII files with file extension avg. The export option for Thermo Avantage data creates for each survey or narrow scan spectrum a file in a directory. CasaXPS converts a directory of these ASCII files to VAMAS format via the Convert option on the File menu or the Convert toolbar button. Once converted to VAMAS format by CasaXPS the newly created VAMAS file is opened in CasaXPS. In the case of data collected from sucrose the VAMAS file is presented as shown in Figure Sucrose1. The righthand pane offers a set of spectra arranged by experimental variable into rows of VAMAS blocks where columns of VAMAS blocks are distinguished by VAMAS block fields indicating the element and electronic state (e.g., C 1s) used to identify photoemission signal.

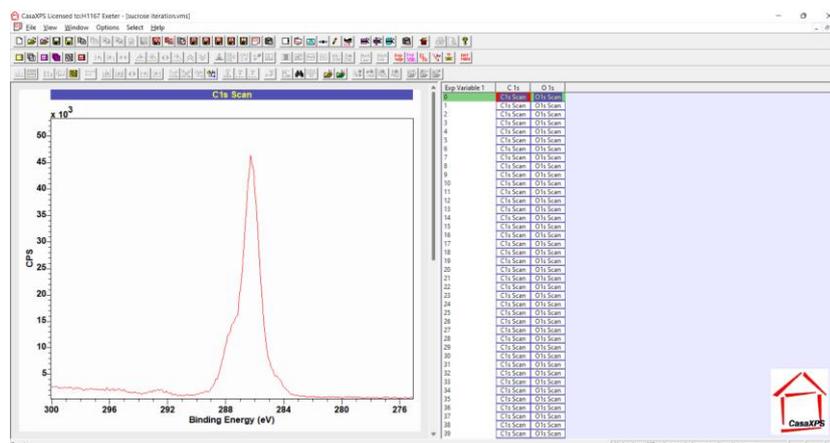


Figure Sucrose1. VAMAS file containing C 1s and O 1s spectra arranged as a sequence of measurements repeated over 101 acquisition cycles for each narrow scan spectrum for these two elements.

A video on the CasaXPS Casa Software YouTube channel illustrates how data files exported from Thermo Avantage as ANSI ASCII (UTF-8) are converted to VAMAS format.



<https://www.youtube.com/watch?v=izoP4LgfsbM>

Figure Sucrose2. YouTube video illustrating how Thermo Avantage data are exported from Avantage and then converted to VAMAS format by CasaXPS.

There are two options for exporting avg files which allow two ASCII formats, namely, UTF-8 and UTF-16. The latter is for an extended character set by using two-byte characters rather than one-byte characters as in the ANSI ASCII standard. If Thermo Avantage data are exported as Unicode UTF-16 format files, then these UTF-16 files must be converted to UTF-8 files by making use of a conversion filter .UTF. The YouTube video demonstrates how to convert a directory structure containing UTF-8 avg files via the conversion filter .DTH. The action of .UTF is identical to .DTH except that avg files in UTF-16 format are converted to UTF-8 before conversion to VAMAS format is performed. Note that UTF-16 format files are twice the size of UTF-8 format files.

Display of Spectra

Sucrose data are displayed in the left-hand pane in Figure Sucrose1. The left-hand pane initially displays via a scrolled-list of display tiles VAMAS blocks that appear in the first row of VAMAS blocks as shown in the right-hand pane in Figure Sucrose1. Only one display tile is visible in the left-hand pane, however the scroll-bar associated with the vertical direction indicates that more than one display tile is currently loaded into the left-hand pane. As the display currently stands in Figure Sucrose1, the left-hand pane allows the display of only one display tile at a time. However, the display arrangement in Figure TEXT2 is an example of exported graphical contents from CasaXPS making use of two display tiles per page of display tiles. The options used to prepare the display format shown in Figure TEXT2 will now be described.

The arrangement of two display tiles, each display tile presenting an overlay of all spectra from a column of VAMAS blocks as seen in Figure Sucrose1 is achieved by making use of the Page Tile Format dialog window. Figure Sucrose3 shows the state for the Page Tile Format dialog window and the display tiles appearance prior to application of the parameters as indicated by the Page Tile Format dialog window. Following application of the Page Tile Format dialog window parameters shown in Figure Sucrose3 the display changes to the state shown in Figure Sucrose4.

The next step to creating the display of data seen in Figure TEXT2 is entering C 1s spectra overlaid in the left-most display tile and O 1s spectra in the right-most display tile. The reason there are two display tiles as seen in Figure Sucrose4 is that the VAMAS file as displayed in the right-hand pane is arranged into two columns and on opening in CasaXPS the default action is to select the first row of VAMAS blocks and display data in these selected VAMAS blocks one spectrum per tile. The same display state for a pair of C 1s and O 1s spectra can be achieved by making the selection of two VAMAS blocks before pressing the toolbar button indicated in Figure Sucrose4. The indicated toolbar button (Display one per tile) acts on the selected VAMAS blocks and for each VAMAS block selected a display tile is created in the left-hand pane and data from the selected VAMAS blocks appear in these tiles, one spectrum per tile. To create a display in which spectra appear overlaid as seen in Figure TEXT2 there are two options which will both yield the same result of two tiles into which C 1s are overlaid in one tile while O 1s are overlaid in the other tile. The first of these ways makes use of a general strategy for overlaying selected data in a display tile. That is, a set of spectra selected in the right-hand pane are overlaid in a tile in the left-hand pane by pressing the toolbar button indicated in Figure Sucrose5. The action of the toolbar button targets the display tile that is the currently active display tile. The active display tile is the display tile for which the title to the display tile is shown in highlight colour as indicated in Figure Sucrose5. The result of pressing the toolbar button is to display overlaid all C 1s spectra in the left-most display tile as shown in Figure Sucrose5. The process of selecting O 1s VAMAS blocks, giving the active state for a display tile to the right-most tile in the left-hand pane, then pressing the toolbar button a second time overlays the O 1s spectra. However, there is an alternative way to create two display tiles at the same time as displaying data. This second option is to make use of a toolbar button that appears on the SIMS toolbar, which

applies to selected VAMAS blocks on a column basis for selected VAMAS blocks in the right-hand pane. Figure Sucrose6 indicates the toolbar button which applies to columns of selected VAMAS blocks. A new display tile is created for each column of VAMAS blocks for which at least one VAMAS block is selected. For each new display tile VAMAS blocks from columns with selected VAMAS blocks are overlaid. Figure Sucrose6 shows two display tiles created by pressing the toolbar button after selecting all VAMAS blocks in the right-hand pane.

The SIMS toolbars and other toolbars can be enabled or disabled using options on the View menu for the Main CasaXPS window.

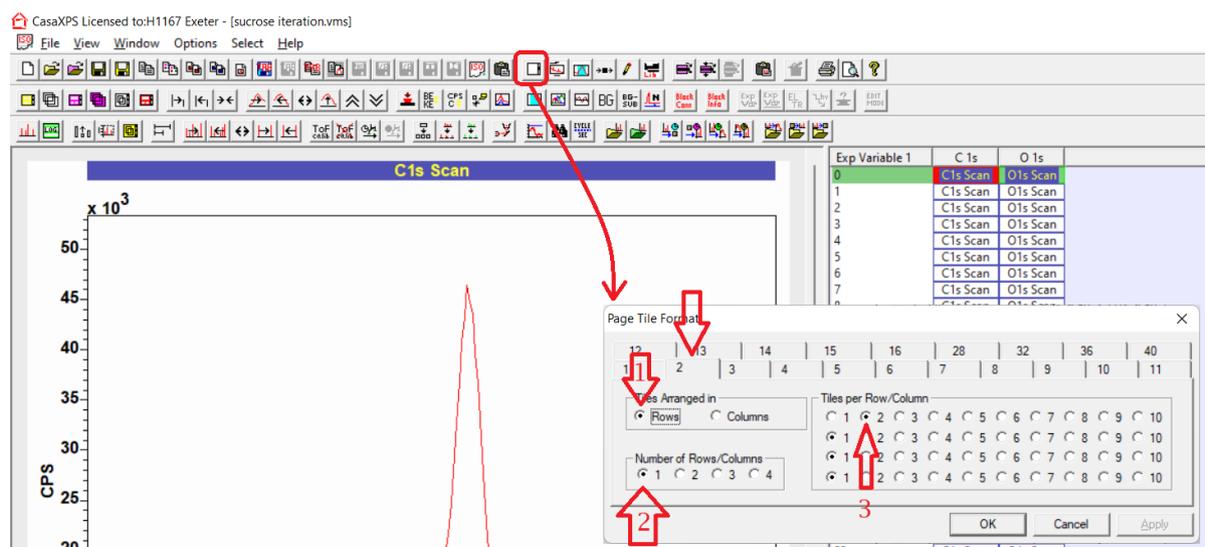


Figure Sucrose3. Page Tile Format dialog window configured to display two display tiles per page of tiles.

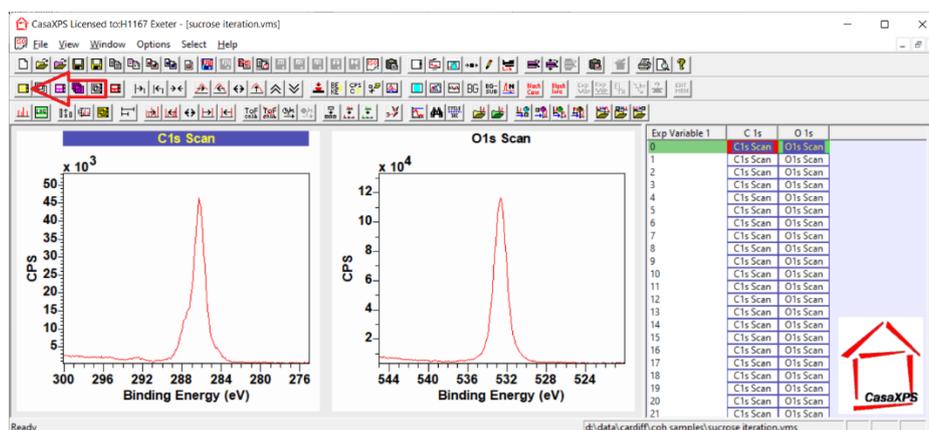


Figure Sucrose4. Result of accepting the Page Tile Format dialog window configuration shown in Figure Sucrose3, labelled 1, 2 and 3, namely, 1) parameters are specified as arrangement in rows of display tiles, 2) the number of rows is specified as one row and 3) the number of tiles in each row is specified to be two tiles per row.

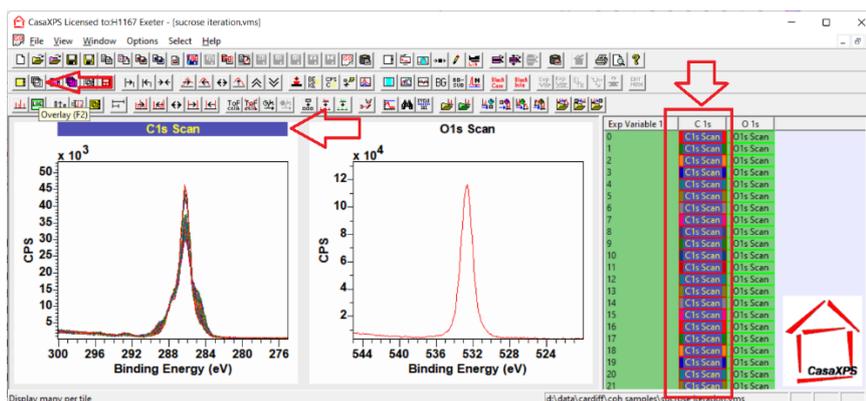


Figure Sucrose5. Overlay (F2) toolbar button makes use of selected VAMAS blocks and displays data from these selected VAMAS blocks in the display tile for which the title string is displayed in highlight colours.

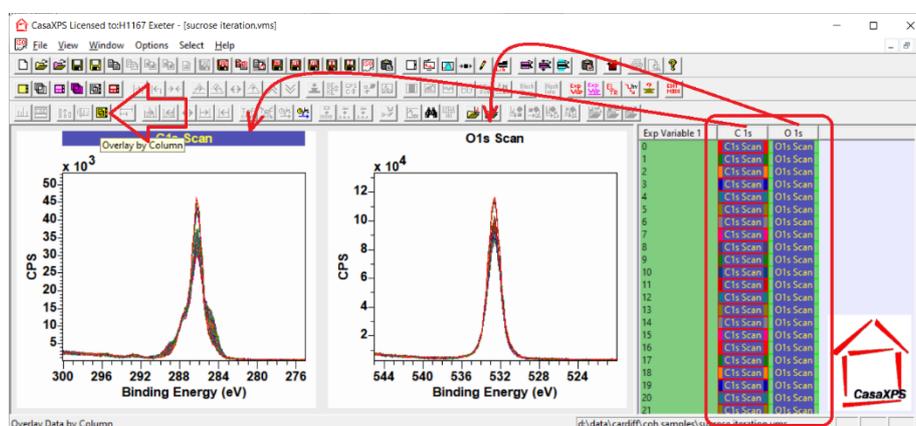


Figure Sucrose6. SIMS toolbar display option used to overlay spectra in display tiles which are created by pressing the toolbar button (indicated by the arrow) for each column of VAMAS blocks with a selected VAMAS block in the right-hand pane. The SIMS toolbar can be displayed using a menu option in the View menu of the CasaXPS main window.

The display of overlaid spectra in Figure Sucrose6 is using a palette of sixteen user-defined colours to display spectra. When 101 spectra per tile are displayed in the active tile, the sequence of sixteen colours is repeated for the display of spectra. The colour assigned to a spectrum when overlaid in the active tile is used to colour the VAMAS block corresponding to that spectrum in the list of VAMAS blocks in the right-hand pane. These colours can be modified using options on the Colours property page of the Tile Display Parameters dialog window (Figure Sucrose7).

Spectra displayed in Figure TEXT2 make use of colours from a predefined colour palette rather than the user-defined palette of sixteen colours shown in Figure Sucrose7. Predefined colour palettes contain 256 colours that are specified through the selection of a radio button on the Colours property page of the Tile Display Parameters dialog window (Figure Sucrose8).

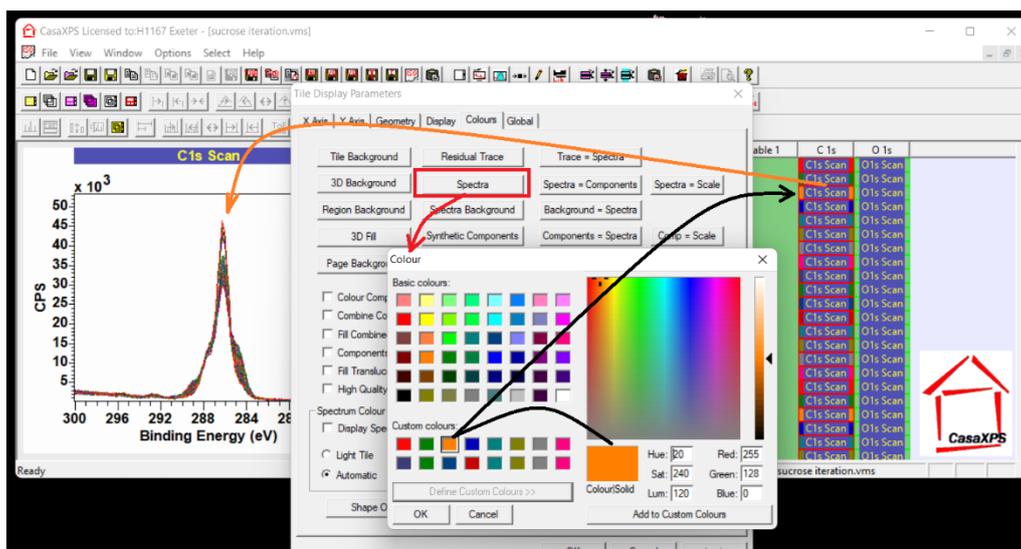


Figure Sucrose7. Editor used to adjust the palette of sixteen colours used to display spectra when overlaid in a display tile. The alternative to making use of sixteen user-defined colours is to make use of predefined colour scales which allow up to 256 predefined colours when displaying spectra.

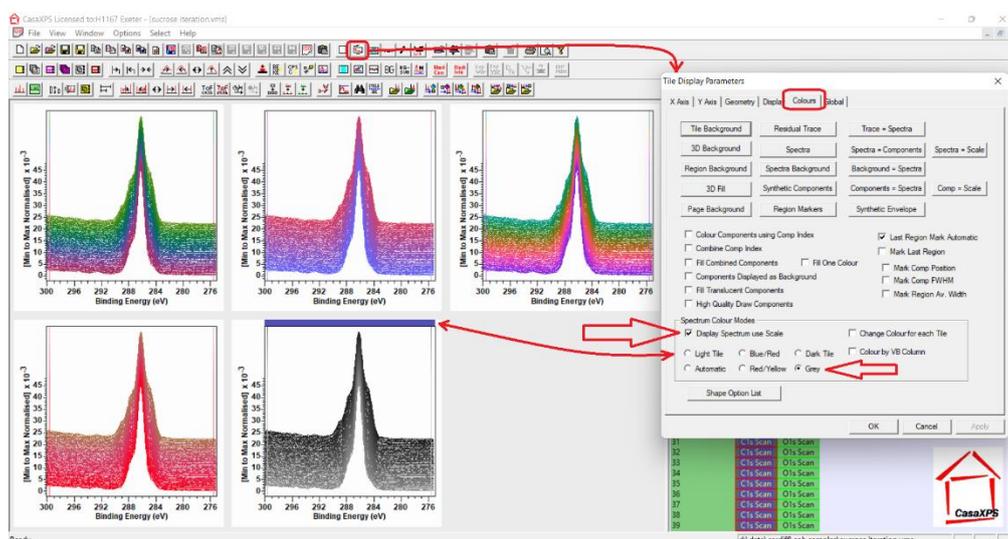


Figure Sucrose8. Predefined colour scales for the display of overlaid spectra are enabled using the tick-box labelled “Display Spectrum use Scale”. One of five different colour scales are available for selection via radio buttons.

Modelling Photoemission and Fitting Curves to Data

The peak model in Figure TEXT3 is an example of modelling data as a set of well-defined mathematically defined bell-shaped curves where each curved is formulated to measure signal associated with a specific photoemission process. Inelastic scattered background contributions to spectra are modelled using one of the many standard curves calculated from data. The paradigm embodied in the model in Figure TEXT3 is theoretically and practically plausible for photoemission from carbon bonded to atoms in different environments. Extrapolation for the use of this paradigm to all photoemission spectra is not always as appropriate, but there are sufficient number of examples where photoemission signal is reasonably modelled using the techniques used in constructing Figure TEXT3 that a detailed description of these features is fully justified. The following focuses on software features that aid and facilitate fitting data as seen in Figure TEXT3.

The peak model in Figure TEXT3 is designed assuming bell-shaped curves are appropriate for photoemission from carbon and that the correct number of components to the peak model which are defined in terms of these bell-shaped curves have been identified. The decision about the number of components to the peak model in Figure TEXT3 is based on the theoretical understanding of the sample (namely, sucrose) illustrated in Figure TEXT1. Implementing the chemistry defined in Figure TEXT1 in terms of the peak model in Figure TEXT3 is performed via the Quantification Parameters dialog window (Figure Sucrose9).

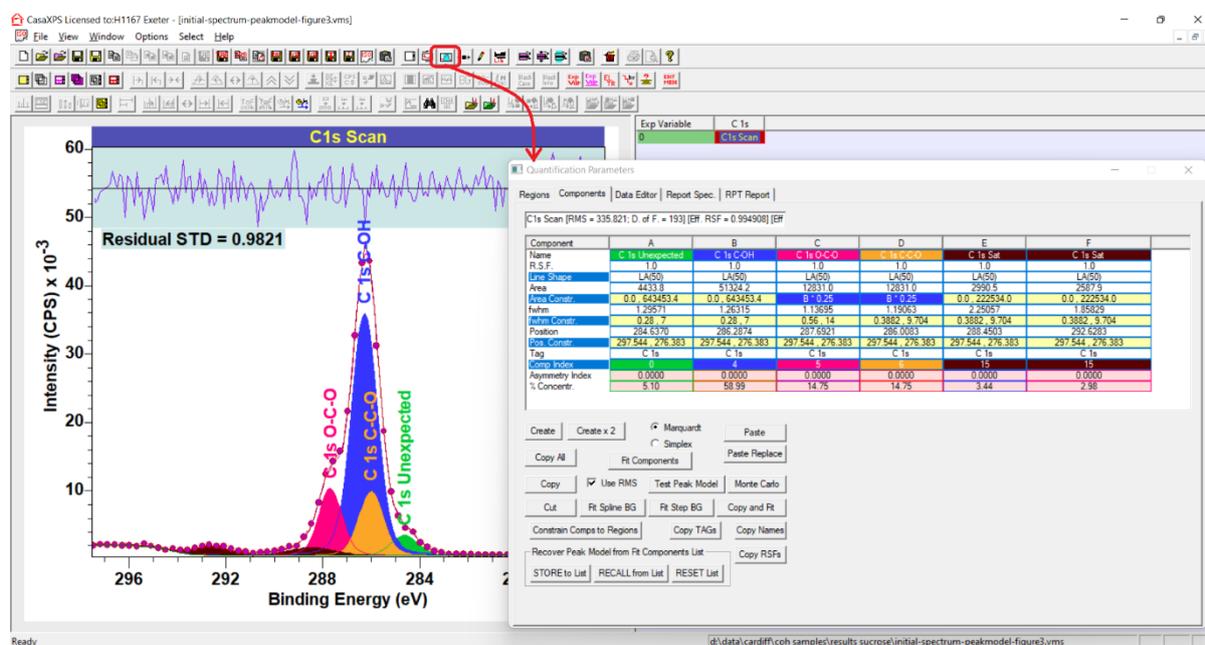


Figure Sucrose9. Wide form for the Quantification Parameters dialog window displaying via the Components property page the table of components for the peak model as shown in the active display tile. Components plotted in the active display tile in the left-hand pane are colour coded according to the Comp Index values using a display option that fills each component curve using one colour for the lines defining the curve and the fill colour indicating the area calculated from the curve. The components plotted in Figure TEXT3 differs from this display mode for components.

Components to a peak model are defined using a table of parameters which can be edited via the Component property page of the Quantification Parameters dialog window (Figure Sucrose9). The current values for parameters are displayed for each component as a column of entries, where each column is assigned a character label. These parameter values can be edited by left-clicking the mouse with the cursor over an entry. If an entry can be edited, then the box containing the current value becomes an edit box. Modified entries are accepted as the current parameter value by pressing the return keyboard key. In particular, the parameter value for each component used in one context to specify the colour used to plot component curves (line shapes) is defined for Figure Sucrose9 to illustrate how components can be plotted using a single colour for the component outline and also the fill colour. The outline colour is specified via Comp Index parameter which is an index value between -1 and 15, where 0 to 15 specify the user-defined colour palette as shown in Figure Sucrose10.

the area constraints appear in Figure Sucrose9 with the colour used to plot the component in column B highlights the relational constraints between components in columns C and D are relative to the area for the component in column B.

Constraints to optimisation parameters are a way to solve problems in fitting curves to data that can best be described as problems caused by lack of an obvious unique fit to data. Constraints attempt to limit adjustments to component parameters and in so doing reduce the number of possible mathematical solutions available to optimisation. C 1s data from sucrose is a good example for which unconstrained optimisation yields a set of outcomes without an obvious correct answer. That is, photoemission from C-O and C-OH does not yield components in a peak model for which the binding energy for these chemical states is well-defined by data alone. For this reason, constraints are important for peak models such as the one in Figure TEXT3. The reason for this importance of constraints is because the components correspond directly to chemical state for carbon in sucrose with similar line shape and relatively small shifts in binding energy due to chemical state.

Since the scenario for C 1s photoemission in sucrose is repeated for many other samples analysed through the use of peak models, a set of keywords are defined in CasaXPS which help define constraints between optimisation parameters across components defining a peak model. For example, a common choice in the initial stages when developing a peak model for carbon is to constrain all C 1s components to have the same FWHM. To achieve a uniform FWHM for a set of components to a peak model, the keywords "link all" can be entered into the constraint field for FWHM in one component. On pressing the return keyboard key, relational constraints for all components of the form that involves the label character at the top of the column (on the Components property page into which the string "link all" is typed) followed by "* 1". Thus, if the component with parameters defined by column C, say, has the string "link all" entered into the FWHM constraints field, on pressing return all other FWHM constraints will be set to "C*1" (Figure Sucrose12). Follow the use of "link all", if the value for FWHM in the component in column C is adjusted, then the FWHM in all components will adjust and have the same value specified for FWHM in column C.

Keywords are used throughout the fields on the Components property page of the Quantification Parameters dialog window. While to this point the focus has been on optimisation parameter constraints keywords are used to aid the development of peak models to include copying line shapes between components and other often performed tasks during the creation of peak models. The following section provides a table of keywords available in CasaXPS 2.3.25.

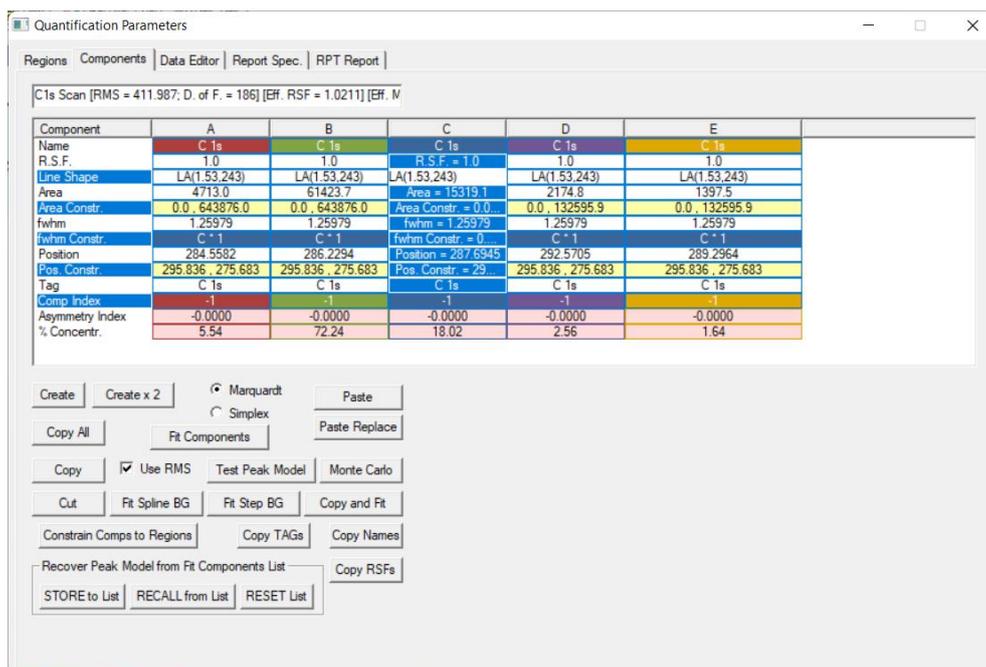


Figure Sucrose12. Components property page of the Quantification Parameters dialog window (wide format) showing the result of entering the key string “link all” into the FWHM constraints field.

Component Property Page Keywords (defined for CasaXPS 2.3.25 and above)

Component Name Field:

#<name string>	If the string <name string> matches an element library name string (e.g. Ti 2p or O 1s) then the entries from the element library are used to update the component parameter fields to include 1) RSF, 2) line shape and 3) TAG string which is set to the <name string> for the component.
=<name string>	Update Component Name Field for all components WITH THE SAME Comp Index with the string <name string>.
==<name string>	Update Component RSF for all components WITH THE SAME Comp Index. Name fields in other components with the same comp index are unaffected.

Component RSF Field:

#	If the string current name field for the component matches an element library name string (e.g. Ti 2p or O 1s) then the RSF entry from the element library is used to update the RSF for the component only.
@	If the string current name field for the component matches an element library name string for doublet-pair peaks (e.g. Ti 2p or Al 2p)

	<p>then the RSF entry from the element library is used to update the RSF for the component adjusting the library RSF according to the ratio $(2j_{\max}+1 / 2j_{\min}+1)$.</p> <p>Adjusting the RSF in this way allows for an element library in which the total RSF for doublet photoemission is available but quantification is performed based on the more intense peak of doublet photoemission.</p>
@-	<p>If the string current name field for the component matches an element library name string for doublet-pair peaks (e.g. Ti 2p or Al 2p) then the RSF entry from the element library is used to update the RSF for the component adjusting the library RSF according to the ratio $(2j_{\min}+1 / 2j_{\max}+1)$.</p> <p>Adjusting the RSF in this way allows for an element library in which the total RSF for doublet photoemission is available but quantification is performed based on the less intense peak of doublet photoemission.</p>
==<RSF>	Update Component RSF for all components WITH THE SAME Comp Index using the value <RSF>.

Component Line Shape Field:

#<line shape>	Update line shape for all components WITH THE SAME Comp Index.
=<line shape>	Update line shape in ALL components.
reset	<p>Resets to empty the list of line shapes previously used and therefore available via the Line Shape dialog window list.</p> <p>The set of line shapes previously used are cached and available for use via the Line Shape Dialog window. Line shapes from data are added to the cache of line shapes when first used and therefore the shape for a line shape is fixed at the time the line shape is defined. Adjustments to data used to create a line shape are not reflected in the line shape previously defined and therefore changes to data (altered background for example) require the reset action to re-initialise a line shape from data.</p>
tp	Creates a line shape with type TP(1,1,0). The position for the component is fixed at the current value by adjusting the position constraint field. The value of the FWHM and

	FWHM constraints is also adjusted to provide a wide interval constraint suitable for tp line shapes used as a linear background component to a peak model.
--	--

Component Area Constraint Field:

-	The current area for a component is reversed in sign and the constraint field is adjusted to account for a switch from a positive component area to a negative component area or vice versa.
lock	Calculate and apply relational area constraints for all components with the same comp index. For each component with the same index the area constraint is made relative to the component into which the lock keyword is applied. Relational constraints are of the form <letter> * <value> where <letter> is the character that appears in the header button on list of components displayed on the Components property page.
lock all	Same action as lock keyword but applied to ALL components, not just components with identical comp index values.
factor = <value>	Applies relational constraints to components with the same comp index as the component to which the factor keyword is applied. For example, factor = 2 applied to two components in columns with header letters A and B via the component in column B creates the relational constraint to the component in column A of the form B * 2. Existing relational constraints are maintained.

Component FWHM Constraint Field:

lock	Calculate and apply relational FWHM constraints for all components with the same comp index. For each component with the same index the FWHM constraint is made relative to the component into which the lock keyword is applied. Relational constraints are of the form <letter> * <value> where <letter> is the character that appears in the header button on list of components displayed on the Components property page.
lock all	Same action as lock keyword but applied to ALL components, not just components with identical comp index values.

link	Sets relational constraints for components with the same comp index of the form <letter> * 1. As a consequence of using the keyword link, components are forced to have the same FWHM as the component into which link is applied.
link all	Same action as link keyword but applied to ALL components, not just components with identical comp index values.

Component Position Constraint Field:

lock	Calculate and apply relational position constraints for all components with the same comp index. For each component with the same index the position constraint is made relative to the component into which the lock keyword is applied. Relational constraints are of the form <letter> + <value> where <letter> is the character that appears in the header button on list of components displayed on the Components property page.
lock all	Same action as lock keyword but applied to ALL components, not just components with identical comp index values.
reverse	Given a set of components with relational constraints with the same comp index, the action of the reverse keyword is to convert positive offsets to negative and vice versa.
regions	Regions are used to define the background to photoemission peaks. Application of the regions keyword adjusts interval constraints of the form <lower>,<upper> to ensure constraints for position limit the position of components to be within the limits defined for each region.
Offset = <value>	An increment is applied to component positions for components with identical comp index. The increment specified by <value> creates a set of equally spaced components.

Component Comp Index Field:

set=<value>	All components with the same comp index are assigned the comp index value specified by <value>.
step	Components with the same comp index are assigned unit increments in comp index starting with the value of the comp index for the component in which the keyword step is entered for the Comp Index field.

group=<value>	Sets of components with the same comp index are assigned the same comp index value after which a unit increment occurs. Arranging components representing doublet peaks sequentially on the Components property page all with the same comp index allows group=2 to assign the same comp index to components from the same doublet while incrementing the comp index for the subsequent doublet pairs.
reset	All components are assigned to the default comp index. The result is all comp index values are set to -1.
lock	Calculate and apply relational constraints to each fitting parameter (Area, FWHM and Position) for all components with the same comp index. For each component with the same index constraints are made relative to the component into which the lock keyword is applied. For Area and FWHM relational constraints are of the form <letter> * <value> while for position constraints are of the form <letter> + <value>, where <letter> is the character that appears in the header button on list of components displayed on the Components property page.

Component TAG Field:

set=<string>	Assigns the same TAG string to all components with the same comp index.
--------------	---

Spectrum Processing

The analysis of sucrose in terms of the peak model in Figure TEXT3 is a relatively simple example of fitting mathematical shapes to XPS data as a means of matching spectral characteristics to chemical state for an atom. The model of using bell-shaped curve to represent photoemission from a particular atom and electron configuration is widely used in XPS, however the use of synthetic line shapes in peak models is open to misuse, mostly because optimisation guided only by bell-shaped components are capable of returning mathematical fits to data that overlook shapes in data that are important to physics but not to the mathematics of minimising a figure-of-merit. The apparently adequate peak model and results in Figure TEXT3 are called into question by the analysis that follows, where spectrum processing is applied to the entire set of C 1s spectra shown in Figure TEXT2 to suggest alternative explanations for spectral shapes in Figure TEXT3.

Sucrose and the repeat measurement by XPS of sucrose are chosen for this manual because the sample chemistry, in theory, is well known (Figure TEXT1) and it is also easy to see from the data displayed in Figure TEXT2 that spectral shapes alter systematically with each additional measurement. While at first sight these evolving spectra appear to be a problem for sample analysis, these changes in spectral form are an advantage provided these changes to spectral shapes are

recognised. For example, observing that sucrose spectra change with each measurement begs the question, is the spectrum selected for the peak model in Figure TEXT3 free from these types of changes? Put another way, are we sure the peak identified as CH-type carbon chemistry really due to contamination of the sample? Answering these questions involves probing these data in Figure TEXT2 using options available on the Spectrum Processing dialog window via the PCA and Calculator Property pages.

The first question to answer when processing data similar to the C 1s data in Figure TEXT2 is, how many different mathematically distinct shapes are required to describe the original spectra? First, it is worth considering what is meant by *mathematically distinct shapes* and how one might measure if two shapes are mathematically distinct. Visually and qualitatively two shapes are mathematically distinct if when normalised and overlaid in a display tile these two shapes do not overlay exactly. Quantitatively, two shapes are mathematically distinct if the angle calculated by calculating the dot or scalar product of these two shapes is nonzero. The greater the angle between two shapes the more dissimilar are the shapes. If the angle between two shapes is 90° then these two shapes have nothing in common. There is an option on the Calculator property page of Spectrum Processing dialog window (Figure Sucrose13) that reports the angle between to spectra overlaid in the active display tile. However, given a set of 101 spectra (which is the data set in Figure TEXT2) determining the angle between pairs of spectra for the entire data set is not feasible. Rather PCA is used to calculate from spectra abstract factors that are shapes which are mathematically equivalent to the original spectra, but represent a sequence of shapes that are determined in terms of the maximum variation in the data set and ordered according to the scale factors for variance. Underlying PCA is computing the eigenvector for a covariance matrix which is formed from matrix elements that are effectively measuring the angle between all spectra in a data set. Thus, PCA is a way to make use of pair-wise analysis of spectra in terms of angles. PCA applied to these C 1s spectra en masse is performed via the PCA property page on the Spectrum Processing dialog window (Figure Sucrose14).

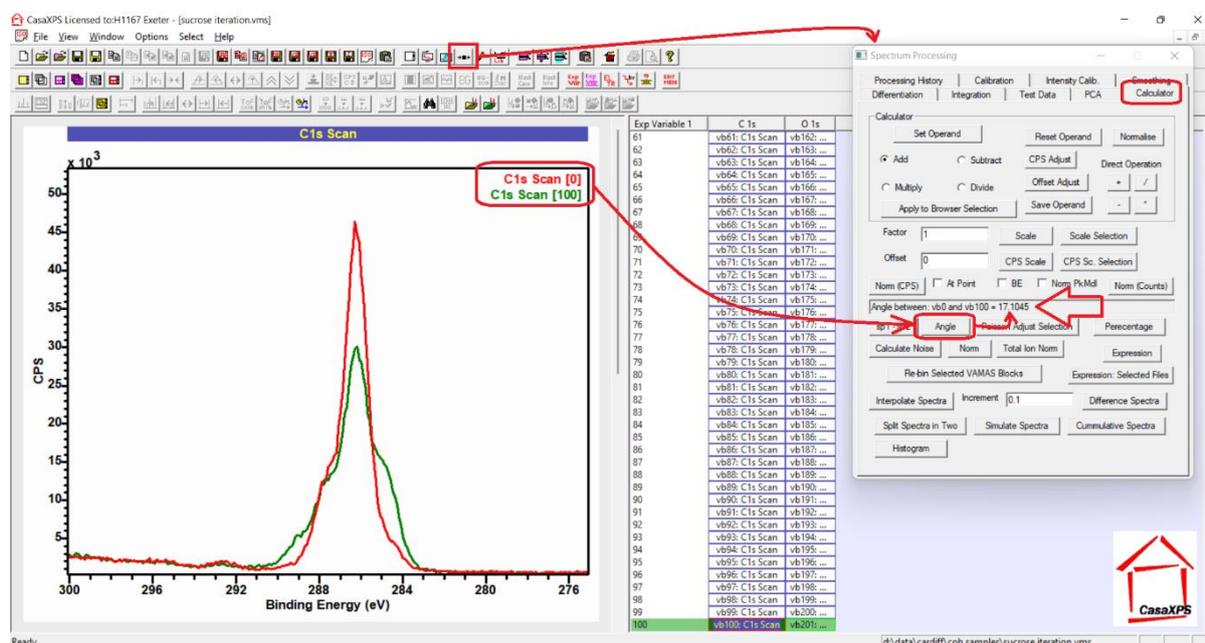


Figure Sucrose13. Calculator property page is used to calculate the angle between two spectra. Two spectra overlaid in the active tile are the inputs to the button labelled Angle, which calculates the angle between these two spectra. The result is displayed on the Calculator property page and also added to the VAMAS block comment for the active spectrum in the active tile.

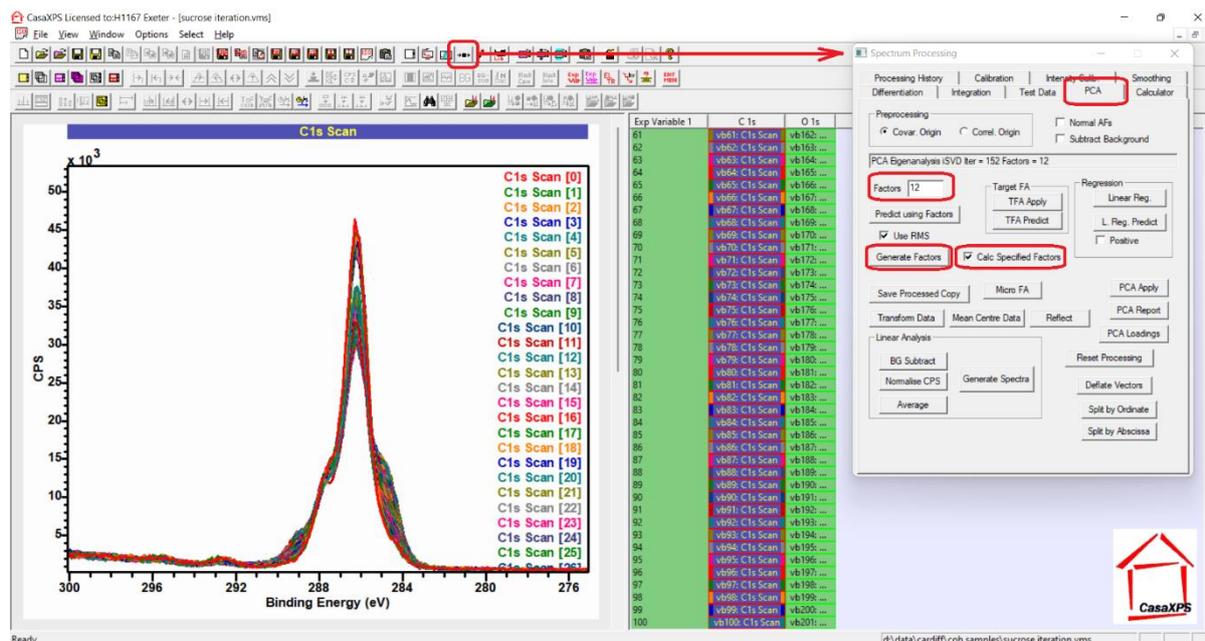


Figure Sucrose14. PCA property page on the Spectrum Processing dialog window.

PCA applied to a set of 101 spectra is achieved by logically forming a covariance matrix which is of size 101 x 101 rows and columns. Determining all eigenvectors, from which abstract factors are calculated, can be time consuming and unnecessary in many cases. These sucrose C 1s spectra displayed overlaid in the active tile in Figure Sucrose14 are a good example of where it is unnecessary to calculate all eigenvectors required for a full PCA. The requirement to calculate an abstract factor is related to whether or not any given abstract factor contains shape information different from shapes expected for noise in spectra. Figures TEXT4 and TEXT5 illustrate that three or two abstract-factors when used to fit spectra both yield approximations that deviate from spectra with residuals and figure-of-merit close to what is expected for any approximation to these pulse-counted intensities within C 1s spectra. The implication from Figures TEXT4 and TEXT5 is that the other 98 or 99 abstract factors only contribute to noise in spectra. Thus, the technique used to calculate abstract factors from a set of spectra displayed in the active tile makes use of the field labelled *Factors* which in Figure Sucrose 14 is set to 12. The fact that the value for *Factors* is 12 means when the *Generate* button (Figure Sucrose 14) is pressed, at most 12 abstract factors are calculated by iterative SVD. If the tick-box labelled *Calc Specified Factors* (Figure Sucrose 14) is ticked then exactly 12 abstract factors are calculated. The time spent calculating abstract factors is reduced from that required for 101 to only 12. For a data set of size shown in Figure Sucrose 14, the saving is less important than if the number of data bins per spectrum is larger and the number of spectra is large. When the problem exceeds a certain threshold in terms of number of data bins and number of spectra the saving in time computing a smaller number of abstract factors is considerable.

Results shown in Figures TEXT4 and TEXT5 represent approximations to spectra in Figure Sucrose 14 obtained by selecting a reduced number of abstract factors (in these two cases 3 and 2 out of the possible 101 abstract factors available for 101 spectra) and constructing linear least squares fits using the chosen reduced set of abstract factors to each spectrum. To create Figures TEXT4 and TEXT5 the spectra for which an approximation based on a specified number of abstract factors must be included in the set of 101 spectra overlaid in the active tile. The number of abstract factors used in the least-squares calculation is specified in the *Factors* field before pressing the button labelled

Predict using Factors (Figure Sucrose15). The current set of spectra displayed overlaid in the active tile are processed by iterative SVD yielding the specified number of abstract factors, after which the processed form of each spectrum is set equal to the approximation to the spectrum calculated using the specified number of abstract factors as shapes fitted to data in a linear least squares sense.

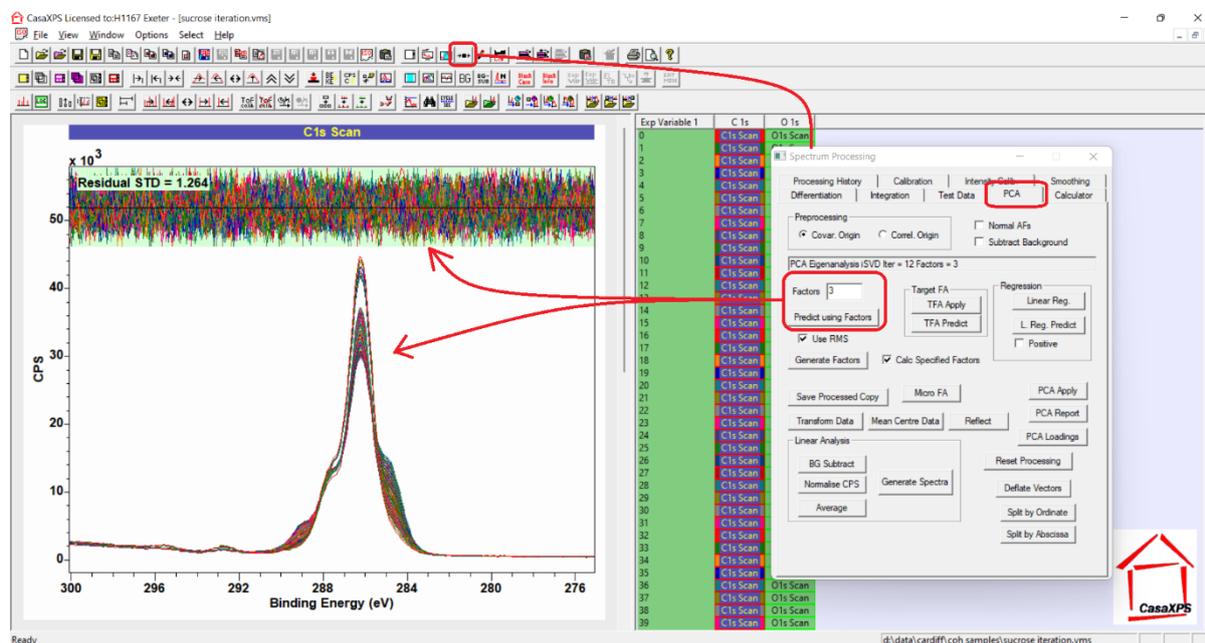


Figure Sucrose15. Display of spectra after processing of spectra using the *Predict using Factors* with the number of factors specified via the *Factors* field equal to 3.

A feature of these two Figures TEXT4 and TEXT5 is the residual plot and the figure-of-merit value displayed with these approximations. These measures for the goodness of fit are available as a side effect of how CasaXPS manages spectral background curves. A background curve is maintained for each VAMAS block. As a rule, background curves are calculated in response to parameters and actions associated with the *Regions* property page of the *Quantification Parameters* dialog window. However, there is a special case for defining a background curve which occurs if no regions are defined on spectra. For spectra without a means of calculating a background curve from regions, when a VAMAS block is reset to the raw data state, following some form of processing, the act of resetting processing for a VAMAS block sets the background curve for the VAMAS block to the raw spectral data. If following the resetting of processed data the VAMAS block is processed, such as calculating an approximation to the spectrum such as *Predict using Factors*, the difference between the processed data and the background curve yields the residual and figure-of-merit seen in these two Figures TEXT4 and TEXT5. Making use of this side-effect feature of background curves allows a comparison between processed data and raw data that is equivalent to the comparison between the total synthetic envelope calculated from a peak model and raw spectral data. The act of predicting spectra based on abstract factors and the residual plot and figure-of-merit is a very useful way of assessing the number of abstract factors applicable to reconstructing spectra for all spectra in a PCA calculation. Ideally one should review the results equivalent to Figures TEXT4 and TEXT5 for all spectra and monitor the figure-of-merit. By so doing it is possible to see if any spectra in the data show through anomalous figure-of-merit to be different from others in the data set. Spotting such anomalies can, for example, help identify outliers in a data set.

The principal object for making use of the actions illustrated in Figure Sucrose15 is to create approximations to spectra within a data set with reduced noise compared to the raw spectra. The

principal reason for creating a set of spectra with reduced noise is to allow the calculation of difference spectra. When two spectra are subtracted the intention is to create spectral shapes that represent systematic shapes in data which can be interpreted as physically meaningful. However, this operation of subtracting two spectra is to amplify noise in the resulting spectral shape. Hence, making use of PCA to eliminate to some degree noise in spectra is an advantage to the analysis that yields results shown in Figures TEXT6, TEXT7 and TEXT8.

The set of curves shown in Figure TEXT6 are calculated from two spectra selected from the set of processed spectra shown in Figure Sucrose15. For this particular case of sucrose data, the results of PCA suggest the entire set of spectra can be approximated by only two spectra. Choosing two spectra from the set of spectra in Figure Sucrose15 is best made by selecting two spectra for which the angle between the chosen spectra is the maximum angle between any two spectra in the data set. This selection of two spectra is relatively easy because the widest angle between spectra occurs between the first and the last in the measurement sequence (Figure Sucrose16). Once the selected spectra are overlaid in the active tile, pressing the button labelled *Difference Spectra* with the field with label *Increment* set to 0.1 results in a new VAMAS file created containing a list of spectral shapes that are computed with varying proportions of the two spectra displayed in the active tile. The experimental variable for the new VAMAS file is in the range 0 to 100, with values incremented by 0.1. These values are an indication of the different proportions for these initial spectra used to compute each difference spectrum in the new VAMAS file. These difference spectra and the VAMAS file with VAMAS blocks ordered by the experimental variable with increment 0.1 is shown in Figure Sucrose17. Scanning through the list of VAMAS blocks in Figure Sucrose17 using the up-arrow and down-arrow keyboard keys allows the visual inspection of these difference spectra and the identification of spectral shapes that are potentially physically significant.

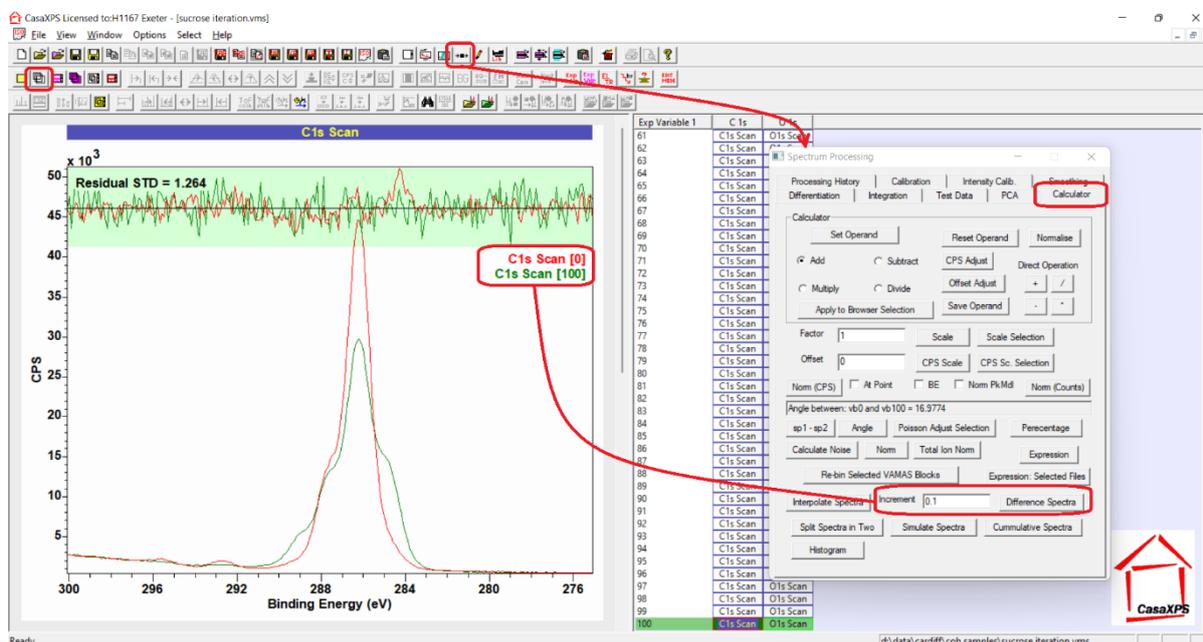


Figure Sucrose16. Two spectra overlaid in the active display tile in preparation for constructing a new VAMAS file containing difference spectra suitable for identifying two spectral shapes of physical significance to the XPS of sucrose. The results of actions invoked via the Difference Spectra button on the Calculator property page with the Increment parameter set to 0.1 are shown in Figure Sucrose17.

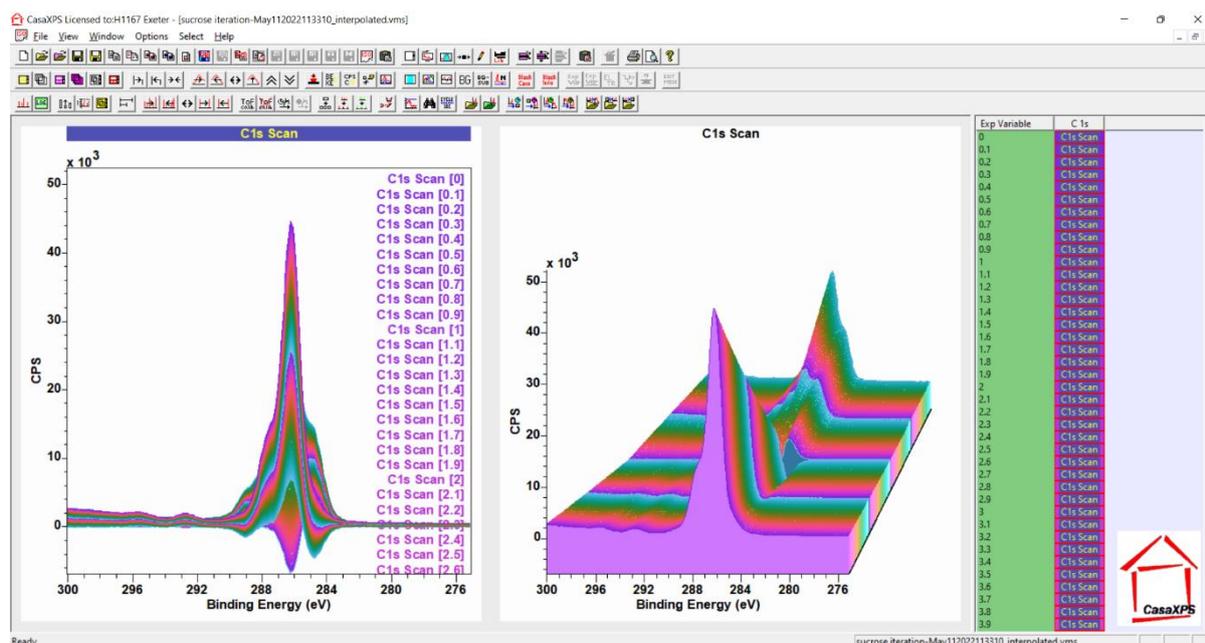


Figure Sucrose17. A new VAMAS file containing the results from actions illustrated in Figure Sucrose16, displayed overlaid via 2D and 3D options available on the Geometry property page of the Tile Display Parameters dialog window. The new VAMAS file contains 1001 difference spectra, which is the result of specifying 0.1 for the Increment parameter in Figure Sucrose16.

The list of difference spectra is visually inspected for likely candidates for spectral shapes that help in understanding evolution in the original C 1s spectra. Figure Sucrose18 illustrates how the *Save Processed Data* button can be used to extract VAMAS blocks from the list and add selected VAMAS blocks to a new column of VAMAS blocks in the right-hand pane. The advantage of the *Save Processed Data* button (Figure Sucrose18) on the PCA property page of the Spectrum Processing dialog window is that a VAMAS block is copied and the element/transition label for the new VAMAS block is adjusted to include a numerical value in the string that separates columns of VAMAS blocks seen in the right-hand pane. By using the *Save Processed Data* button mechanism, selected VAMAS blocks can be separated from the general list of difference spectra and therefore can be selected and copied to other VAMAS files without difficulty.

Once two spectral forms of interest are identified, copying these two VAMAS blocks to a file containing the original spectra (Figure Sucrose19) allows, by using these two spectral shapes in a linear least-squares fitting to the original spectra, the appropriateness of these two spectral shapes to be assessed. Since in the case of these sucrose data all C 1s spectra can be approximated using two abstract factors, if these two spectral shapes are selected correctly, then when fitted to data the result should be a good fit to all spectra. An example of data fitted with two spectral shapes is shown in Figure Sucrose20. By these means the results shown in Figure TEXT8 are created. Simply because two spectral shapes fit all spectra in the original data set, this does not mean that the most physically meaningful spectral forms have been identified. Rather, the fitting as described demonstrates that the mathematics required to fit data has been correctly applied when creating these spectral shapes. Sample knowledge and experience of XPS is more important in identifying the best pair of spectral shapes than mathematical solutions obtained by these means.

The steps illustrated in Figures Sucrose19 and Sucrose20 are used to create the results shown in Figure TEXT8.

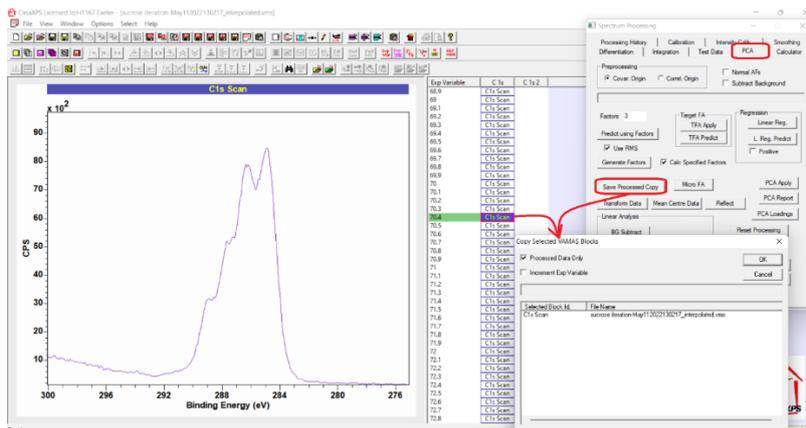


Figure Sucrose18. An example of a spectral form identified within the list of difference spectra as a possible candidate for physically meaningful information about the sample. The Save Processed Copy button is used to copy the selected VAMAS block to a new column as displayed in the right-hand pane.

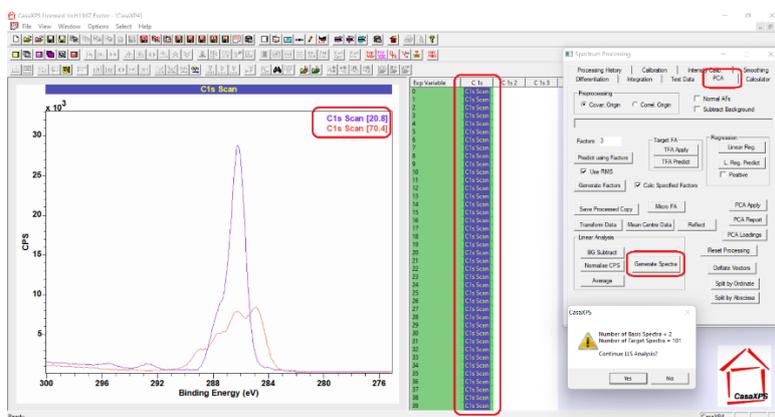


Figure Sucrose19. Example of a VAMAS file created by copying VAMAS blocks from the original set of C 1s spectra added to two VAMAS blocks created from difference spectra. Overlaying the two difference spectra in the active tile and selecting in the right-hand pane VAMAS blocks containing the original C 1s spectra allows the use of the Generate Spectra button to create linear least squares fits of the original C 1s spectra with respect to the two spectra overlaid in the active tile.

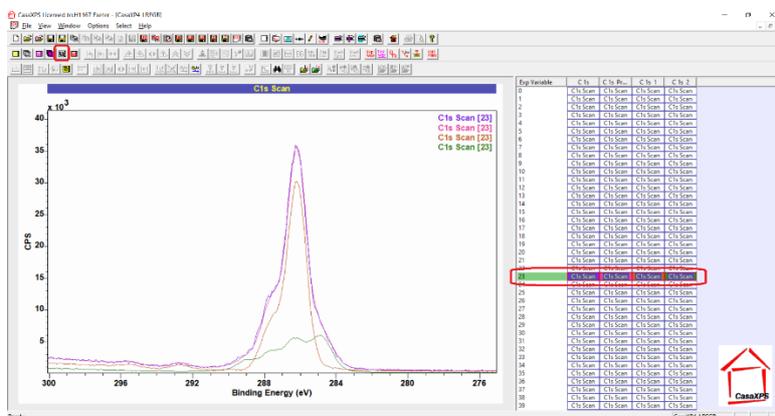


Figure Sucrose20. The result of the steps illustrated in Figure Sucrose19 where the spectrum measured after 23 repeat acquisition cycles is fitted by the two spectral forms calculated from difference spectra.

The two peak models proposed in Figure TEXT7 are used to construct the peak models in Figure TEXT9. For each component defined in Figure TEXT7, the same components are added to spectra in Figure TEXT9, where the two sets of components in Figure TEXT7 are scaled making use of the fits to data shown in Figure TEXT8 so that when added to spectra these components fit each spectrum. The fact that two spectral shapes fitted with subsets of components are the basis for the models in Figure TEXT9 is emphasised by filling each component with grey scale colours. These colours are specified using the palette of colours mechanism, where sixteen colours are associated with Comp index values, however the precise way a grey scale is made available when displaying components in a peak model is explained in the following.

Colours used to display components are individually specified using the methods illustrated in Figure Sucrose10 and Figure Sucrose11. However, it is also possible to populate the custom colours palette by making use of the colour scales available to the display of spectra (Figure Sucrose8). Component colours can be defined by first selecting one of the colour scales available on the Colours property page of the Tile Display Parameters dialog window, then selecting the button labelled *Comp = Scale* (Figure Sucrose21).

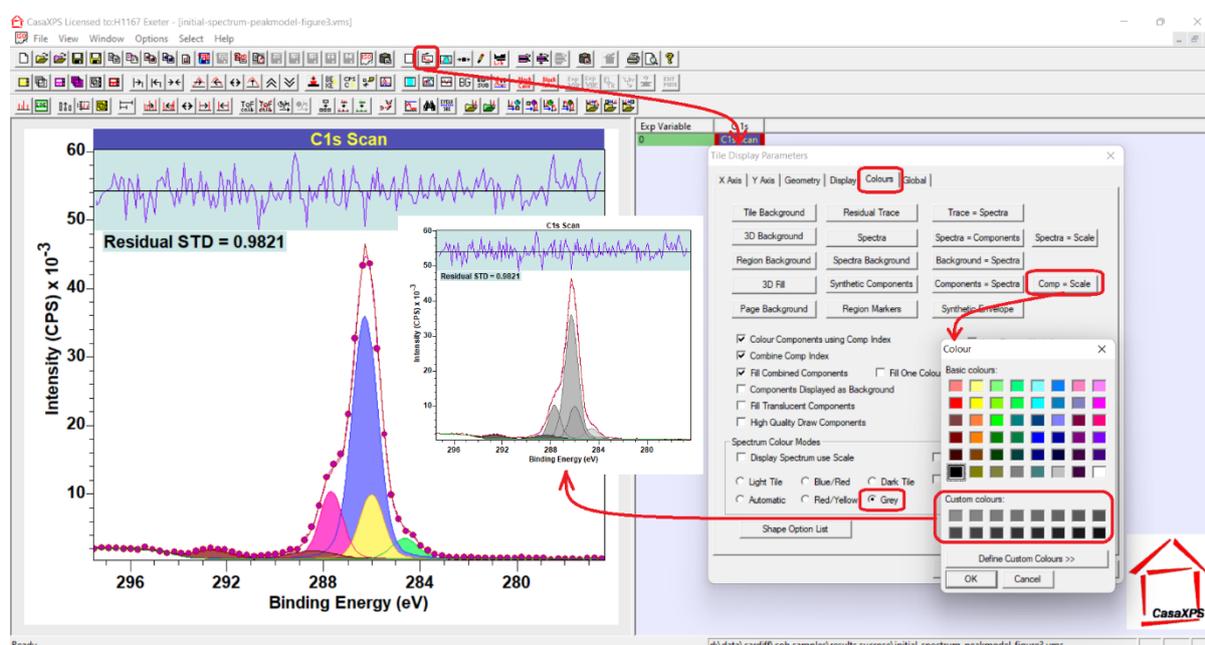


Figure Sucrose21. Colours property page indicating the use of *Comp = Scale* button.

Peak Models and Constraints

Technical Concepts useful to Fit Curves to Data

The reason a peak model is constructed for XPS data is to measure quantities which suggest scientifically meaningful information about a sample. The limiting factors in achieving this aim are the choices made for components used in a peak model and the ability of mathematical optimisation to translate these choices into output that make scientific sense. In the following list of dos and don'ts the objective is to highlight aspects of peak models which make a difference to outcomes. However, to appreciate the need for these points it is useful to see examples of where peak fitting is successful and, equally as important, to see where peak fitting is unsuccessful. Rather than providing examples of XPS data as measured, since there is always the possibility real data contains artefacts that introduce doubt into the mathematical basis for a discussion, data used in these examples is constructed from a single XPS peak measured from Nylon-6 which is modified through shifts in

binding energy only. Various combinations of these shifted N 1s data are scaled and combined to provide examples of data envelopes for which the mathematical definition can be considered the precise physical definition for spectra. To be specific, the outcomes determined by peak models and optimisation are not in doubt unless doubt is introduced by the peak model or the ability of optimisation to return the correct result.

The importance of defining a peak model in an informed way is illustrated making use of data constructed from three N 1s spectra all identical with the exception that each N 1s peak maximum is offset in energy from each other. These N 1s spectra are plotted beneath peak models shown in Figure SPM1, Figure SPM2 and Figure SPM3. Spectra are scaled in the proportion 2:5:1 as shown in the header text to each plot in each of these three figures then summed to form the overall data envelope against which peak models are fitted.

The problem for optimisation posed by data of this nature is the correlation of signal from underlying spectra where each of these three basis spectra can contribute to signal used in the construction of the data envelope over a range of data channels. The reverse process achieved by fitting curves to data involves separating signal into three components. However, without guidance from constraints to optimisation parameters and the constraint imposed by line shapes, the result is uncertainty for binding energy and relative intensity when measured by components. Three examples of peak models illustrate possible types of errors due to choices made for peak models. Figure SPM1 is an example of a peak model in which all optimisation parameters (area, position and FWHM) are determined by optimisation of the chi square statistic. The initial example involves line shapes (Figure SPM1) arbitrarily chosen to be a pseudo-Voigt function GL(40). The intention in using GL(40) is to compare how optimisation performs given a peak model without constraints (Figure SPM1). A peak model where line shapes are very close to the true line shapes for data (Figure SPM2) which involves optimisation through adjusting area and binding energy without the use of FWHM is designed to demonstrate how line shapes close in shape to the true line shape alters the outcome for optimisation. Finally, the peak model in Figure SPM1 is constrained to force all FWHM to be identical which results in the outcomes illustrated in Figure SPM3. These three examples demonstrate some general trends when fitting data by means of specific situations, namely, that unconstrained fitting of components based on synthetic line shapes GL(40) (Figure SPM1) fails to return the mathematically certain composition of the data envelope, while appropriate use of constraints yields improvements in outcomes. The best outcome is achieved for precise line shapes calculated from background subtracted data (Figure SPM2). The reason Figure SPM2 yields the best outcome is only area and energy are calculated by optimisation coupled with the line shape is based on data and therefore as close as possible to the true line shape. Since the line shape is fixed in shape, this implies the FWHM cannot be adjusted during optimisation, therefore the solution is guided to a minimum chi square by rigid and correct shapes. Figure SPM3 is the next best outcome, which is achieved by forcing the FWHM for all three components to be identical. Identical FWHM limit the way GL(40) synthetic shapes can broaden or narrow. While clearly advantageous in this example, the reason this is so powerful is it is applied to a data envelope constructed from peaks that are indeed all the same width. Photoemission spectra include many examples of peaks from the same element and electron configuration which broaden or narrow depending on chemical state. The use of constraints such as those shown in Figure SPM3 must be supported by evidence rather than applied as a generally applicable rule.

A point worth noting about the best case for outcomes shown in Figure SPM2 is that the fit of these three components is not exact in a numerical sense. Few would object to the statement that Figure SPM2 implies the ratio for the underlying components is 2:5:1, but the minor numerical differences

from these integer values are due to inaccuracy in background curves that are integral to any peak model. Also, non-linear optimisation searches for the best fit for a peak model rather than calculating the exact fit mathematically, which is the case for linear least squares fitting. When presented with three spectral shapes that were used to construct the data envelope in Figures SPM1, SPM2 and SPM3, linear least squares optimisation calculates the precise solution, limited only by numerical machine precision errors and algorithm used to perform matrix inversion (Figure SPM4). The success achieved in Figure SPM2 and Figure SPM4 is obtained by making use of curves that are accurate representations of the underlying curves used to construct the data envelope. The optimum success for curve fitting shown in Figure SPM4 demonstrates that least squares can work extremely well provided the exact line shapes and background curves are known. In the case of Figure SPM4 background curves and bell-shaped curves are integral to the spectral components. While successful, Figure SPM2 results are compromised slightly by the background approximations involved in the peak model, which is arguably best seen in the shape of the residual in Figure SPM2. A close inspection of the residual plot indicates high quality fit at the peak maxima but poorer quality at the points where the background in the peak model meets data.

While these examples help to appreciate certain aspects of fitting curves to data, the list of dos and don'ts are designed to expand on these ideas and provide a basis for anyone experimenting with XPS data.

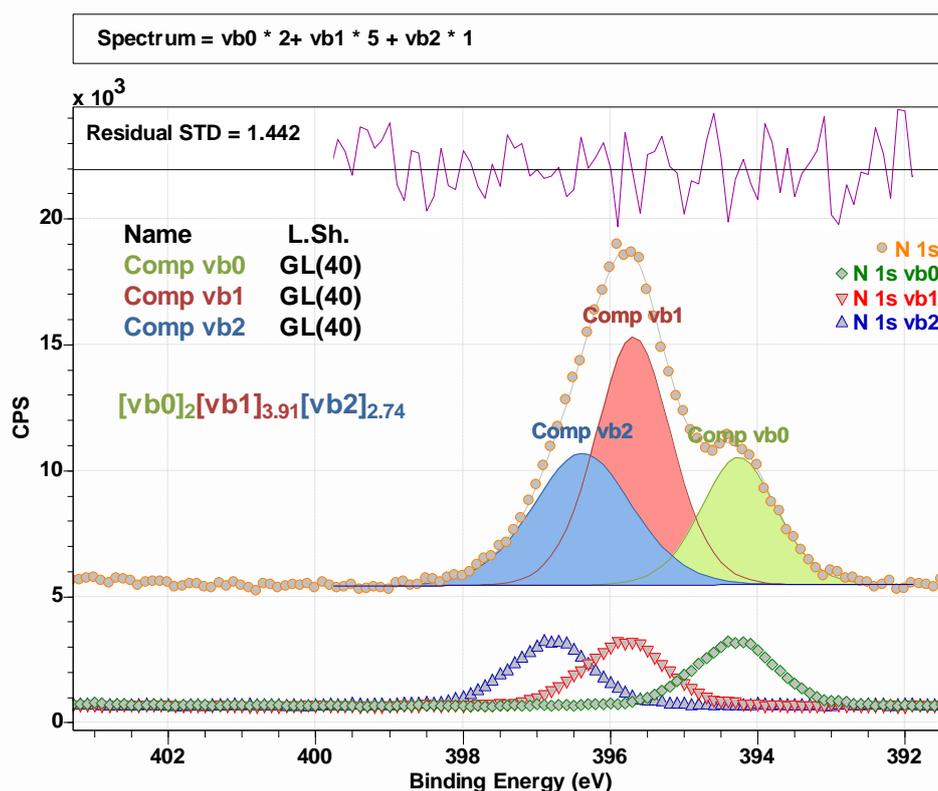


Figure SPM1. Constructed spectrum fitted with three components without the use of any relational constraints or interval constraints that alter the outcome for optimisation. The constructed spectrum is designed to have three peaks of identical shape, off set in energy and summed in the proportions 2:5:1. The outcome of optimising area, position and FWHM fitting parameters does not return the expected area ratio for these three components.

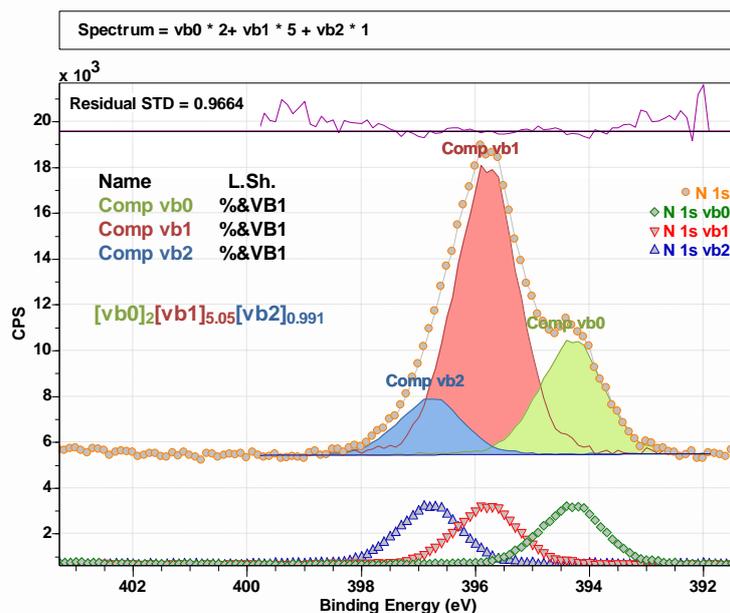


Figure SPM2. Constructed spectrum fitted with three components without the use of any relational constraints. The line shape used is calculated from the spectrum N 1s vb1 by subtracting a linear background and interval constraints are applied to FWHM to prevent changes to the shape for each component. The constructed spectrum is identical to the data in Figure SPM1. The outcome of optimising area and position with fixed FWHM fitting parameters returns the expected area ratio for these three components. Loss of precision in this case is due to inaccuracy in the line shape due to the selected background and energy interval over which the background is computed for both fitting to data and the calculation of the line shape %&VB1.

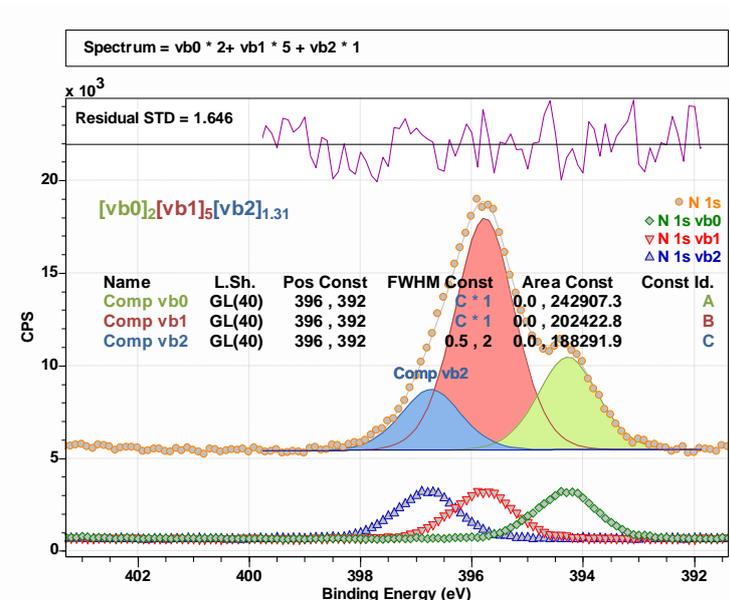


Figure SPM3. Constructed spectrum fitted with three components without the use of any relational constraints for area and position, but where relational constraints ensure all components when optimised have identical FWHM. The constructed spectrum is designed to have three peaks of

identical shape, off set in energy and summed in the proportions 2:5:1. The outcome of optimising area, position and FWHM fitting parameters does not return the expected area ratio for these three components, but making use of relational constraints for FWHM the outcome is superior to the result in Figure SPM1.

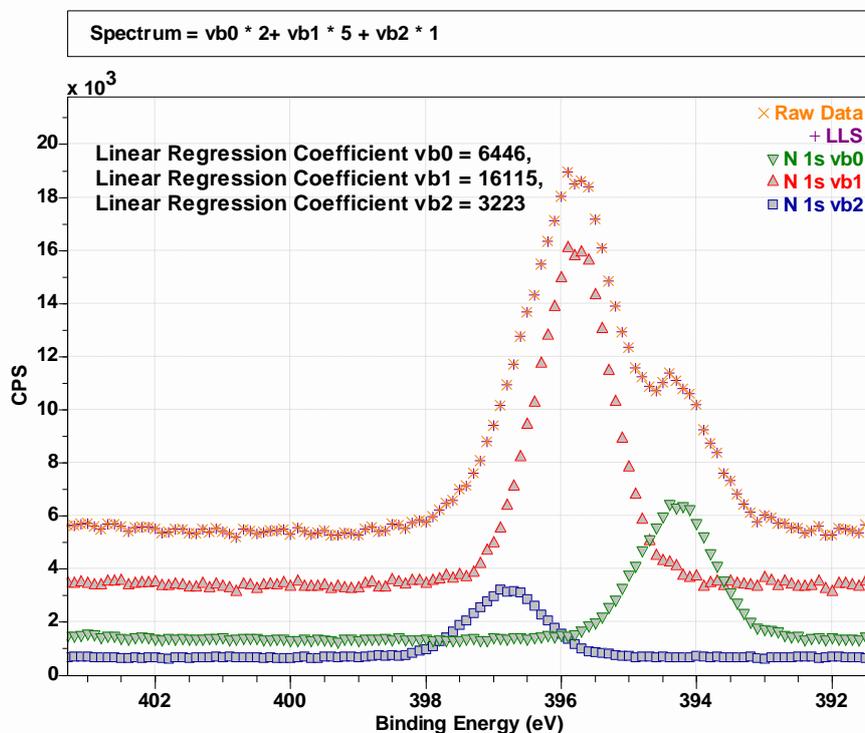


Figure SPM4. Linear least squares fit to data of three component spectra corresponding to the three spectra shown at the bottom of Figures SPM1, SPM2 and SPM3. The linear regression coefficients are the scale factors applied to each spectrum to achieve the quality of fit implied by the overlay for the symbols corresponding to raw data and LLS curves. Since these component spectra used in the linear least squares fit, the coefficients as computed by linear least squares are exactly in the proportions 2:5:1.

Details that should be Considered when Fitting Curves to Data

When attempting to fit data with a set of curves it is important to have more with which to work than simply spectra. Specifically, one needs to know the characteristics of the instrument used to collect spectra, should have an understanding of the sample and have an understanding of the physics involved in measuring the sample by XPS. While full knowledge for all three of these points suggests no analysis need be performed because all aspects of a measurement and sample are already perfectly clear, for practical XPS the objective in performing a measurement is to expand upon initial knowledge resulting in an improvement in all three of these prerequisites for successful XPS. Planning and performing experiments that further an understanding for any or all three prerequisites are recommended, if possible.

Design experiments that in so doing create enough data to enhances an understanding for the points described above.

Experimental design is a somewhat broad subject for which specific advice generally applicable to samples of a disparate nature is difficult to address. However, some basic thoughts may aid the process.

Energy resolution is a measure of how precisely an electron of a given energy can be measured. Low energy resolution implies electrons with a relatively wide range of energies for electrons are all allocated to the same data bin within a spectrum. High energy resolution implies the range of energies for electrons assigned to the same data bin in a spectrum is relatively narrow.

All instruments offer several operating modes which alter the energy resolution, meaning different operating modes change the ability to allocate an energy data bin into which an electron, when recorded, is placed.

Determine or request information about the energy resolution used to collect spectra.

Determining the energy resolution for an instrument is often an advantage when analysing data and the energy resolution for the instrument should be reported in the experimental section for XPS data included in a publication.

Energy resolution for an instrument is a consequence of multiple factors relating to instrument design. One should be aware that for some instruments, the energy resolution may not be constant over the entire energy interval used to measure spectra. Notwithstanding this observation, one means of measuring the energy resolution of an instrument is to make use of valence band data that includes the Fermi edge from a clean silver sample. The resolution of the Fermi step for clean silver is one measure for the energy resolution of an instrument. Another measure is to use the FWHM measured from Ag 3d_{5/2} to monitor the energy resolution for photoemission typical of materials with low atomic number. The FWHM measured by these two means differs because in the case of Ag 3d_{5/2} the photoemission line width is determined by the physics of photoemission which is inherently broader than the processes that determine the sharpness of the Fermi edge drop in signal.

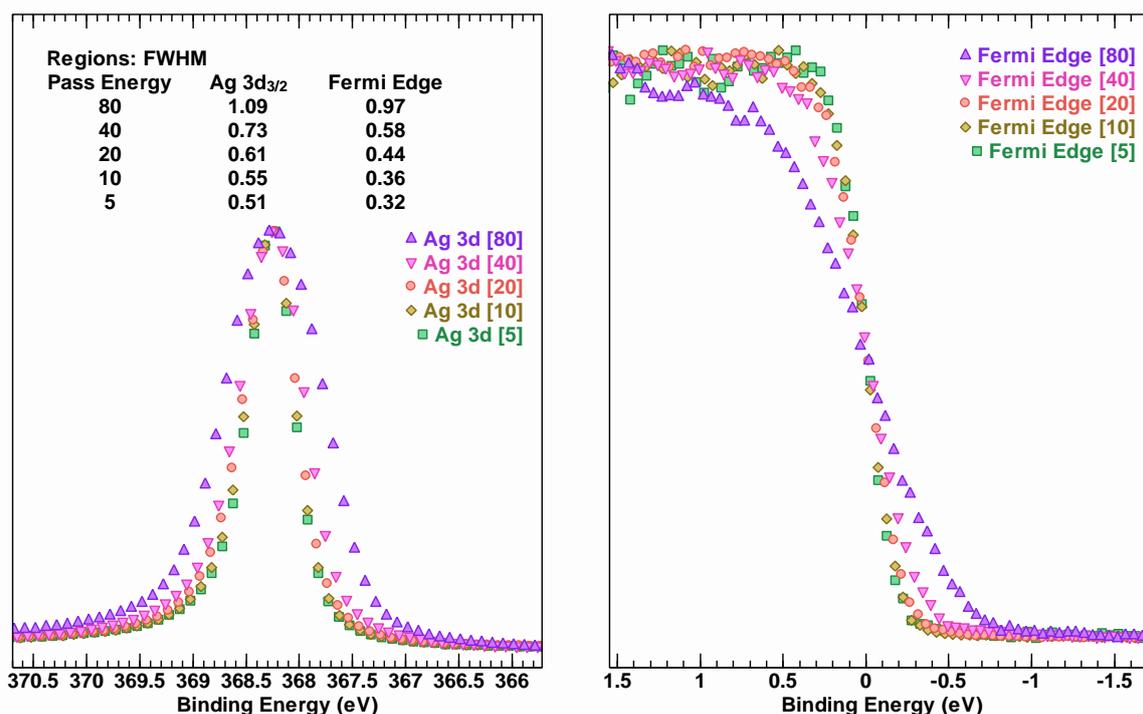


Figure ER1: Comparison between photoemission and Fermi edge FWHM for an instrument operating with different pass energies. Both sets of data demonstrate the general trend in energy resolution for these different operating modes. Note that Fermi edge data yield narrower FWHM than Ag 3d_{5/2} because of the difference in natural width for Ag 3d_{5/2} and the instrument independent drop in signal at the Fermi edge for silver.

Energy resolution is important for samples with narrow photoemission peaks, but not all samples yield narrow peaks. Deep core level electrons or photoemission of electrons from elements in higher oxidation states or from heavier elements tend to produce broader FWHM, in which case monitoring energy resolution is less advantageous. However, one of the most common photoemission lines to which peak models are applied is C 1s. Peak models for C 1s are often open to criticism for including components within a peak model that do not respect the physical processes which determine FWHM for photoemission as measured. Regardless of how narrow a photoemission line is at point of source, an instrument energy resolution places a lower limit on the width for measured peaks and so no component peak in a peak model can meaningfully be fitted to data with FWHM less than the energy resolution of the instrument used to measure that peak. By contrast, it is often the case that sample and chemical state is more influential on FWHM for photoemission peaks than instrumental resolution. Figure ER2 is an example of a peak model applied to C 1s data collected from Nylon-6 sample using two pass energies, namely, pass energy 20 eV and pass energy 40 eV. For all four components used to model chemical state for Nylon-6 the FWHM exceeds the FWHM limit for these two pass energies for this particular instrument. Indeed, these two peak models in Figure ER2 are identical and when fitted to these two C 1s spectra there is little difference in terms of outcomes for FWHM despite the pass energy difference involved. Changing the pass energy from PE 40 to PE 20 for Ag 3d_{5/2} as shown in Figure ER1 does affect FWHM. The FWHM and lack of change with pass energy for components in Figure ER2 suggests factors other than instrumental broadening of photoemission lines are responsible for the widths for components seen in Figure ER2.

An important point to observe about the peak model applied to data in Figure ER2 is not all FWHM are necessarily identical. While developing a peak model for data similar to these C 1s spectra, it may initially be useful to link FWHM in terms of relational constraints that force all components to attain identical FWHM values, the ultimate peak model may contain FWHM that are different. The choice for constraints involving FWHM have implications for line shapes used in components. Note that C=O C 1s signal in Figure ER2 are modelled with different line shape to line shapes used for other components. Line shape parameters alter fitting parameters such as area and FWHM when peak models are fitted to data.

Understanding the energy resolution helps to spot potential issues with components to a peak model. The results tabulated for FWHM measured from clean silver in Figure ER1 provide evidence that a peak model in Figure ER3 with three components includes a component that is mostly likely fitting noise rather than signal. The signal to noise in these data in Figure ER3 is too low to support a peak model design with either three or two components, but the mere fact that the FWHM for one component is comparable in width to the FWHM achieved for pass energy 20 eV for silver Fermi edge data suggests Al 2s photoemission is unlikely to be responsible for the feature fitted and labelled Noise in Figure ER3.

The conclusion from energy resolution considerations is that FWHM for components to a peak model are important indicators for rejecting incorrect peak models regardless of how well the model may fit data. Thus, do make use of FWHM results to assess the validity of a peak model.

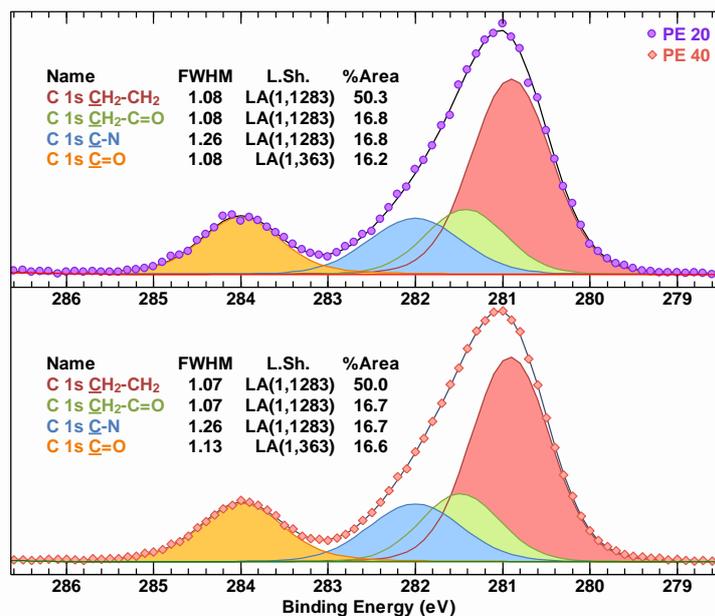


Figure ER2: Nylon-6 C 1s spectra modelled using identical peak models which fit pass energy 20 eV spectra with similar results as when fitted to pass energy 40 eV measured sequential from the same sample. These data when fitted with the peak model both indicate that the same FWHM for three components are close to equal, while C-N C 1s signal has a FWHM that differs from these other three.

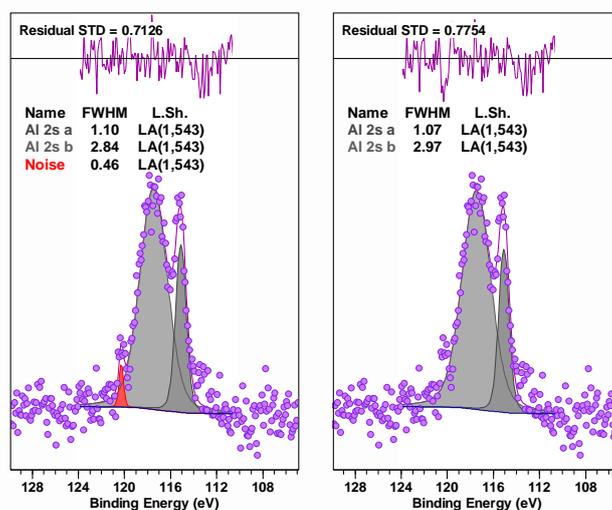


Figure ER3: These data are measured from a sample with an instrument with characteristic for energy resolution similar to the one used to collect data in Figure ER1. The pass energy for these Al 2s data is 20 eV, therefore attempting to fit a component with FWHM 0.46 eV is doubtful given that the Fermi edge width measured using pass energy 20 eV for silver is 0.44 eV. The most likely explanation for the component labelled Noise is some systematic noise has contrived to create data that is not of chemical origin. By contrast the other two components are almost certainly associated with chemistry within the sample measured.

Consider stoichiometry when assessing sample chemistry, but do not place too much credence on atomic concentration results from XPS.

XPS is a surface sensitive technique. Signal as measured from a given atom depends on position within the volume of material at the surface from which photoemission without energy loss is recorded. For heterogeneous samples quantification results presented as atomic concentration often yield values at odds with the true chemistry of a sample. A well-known example which demonstrates anomalies in chemical composition measured by XPS is measuring a film of germanium oxide on elemental oxide. If low kinetic energy photoemission (Ge 2p) from germanium is used to measure the proportions of oxide to elemental forms the ratio differs markedly from if high kinetic energy (Ge 3d) photoemission is used. Kinetic energy dependence of photoemission signal as measured by XPS similarly influences photoemission from different elements with different kinetic energies, therefore a basic application of standard quantification procedures for most samples should be considered with suspicion, even when the results appear to make sense.

Homogeneous materials are less prone to quantification anomalies because the unit volume sampled by XPS becomes the dominant influence on quantification results. Variations in volume analysed due to collecting signal at different kinetic energies are normalised through the use of transmission correction and escape depth correction, hence for homogeneous samples quantification by XPS is more plausible.

If samples analysed are known to be homogeneous then comparison of intensities from components in a peak model provide support for chemical state assignments. This aspect of peak models is illustrated in Figure ER4 making use of C 1s and O 1s data from a sample presented as poly (ether ether ketone) (PEEK). Chemical state for carbon and oxygen in PEEK provides a good example of how XPS reveals consistent results from both C 1s and O 1s. Chemical state for carbon within PEEK includes ring arrangements for these atoms where these rings are linked by oxygen atoms and carbon which is double bonded to an oxygen atom. The repeat units for PEEK imply the unit volume measured will include nineteen carbon atoms and three oxygen atoms. Carbon appears bonded in four different environments assigned \underline{C} -C, \underline{C} -C=O, \underline{C} -O and \underline{C} =O. Oxygen appears bonded in two different contexts, namely, \underline{O} -C and \underline{O} =C. The number of times these chemical bonds appear in PEEK implies chemically shifted carbon should manifest in XPS signal with four components available to a peak model which are in intensity ratio 12:2:4:1 corresponding to \underline{C} -C, \underline{C} -C=O, \underline{C} -O and \underline{C} =O, respectively. Similarly, for oxygen two components with intensity in the ratio 2:1 corresponding to \underline{O} -C and \underline{O} =C, respectively. Thus, when constructing a peak model for PEEK, bell-shaped components featuring in models for both C 1s and O 1s spectra must (according to theory) appear in the ratio 4:1 in a model for C 1s and 2:1 in a model for O 1s.

The peak model for O 1s in Figure ER4 represents strong evidence that the material generating these data contained two chemical states for oxygen. In particular these two O 1s component peaks are in the correct 2:1 proportion for PEEK and the relative difference in energy for these two components indicated in Figure ER4 is compatible with C-O and C=O type chemistry. A satellite structure, which is expected for polymers with carbon rings, is all that is required for the peak model to return the ratio 2:1. Specifically, no optimisation constraints in terms of relational constraints between component area parameters are required to achieve the O 1s result in Figure ER4. By contrast, the peak model for C 1s in Figure ER4, while goodness of fit for the C 1s model suggests the solution is a feasible solution for C 1s signal measured from PEEK, the same level of confidence in the result as O 1s is not achieved. The reason the peak model represents a feasible solution for C 1s signal from PEEK is, to achieve the relationships between components in the peak model given the degree of correlation for C 1s signal, optimisation parameter relational constraints were necessary to enforce the

characteristic for components for C 1s as seen in Figure ER4. When models are fitted to data, different choices for optimisation parameter-constraints yield different outcomes and each outcome represents a different solution which must be tested against sample knowledge to separate mathematically correct solutions from physically feasible solutions.

Use constraints when constructing peak models for correlated signal.

Parameter constraints are required for two reasons. If the true line shapes and shapes for background signal are not well understood, parameter constraints are used to impose relationships between component peaks to compensate for these deficiencies and allow optimization to find an adequate reproduction of data with physical meaning. The second situation where parameter constraints are required, even if the true background and true line shapes are known, is the case in which a peak model is used to analyze highly correlated signal from different photo-emission sources. Optimization algorithms guided by a single figure of merit applied to highly correlated signal struggle to return meaningful data reproduction without constraints to enforce a necessary condition for a particular analysis as shown in the case of C 1s in Figure ER4. Expecting optimisation to distinguish between signal assigned to $\underline{C-C=O}$ and C-C as illustrating in Figure ER4 without constraints is unrealistic.

The use of parameter constraints represents user bias and therefore an overuse of constraints limits the independence of information extracted from a peak model such as shown for C 1s in Figure ER4. The intensity ratios are achieved by imposition of area relational constraints based on the theoretical chemistry for carbon in PEEK. The constraints used for the C 1s peak model in Figure ER4 are such that the only conclusion drawn from the peak model is that these data are consistent with the hypothesis that the sample is PEEK.

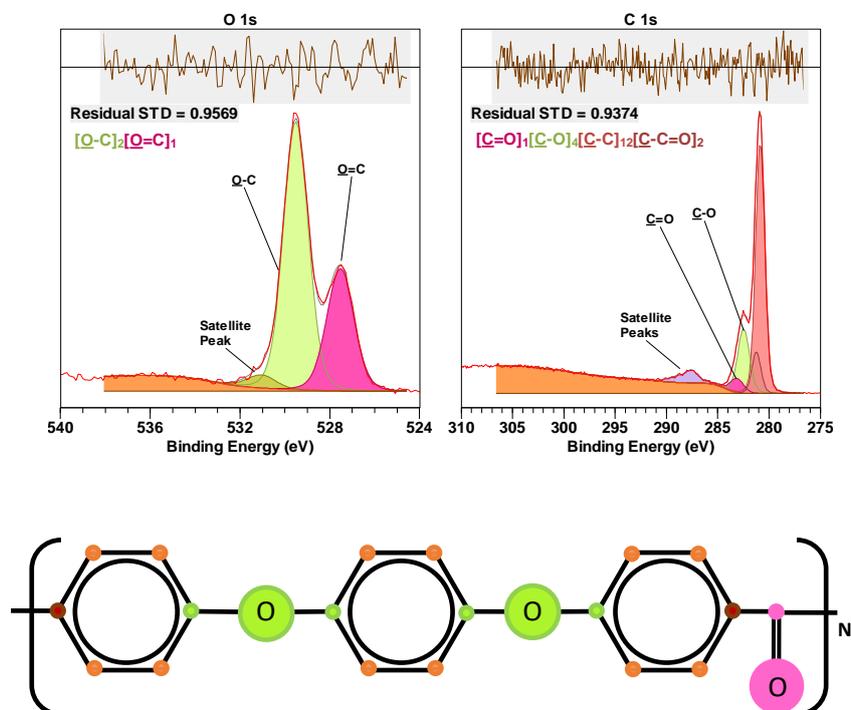


Figure ER4: O 1s and C 1s spectra fitted with peak models that support an assignment for components in these peak models consistent with the expected properties of PEEK.

CasaXPS Review of Constraints to Peak Models Analysis

A feature of the figures used to examine the role of constraints in peak models is the use of colours and symbols to plot and emphasise differences in data when overlaid in a display tile. By default, both colours and symbols are allocated to the display of data according to the order with which VAMAS blocks are selected in the right-hand pane prior to invoking the overlay toolbar button. There are options on the Colours property page (Figure CPM1) that alters the default behaviour for allocating colours to data displayed in tiles, but for the purposes of the following discussion it will be assumed that the default behaviour is active.

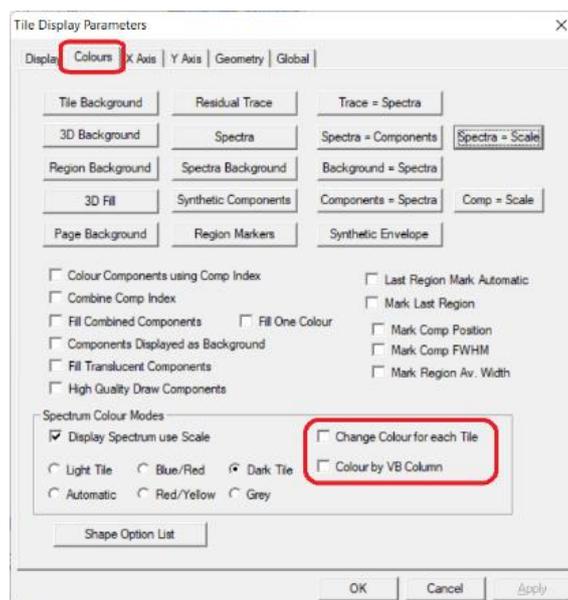


Figure CPM1. Options on the Colours property page which alter the default behaviour for assignment of colours and symbols to spectra displayed in the display tiles.

The display of overlaid spectra using symbols is illustrated in Figure CPM2. Colours for data are defined using *Display Spectrum using Scale* where the selected colour scale is labelled *Dark Tile*. The selected colour scale is designed for display tiles for which the tile background colour is assigned a dark background colour. These Ag 3d and Fermi edge data are displayed using symbols. The boundary to each symbol is drawn using the assigned colour. When *Display Spectrum using Scale* is active, the interior to each symbol that is defined in this example (symbols that appear on the *Editor for Points Draw Shapes* dialog window coloured green (Figure CPM2)) is a colour that is derived from the exterior colour for each shape. If *Display Spectrum using Scale* is unticked, then symbol exterior colours are assigned to data based on the Custom Colours palette and interior colours are assigned a colour based on the Custom Colours palette defined by the *Synthetic Envelope* colours. The example in Figure CPM2 also illustrates how a Custom Colours palette can be defined for spectra using the *Spectra = Scale* option which fills the set of custom colours using colours derived from the colour scale labelled *Dark Tile*.

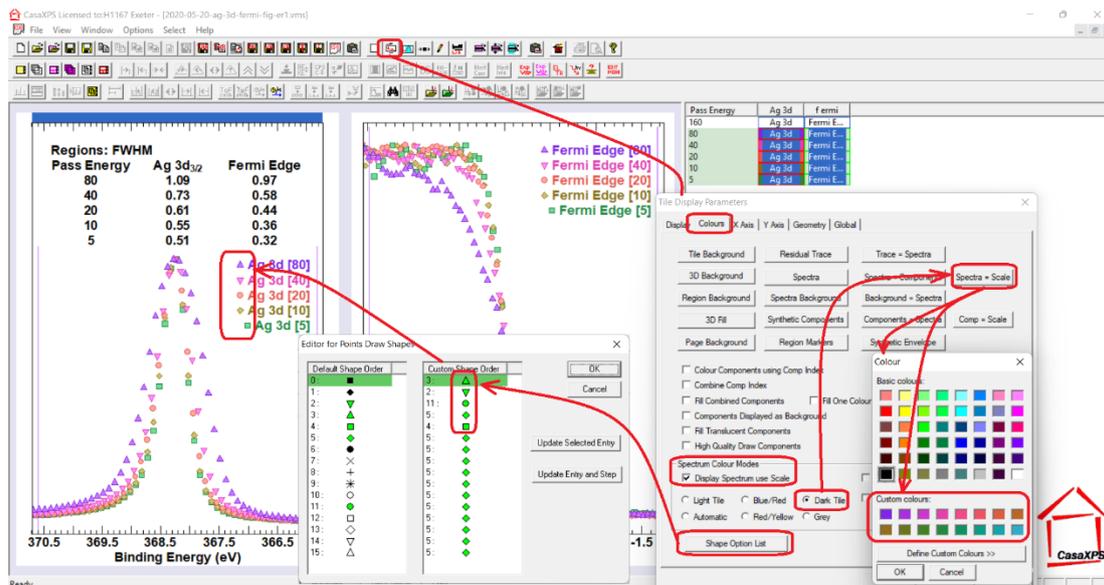


Figure CPM2. Colours and Symbols are defined via options on the Colours property page of the Tile Display Parameters dialog window.

An annotation option new to CasaXPS 2.3.25 is the introduction of the tick box labelled *Word Wrap Text* (Figure CPM3). When text is added to a VAMAS block using *Word Wrap Text* together with the *Draw Line* tick-box ticked and then the Annotation History property page is made the top-most property page on the Annotation dialog window, three move boxes are made available for adjusting the display of text. Two markers indicate the interval in the display over which text is drawn. If text is wider than the interval so defined then additional lines are used to display said text. The third marker allows the positioning of a line drawn from the centre of the text as drawn to the point in the display defined by the third marker (Figure CPM3).

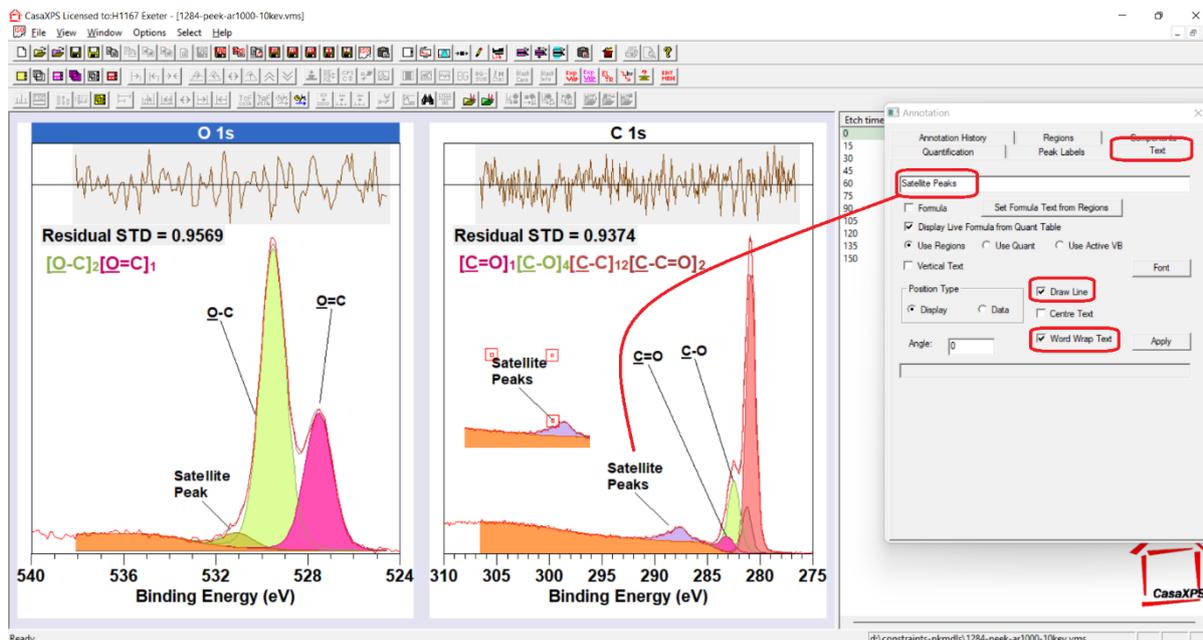


Figure CPM3. Text property page of the Annotation dialog window illustrating the options used to create and position text with a line where the string “Satellite Peaks” is drawn over two rows of text.