LICSS 3.0 Help

Contents

Overview

LICSS is a lightweight chemical spreadsheet for MS Excel for Windows easily capable of handling sheets with hundreds of thousands of molecules

You can enable Structure Display for a Workbook with Worksheets and Charts (embedded or as a separate Chart sheet). You need to be connected to the network/internet only if you don't already have the necessary JChemPaint visualisation and CDK files (ie: LICSS has not been installed). Worksheets are marked as enabled by setting the sheet name “Smiles” (by selecting the entire column containing the Smiles range and selecting <Ctrl Shift M>). Associated charts are enabled automatically.

When enabled for structure display, worksheets will display structure when a cell containing a Smiles string is selected or, when a chart is active, when hovering over chart points. Selecting <Ctrl R> (or running the macro ShowRowStructure) will also display the structure for the Smiles string on the selected row (or the structure for selected Smiles String if the Smiles name is not set, even on a *non-enabled* worksheet). Selecting <Ctrl T> will toggle between showing all structures and none. If you want to use a column to provide a window title for the structure window, select <Ctrl Shift Z>.

To Return a Smiles String from the Structure Edit window, select toolbar button: "Save Contents and Return to Application"

Selecting the LICSS Programs worksheet tab will display a menu of LICSS programs which can be run for enabled sheets including substructure and similarity searching; Selecting <Ctrl Shift R> will also display this menu. When using the structure editor to draw query structures, the atom Xe will match any atom and the atom He will match any atom except H and C. Alternatively you can use a range of shortcut keys

New Excel formulas: GetCDKDescriptor will calculate any CDK Molecular Descriptor or Mol Formula or Mol Weight, IsSubStructure and IsSimStructure will determine whether one Smiles string is a substructure of/is similar to another. ReformatSmiles will reformat a Smiles string and GetSmilesFromStructure pops up the JChemPaint Editor and returns a Smiles string for the structure drawn

Enabled worksheets (known as LICSS sheets) are fully independent of the creator Workbook or any AddIns

Return to Contents

New Features for Version 3.0

1. Version 3.0 of LICSS has been recoded to work as an Excel AddIn (rather than require code to be copied to all enabled workbooks). This has certain advantages:

   - Workbooks may be saved in the XLSX format – ie: they don't need to be macro-enabled
   - There are no longer any security issues with the Trust settings – all that is required is to enable Macros
   - Workbooks don't need to be closed and reopened after enabling before working – they are activated automatically
   - Users don't need to individually choose worksheets/charts to enable – if the Smiles name on a worksheet is set with <Ctrl Shift M>, that worksheet and all its associated charts are automatically enabled
   - Occasional Excel crashes due to closing/reopening activated workbooks are circumvented

2. Multiple columns of Smiles are now supported when displaying all structures on worksheets
3. A new feature to enumerate libraries has been introduced
4. Export SD File has been improved to produce files compatible with a larger range of programs
5. All the batch processes working on individual structures (substructure searching etc) have been recoded to use the new direct Java call mechanism used for Smiles visualisation. This makes the routines rather faster particularly on large numbers of compounds

Return to Contents

Setting up Worksheets and Charts

Make sure the AddIn 'LICSSV3.0.xla is loaded in Excel (this should have been done for you if you ran the installation program: "LICSS-3.0-FullInstaller.exe"). You can do this manually using Excel-Options or Tools-AddIns depending on your version of Excel (or, for a 'one-off' session just open LICSSV3.0.xla). Select <Ctrl E>. If you have already used the LICSS 3.0 system you will be prompted immediately with the dialog box:

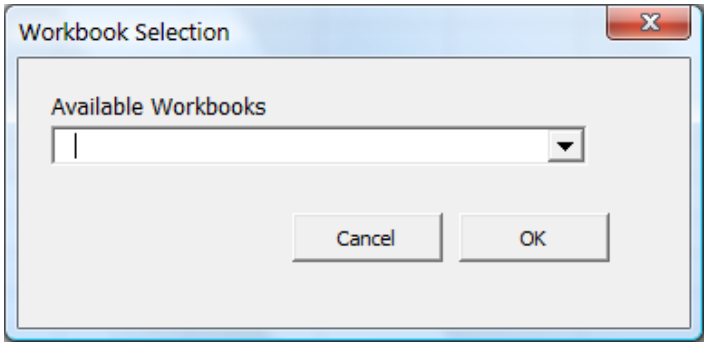

Alternatively, if you have a company LICSS setup, the necessary files will be downloaded automatically from your network. Finally, if yours is a 'stand-alone' installation, you will be prompted to download the necessary files from the internet. A 'yes' to this prompt will result in the program automatically downloading the necessary support files from the LICSS Google Project site (can be slow/fail depending on network speed/security settings). In any case, you will ultimately be presented with the dialog box above.

Select the workbook you want to enable in the 'Available Workbooks' dropdown. On pressing OK, the workbook is enabled and a success message is displayed. At this stage it is good practice to save the workbook. You can now optionally choose to enable particular worksheets by selecting the column containing Smiles strings (for that sheet) and pressing <Ctrl Shift M> to label the range as "Smiles"; you can also do this later if preferred. Any sheets with Smiles ranges labelled will now be activated for structure display. You may also want to set the column from which the structure window titles are created: select <Ctrl Shift Z> and follow the prompt. You can reset the various Smiles ranges as often as you like with <Ctrl Shift M>. This can be useful if you only want to operate on a subset of structures.

Charts associated with particular worksheets are enabled automatically; however you may need to re-run the ActivateStructureDisplay macro before the visualisation is activated: select <Ctrl Shift A>

You can update workbooks created with previous versions of LICSS by using <Ctrl E> as above. Save them in xlsx format to remove any previous code (no longer necessary).

The Smiles strings needed for LICSS can be obtained from many chemically-aware programs. LICSS has the ability to convert chemical names to Smiles and also has a formula GetSmilesFromStructure() which allows you to draw a structure and return the Smiles for it.

Return to Contents

Importing and Exporting SDF Files

SDF Import

SDF Files may be imported from an enabled worksheet, either into the worksheet itself or a new worksheet.

On selecting Import SDF File from the main LICSS Programs Menu, you will be prompted for an SDF File to import. After file selection, the SDF File is converted to Smiles format by the program then imported either into a new worksheet (default) or into the worksheet from which the program was run. In the latter case, you will also be prompted for a cell for the top-left of the new datablock.

When SDF Files are imported, the identifier becomes the first column, the Smiles data becomes the second column and all other data fields form the subsequent columns.

A common sequence is to chemically-enable a new worksheet then activate it followed by SDF File import into the worksheet. This is equivalent to creating a new LICSS spreadsheet from an SDF File. Before import of the SDF File you should name the 2nd column of the blank worksheet "Smiles" by selecting the entire column then <Ctrl-Shift-M>. This ensures that you can use the LICSS programs tab and that the new worksheet is ready for structure display after SDF File import (select cell A1 when prompted for a cell to import data into).

SDF Export

On selecting Export SDF File from the main LICSS Programs Menu, you will be prompted for an SDF File to export. After file selection, the program exports the data as an SDF file to the directory you have chosen. Data is chosen for export based on your range selection in the active sheet. If a single cell is chosen, the entire region around the active cell is used otherwise, the selection is used. In either case, the first row must contain column headings.

Return to Contents

Structure Display

Structures may be displayed for enabled worksheets by selecting the individual Smiles strings. A popup Java window shows the structure; this window may be resized and/or repositioned and will maintain the new settings in the current session.

For enabled charts, hovering over chart points will display the structure window as described above. Note that if the worksheet is filtered, for example by substructure, the chart will only display (by default) the filtered points and their corresponding structures if hovering over them.

For non-enabled worksheets (but in enabled work**books**), selecting <Ctrl R> will display the structure window for any **selected** Smiles string. For enabled sheets, <Ctrl R> will display the structure for the first Smiles string on that **row**.

<Ctrl T> will toggle all structure display on/off. In this case, the structures for all visible cells are displayed in each Smiles column (and Smiles strings hidden). Although it is possible to edit Excel whilst structures are displayed, there is a limit to what operations are consistent with successful all structure display and, in general, you should toggle all structure display off where possible. Thus, for example, you can filter a worksheet and then toggle all structure display on (and structures will display correctly for filtered data). However, if you try to filter **whilst** all structure display is on you will get some unpredictable results. If you navigate to a different sheet or attempt to save a sheet with all structure display on it will automatically be toggled off.

The structures displayed with all structure display toggled on are pictures, and may be copied to the clipboard and pasted into other applications before toggling all structure display off.

Return to Contents

LICSS Programs Menu

Is displayed when the LICSS Programs Tab is selected from a LICSS-enabled worksheet (or by selecting <Ctrl Shift R>)

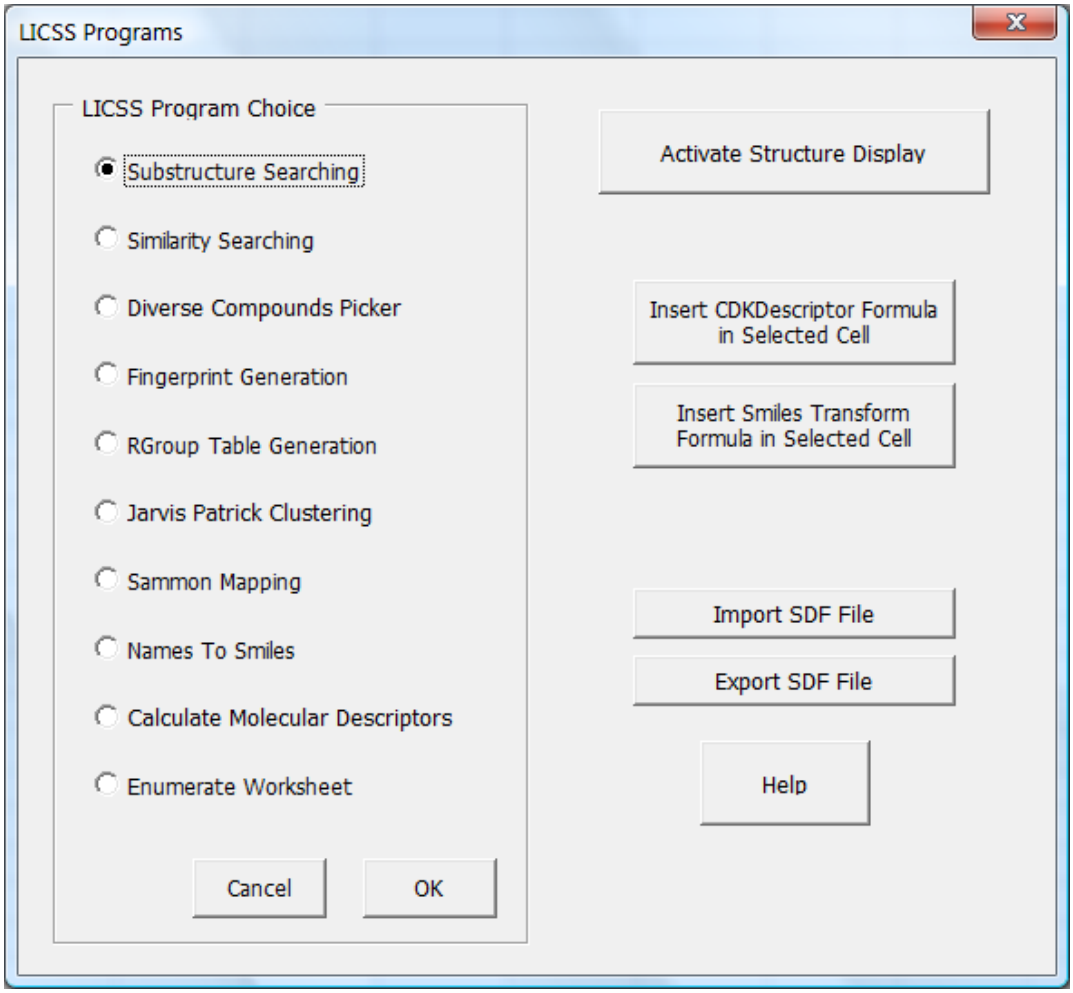

Select one of the options on the Left + OK to carry out a LICSS operation (most options operate on the 'Smiles' range except Name to Smiles [current selection] and Enumerate Worksheet [worksheet set up for enumeration]). Alternatively, Enter a CDK Descriptor or a Smiles Transform formula into the current selection, Read or Write SDF files, Activate Structure Display or display Help with one of the buttons on the Right

Return to Contents

Substructure Searching

On selecting Substructure searching from the main LICSS Programs Menu, if you have previously calculated fingerprints for the compounds on the sheet, you will first be asked whether you want to carry out a fingerprint search. Select 'Yes' for big sheets with > 5000 compounds. You will then be presented with the dialog box:

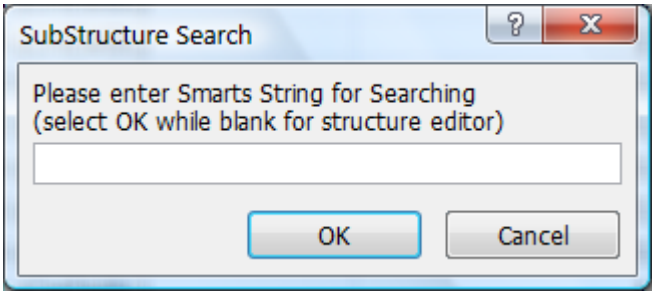

You can enter Smiles or Smarts at this prompt (except when you are carrying out a fingerprint search first). Normally, however, just select OK to be shown the main JChemPaint edit window:

Select when finished drawing

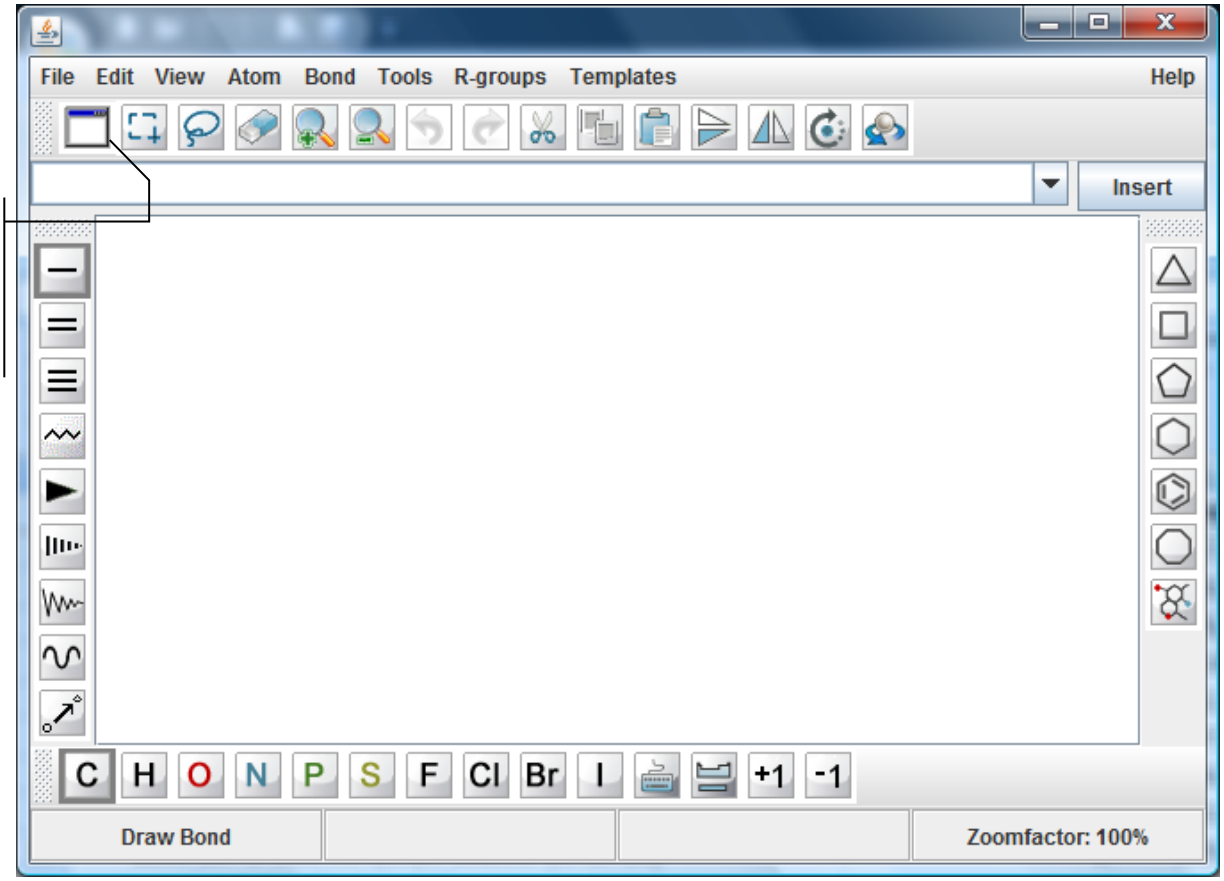

You can then use all the drawing tools as usual. The Xe atom will match any atom and the He atom will match any except Hydrogen and Carbon. If you want to carry out a partial or exact search, put in explicit bonds to Hydrogen as required. You can hover over atoms to be changed from Carbon and just type the atom required (eg N/O).

When drawing is finished select the top-left tool button indicated (tooltip: Save Contents and Return to Application).

A window displaying the search progress appears:

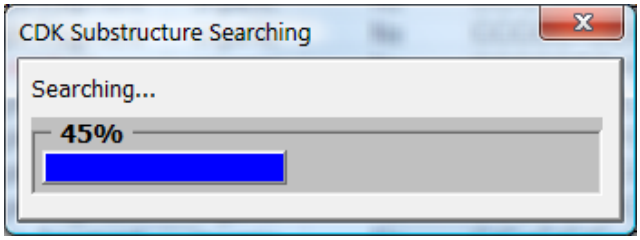

If you close this window the search is aborted. When complete, a new column is created for hits (1) or misses (0). A title cell is added: SSS: + the Smiles string used for searching. You can then filter on the search results using Excel's normal filtering facilities.

For sheets containing fingerprints, Substructure searching on several hundred thousand structures is eminently feasible in a short time.

Return to Contents

Similarity Searching

On selecting Similarity searching from the main LICSS Programs Menu you will be presented with the Similarity Search Settings dialog. Select the Tanimoto coefficient (between 0 and 1) you would like to use as a cut-off. CDK Fingerprints are such that you may wish to use a relatively low value between 0.5 and 0.7 to obtain many hits. On, Okaying this box, you will be presented with a Smiles entry window which, if OK'd when blank will pop-up the JChemPaint editor after which the Java program will run as with Substructure Searching. A new column is inserted to the right of your Smiles names with 0 (for miss), 1 (for hit) and a column title with the Smiles string searched for @ level of similarity chosen

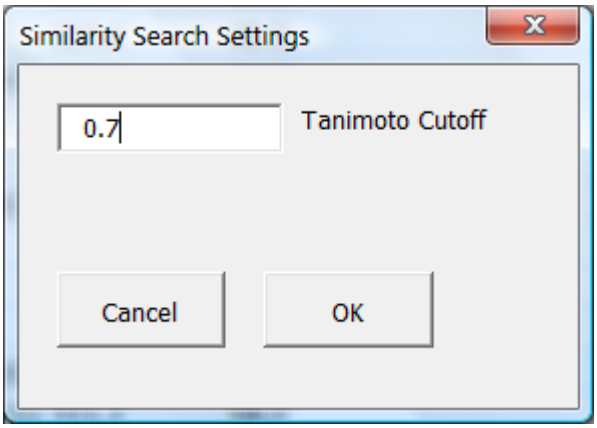

Return to Contents

Name to Smiles Conversion

This program converts chemical names or identifiers to Smiles strings. On selecting this option you will be prompted for which names you want a conversion (default is current selection). A second prompt will ask you to choose between OPSIN or CIR for conversion. OPSIN is designed primarily for IUPAC names (which should exactly follow IUPAC rules including all spelling/capitalisation). CIR is a web service which will convert a variety of formats (including CAS numbers) to Smiles. This option will only work for compounds with entries in the CIR database. Note also that internet access is required for this to work.

Output Smiles strings are put in the Smiles column (if it exists) or (if not) the column to the right of the source names.

If you are using the CIR service and you need a proxy for internet access this should be entered as a LICSS Default (with <Ctrl Shift D>: using proxy_address:proxy_port" as the "Proxy value.

Return to Contents

Fingerprint Generation

The Java program to calculate fingerprints is launched with a progress window (as in Substructure Searching). Fingerprints are stored in a column to the left of the Smiles column. This program may also be conveniently used to calculate fingerprints for new compounds on an existing, enabled sheet – only the new compounds are calculated.

Fingerprints are stored as Hex strings corresponding to 512 bit fingerprints. They are used to speed up Substructure Searching (perhaps most useful on sheets > 2000 compounds)

Return to Contents

R Group Table Creation

This program analyses a compound set as a core structure with a substituent table. On selection, you will first be prompted whether or not you want to include compounds in which the core is not found and whether you want Kekule structures used (means double bonds from fused rings/bicycles are preserved). You will then be prompted to draw the core (as in Substructure Searching). Every position bearing implicit Hydrogen atoms is used as a potential site for substitution. If you want to block a position from substitution, draw an explicit bond to H.

Upon completion of the Java program (with status window as in Substructure Searching), a new sheet is created in the sheet with the Core molecule and substituent table. Substituent connection points are marked with a Xe atom. each Smiles string may be individually visualised without this activation by merely selecting it and choosing <Ctrl R> (or just select <Ctrl Shift M> as normal with the Smiles strings selected to mark them for visualisation

If you wish to calculate an RGroupTable for very large sheets it is worth breaking up the task in to batches of, say 50,000 structures.

Return to Contents

Diverse Compound Picker

This option is used to pick a diverse set from a bigger compound set.

On selecting this option you are first presented with the Compound Selection Dialog:

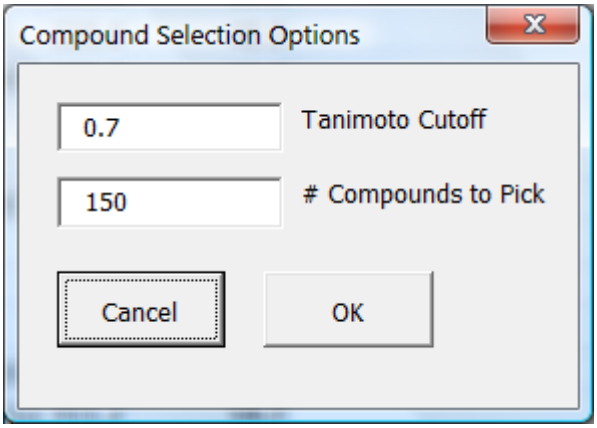

The # Compounds to Pick box defaults to approx 10% of the set (to the nearest 50). You can set this box to the number required. Note that although the set picked for a desired number of hits will be diverse, it will not necessarily be representative as the compounds are picked at random and may not fully sample the chemical space represented by the full set. If you want to pick the full set of compounds at a particular similarity threshold, enter 0 in this box.

The Tanimoto cut-off ensures that no two compounds picked will be more similar than represented by this value. Note that if a very low value is used (< 0.5), then it may not be possible to pick sufficient compounds.

Once the OK button is pressed, the Smiles strings are written out to file and the Java program launched with a status window as in Substructure Searching giving the number of compounds which have been tried.

A Column title is produced showing the number of compounds chosen for picking and the similarity threshold. If it was not possible to pick sufficient compounds (either because 0 was used a number to pick or because the similarity threshold is too low for the number chosen), then the column title will reflect this by prefacing the number actually picked with the text: “Max Pick:”

Return to Contents

Jarvis Patrick Clustering

On selecting this option you will be presented with the Cluster Settings dialog box:

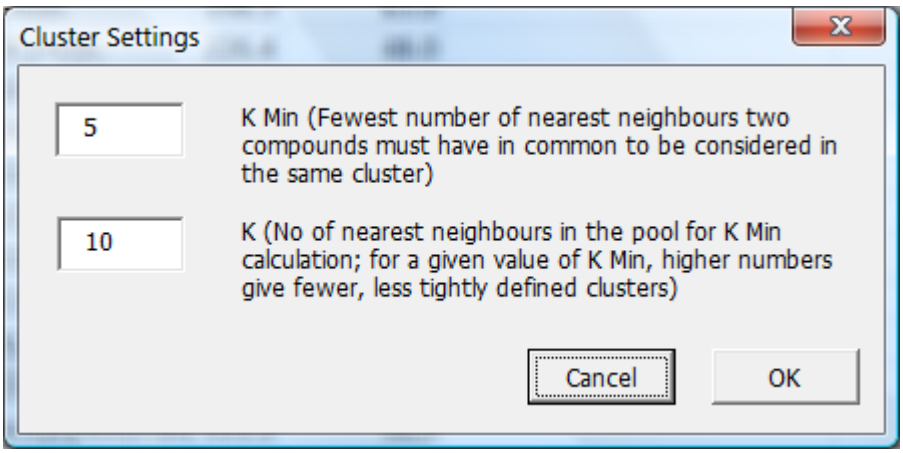

The default settings don't normally need changing. If you want fewer clusters, increase K.

Selecting OK launches the Java program with a status window similar to that for Substructure Searching. Fingerprints are first calculated then the clustering algorithm runs. A full table of inter-compound distances must be calculated and the algorithm scales with N^2. There is therefore a practical limit of *ca* 5000 compounds for clustering using this algorithm (depending on machine memory available)

Once clustering is complete, a new sheet is created with all the clusters described. A new column is also inserted into the original sheet describing which cluster each compound belongs to.

Return to Contents

Sammon Mapping

Sammon mapping is a technique for projecting fingerprint multi-dimensional space into two dimensions. The results can be displayed using Excel's scatter plot facility whence similar compounds cluster together.

On choosing this menu option, you will first be presented with the Sammon Options dialog:

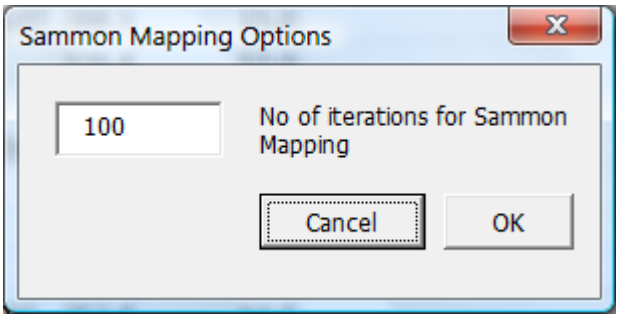

The default value is normally fine. As with Jarvis Patrick clustering, Fingerprints are first calculated and then the Sammon Mapping algorithm runs with a practical limit of about 5000 compounds. As in Substructure Searching a Java status window pops up displaying calculation progress.

Results are presented in Excel as Sammon X and Sammon Y coordinate columns.

Return to Contents

Keyboard Shortcuts

Shortcuts with no menu equivalent

<Ctrl R> - display Row structure (for enabled sheets)/Structure for selected cells (non-enabled sheet in enabled workbook)

<Ctrl T> - toggle display of all structures on/off

<Ctrl Shift A> - (re) run the ActivateStructureDisplay macro

<Ctrl Shift M> - set the sheet name “Smiles” to the currently selected range

<Ctrl Shift Z> - select a column for title of structure display window

Shortcuts with equivalent menu commands

<Ctrl Shift S> - carry out substructure searching

<Ctrl Shift I> - carry out similarity searching

<Ctrl Shift R> - display the Run Menu

<Ctrl Shift E> - pick a diverse subset of compounds

<Ctrl Shift F> - calculate Fingerprints for each structure on the spreadsheet to speed up substructure searching

<Ctrl Shift T> - generate an RGroup Table for any enabled sheet

<Ctrl Shift C> - carry out Jarvis-Patrick Clustering for any enabled sheet

<Ctrl Shift O> - carry out Sammon Mapping for any enabled sheet

<Ctrl Shift N> - carry out Names to Smiles conversion

<Ctrl Shift P> - carry out Molecular Descriptors calculation

<Ctrl Shift H> - output a window of all the shortcut keys described above

Return to Contents

Calculate Molecular Descriptors

This option leads to a prompt for which CDK Descriptor to calculate and this is followed by writing of Smiles strings to file and activation of the Java program to calculate the descriptors. As with Substructure Searching, a Java popup window informs you of progress.

The results of all descriptors are returned in a new column as text, preceded with the name of each descriptor type (some types return several values). The text can be removed by excel using search and replace if the raw data is required (or a Molecular Descriptor Formula can be used)

Return to Contents

Insert Molecular Descriptor Formula (and other new Excel formulas)

This menu option enters a formula: GetCDKDescriptor in to the current selected cell. You can also enter the formula manually, but the menu looks after the parameter syntax for you. Note that some descriptors (such as BCuts) return arrays rather than single values; the last formula parameter is the array index you want to return – if entered *via* the menu, this is assumed to be 1. There is a similarly-used option to insert a Smiles Transform formula – use of this is described in the section on enumeration

There are also formulas for IsSubStructure(query_smiles, target_smiles) and IsSimStructure(query_smiles, target_smiles, cutoff_value), each returning 1or 0 for hits/misses respectively.

A formula ReformatSmiles(Smiles, kekule, expHs) is available which takes a Smiles string as first parameter and reformats it according to the two remaining parameters. Kekule (true/false) determines whether you want kekulised structures or aromatic Smiles (C1=CC=CC=C1 vs c1ccccc1). ExpHs determines whether you want Smiles returned with explicit hydrogens recorded (1 for 'No', 2 for 'All' or 3 for 'As drawn').

If you want to sketch a structure and find out the equivalent Smiles string enter the formula: GetSmilesFromStructure(). A structure edit window will pop up and, on selecting the transfer button (top left) the Smiles string equivalent to the structure drawn will be shown

Return to Contents

Worksheet Enumeration

LICSS can enumerate many libraries where each attachment is made *via* a single/double bond. The actual enumeration is done by direct string manipulation of the Smiles strings and is very fast. To make use of this feature you need to set up a LICSS activated worksheet with columns for Scaffold(s) and associated RGroups.

The scaffold column heading should be "Scaffold" and each RGroup column should be headed one of "Ra", "Rb", "Re", "Rf", "Rg" and "Rh" (the missing "Rc" and "Rd" don't correspond to element symbols). The scaffolds should each be a single Smiles string where the attachment points are indicated with the elements "Ra" – "Rh" as above. Each Smiles string for each RGroup should show the attachment point as a "Xe" atom.

For example:

Scaffold	Ra	Rb	Rf
c1c([Ra])c([Rb])c([Rf])cc1	C[Xe]	CC[Xe]	CCC[Xe]
	C(F)[Xe]	C(F)C[Xe]	C(F)CC[Xe]
	C(F)([Xe])Br	C(F)C([Xe])Br	C(F)CC([Xe])Br

With the worksheet set up like this just select the 'Enumerate Worksheet' option from the LICSS Menu and, after choosing a cell for the enumeration output, the rest of the process is automatic[1].

The required scaffolds may be generated easily, either manually, or by using the GetSmilesFromStructure() formula. For the RGroup arrays, typically a user will import a Smiles or SDF file of reagent structures. These need to be manipulated to create the [Xe] (or [Ra] etc..)-containing Smiles strings. To do this a special formula "SmilesTransform()" may be used. This is most easily accessed by using the option: "Insert Smiles Transform Formula in selected cell" from the LICSS main menu. After selecting the Smiles string you want to transform you will be presented by the Structure Editing window. Draw one or more substructures with attachment points as "Xe" atoms (or use "Ra" etc for scaffolds). When you choose the "Save contents and return to application button (top left), the Smiles string will be transformed with the substructure drawn replaced by [Xe] (or [Ra] etc). If you don't attach a Xe or Ra, etc atom the substructure will just be deleted without being replaced. The reason multiple substructures are allowed is so that reactions which can use multiple types of reagent (eg Br-, Cl-, CF3S03-) can be dealt with in one pass. In the usual Excel way, the formula may be copied down and entire column. Before using the resulting Smiles column, the whole should be 'Copied' and 'Pasted-Special Values' to remove the formula dependency.

Another mode for SmilesTransform exists: if the Smiles submitted has a single atom with a +1 charge (right-click the atom and choose charge +1), then that atom is converted to an Xe atom directly. Note that higher charges +2 etc give Ra etc atoms. This option may be used when only a single atom of the structure needs replacing during enumeration (eg hyrdrazone formation from ketones) but

[1] The actual enumeration consists of making all possible combinations of Rn-Xe bond. However, quite complex libraries can be enumerated by prior manipulation of the Scaffold and Ra… Smiles strings for which the SmilesTransform function is used as described in the main text (bond orders and atom types may be changed at this stage and partial substructures eliminated). Rings may also be formed by use of *two* Ra groups in the same scaffold to match *two* Xe atoms in the same Ra group.

the atom alone is not sufficient to uniquely identify the required position (you can also use the template method below to achieve this).

For more complicated enumeration cases the SmilesTransform function may be used with a template structure to show desired atom/bond transformations. We illustrate this for a Diels Alder Reaction.

For the scaffold, we enter two structures into the JCP editor: firstly, the diene we want to use as an SSS query against the various reagents and, secondly, the way we want this part of the molecule to be transformed by the reaction. Atom mappings are shown with charges +1, +2 etc. The bonds to Ra/Rn[2] in the RHS structure indicate we want these bonds to be added to the transformed reagents. Note that the required charge mappings can most easily be generated by clicking the '+1' button underneath the editing window and then multiple clicking on each atom to be mapped.

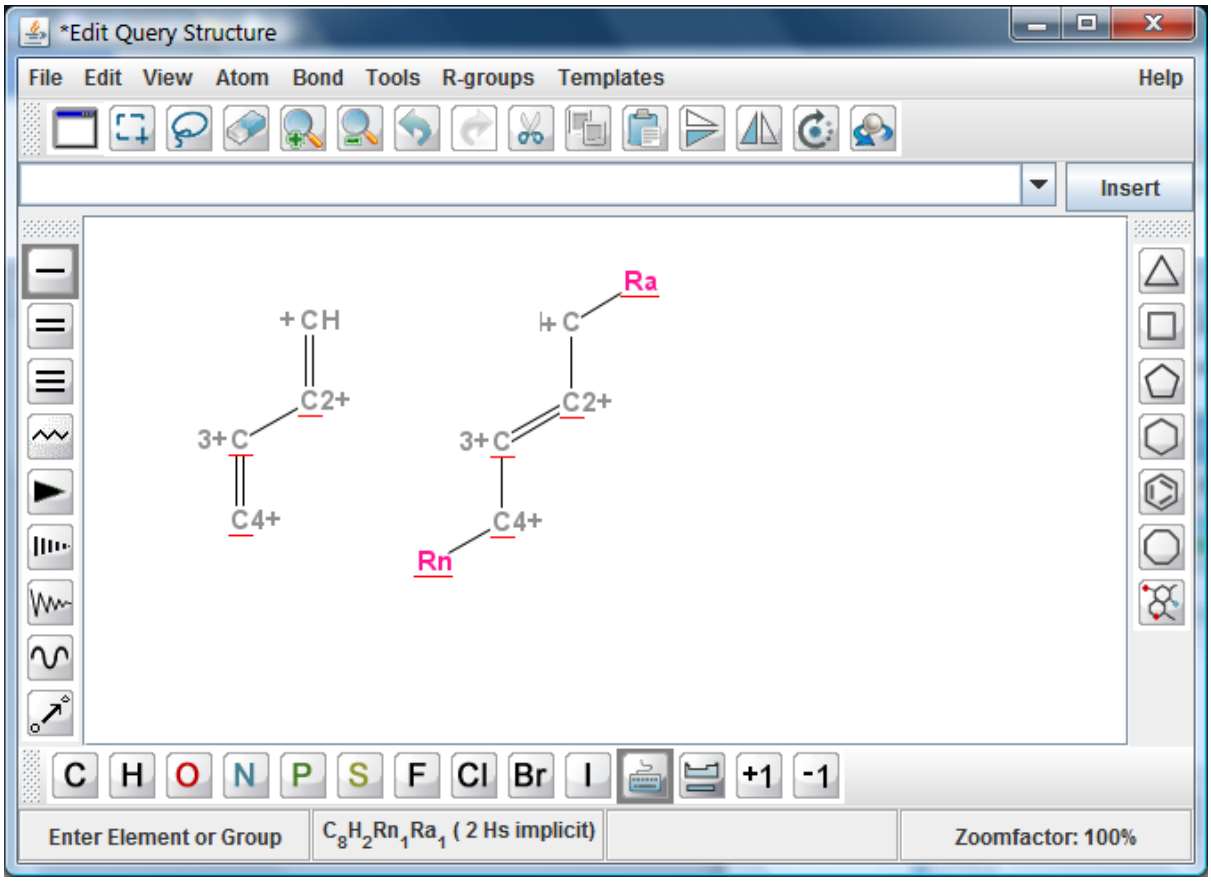

With these structures we get the formula:
=SmilesTransform(A1,"[C+1]=[C+2][C+3]=[C+4].[C+1]([Ra])[C+2]=[C+3][C+4][Rn]", 1) and, if A1 contains the reagent shown below (red box: left) it is transformed to the Smiles shown below (red box: right) and represented in the structure diagram.

[2] The second atom to make a bond linking scaffold and R Group (leading to ring formation) is always designated Rn in both scaffold and R Group

CC=CC=CC	C(=CC(C)[Ra])C(C)[Rn]

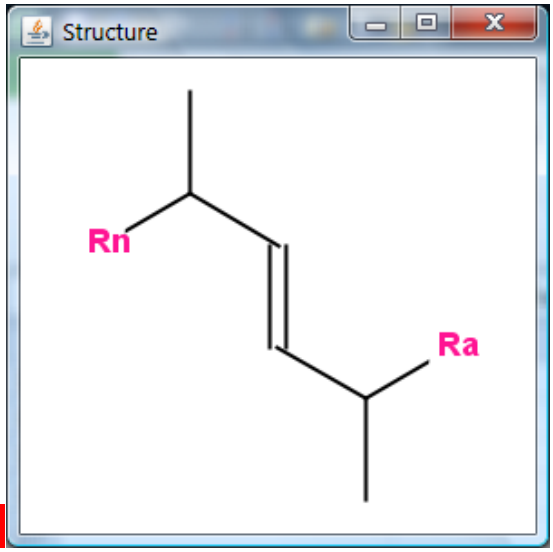

Scaffolds like this may now be enumerated with alkenes transformed by a similarly generated formula to Smiles of the type: C([Xe])CC[Rn] to enumerate to Diels Alder products. Note that mappings used for this method must contain at least three mapped atoms.

If you wish to specify a charge on mapped atoms in the transformed structure then this is done by setting the isotope of the atom in question to a higher or lower isotope (the difference in mass from the most common isotope representing the charge. This is done easily from the structure editor by right-clicking on an atom and selecting 'Properties'. The 'Isotope no' is then adjusted to a number equivalent to the desired charge.

In all modes of use, the SmilesTransform function can return multiple results if more than one substructure mapping is found. By default, only the first of these is returned but, if you wanted a different transform product from that derived *via* the substructure match used, change the index number '1' at the end of the formula to '2' or more until you get the product you wanted.

Return to Contents