

A refnx introduction

Aim

Familiarisation with the refnx GUI, using it to co-refine multiple contrast datasets.

Background

refnx is designed to be used either from a standalone GUI, or a *Jupyter* notebook. In this tutorial you will use the *refnx* GUI to analyse NR datasets from a solvated polymer film that was deposited onto a Silicon wafer. If you get to the end quickly there is an accompanying *Jupyter* notebook that demonstrates the same analysis.

The film has the following structure:



Figure 1

The neutrons are incident onto the sample through the Silicon crystal, this is called the *fronting* medium. In this system the *backing* medium is the semi-infinite water reservoir. We measured the polymer film against D₂O, HDmix (SLD $\sim 3.47\text{e-}6 \text{ \AA}^{-2}$), and H₂O. The SLD of the polymer material is expected to remain unchanged in these different contrasts, but there will be solvent penetration into the film, changing the overall SLD of the layer. We expect the thickness and roughnesses of each of the layers to be unchanged in the different water contrasts.

Good reference sources

- The refnx documentation, <https://refnx.readthedocs.org>
- The refnx paper, <https://github.com/refnx/refnx/blob/master/paper/Nelson2019.pdf>
- YouTube tutorials, <https://www.youtube.com/channel/UCvhOxwZsdFMGqSzasE0ZSOw>
- refnx repository, <https://github.com/refnx/refnx>

First steps

Start an IDAAS session.

Obtain the tutorial data. Select “Applications → Examples → NTC 2021 Course Material”. This will open the NTC2021 directory which holds the “Practical9_RefNX_Demo” subfolder. Copy that directory onto your IDAAS desktop.

Alternatively you can use a browser (“Applications → Utilities → Firefox Web Browser”) to download the tutorial repository at <https://github.com/refnx/tutorial/archive/refs/heads/main.zip>. Unzip it onto the Desktop.

Open a terminal session (“Applications → System → Terminal”)

In the terminal window start the GUI by typing `refnx` (and pressing enter)

Different GUI parts

- Each GUI control has a tooltip. Hover the mouse pointer over each of the items in turn. Can you find the controls for “Insert a Layer”, “Remove a layer”?
- Select the SLD tab, do you see the SLD profiles associated with the models?
- Explore all the main menu items, can you figure out how to save a fit curve?
- Select “misc → SLD calculator”, try working out the SLD for D₂O (mass density = 1.105g/cm³). Can you estimate the SLD of T₂O (H[3]2O, Tritium Oxide), assuming the same molecular volume as H₂O/D₂O?

Estimated molecular volume of H ₂ O / D ₂ O: _____ Å ³
Estimated SLD of T ₂ O: _____ Å ⁻²

- Explore the lipid browser “misc → lipid browser”. This contains a selection of lipids that are relevant to scattering from biological systems. Once you’ve finished close the lipid browser

The **Data Tree** tab contains holds the information for all the datasets being modelled. At the moment there are no loaded datasets, only a ‘theoretical model’ used for making initial simulations. The tree can be expanded and collapsed to increase the level of detail.

- Try navigating the data tree using the cursor/arrow keys to expand and collapse nodes (it’s much quicker than a mouse). Try selecting a Parameter value and press the space bar to select/deselect whether a Parameter will be fitted or not.
- Expand the ‘theoretical model’ fully, what is the thickness of the slab with SLD=3.47?


Loading data

- Press **Ctrl-L** (ctrl and L). This brings up the file open dialogue
- Load ‘e361r.txt’, ‘e365r.txt’, ‘e366r.txt’ from the ‘corefinement’ directory in the simulation repository. Try loading by using the “Data → Load Data” menu items.
- Hide the e365r and e366r datasets by using the ‘display’ checkboxes, the curves corresponding to those datasets should be hidden

Setup the model

- Expand the ‘e361r’ dataset in the data tree, so that you can see the full structure of the model.

The dataset is given a default model of a “fronting medium | slab | backing medium”. We need to add an extra slab layer to the system.

- In the e361r model select the slab in between the fronting and backing medium^a. Press **Ctrl+=** (ctrl and =). This is a keyboard shortcut for adding a Component. Alternatively you can press the  button. The following dialogue should appear.

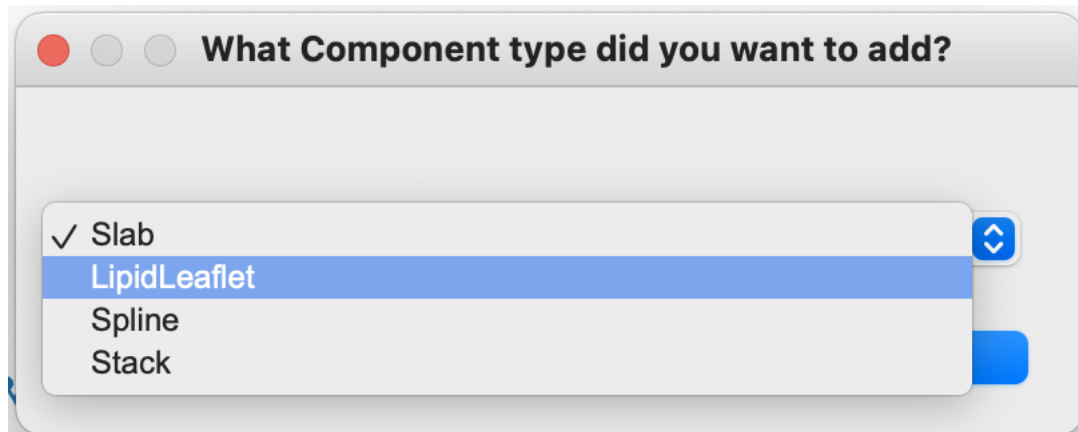


Figure 2

- Add a Slab Component (it should be pretty obvious how to add other types of Components now).
- Once you’ve added the Slab component double-click where it says ‘Slab’ in the model, this will allow you to give each of the components a name. Label the first Slab (after the fronting medium) ‘sio2’, and label the second Slab, ‘polymer’. We can also rename the fronting and backing medium to ‘silicon’ and ‘water’ respectively.

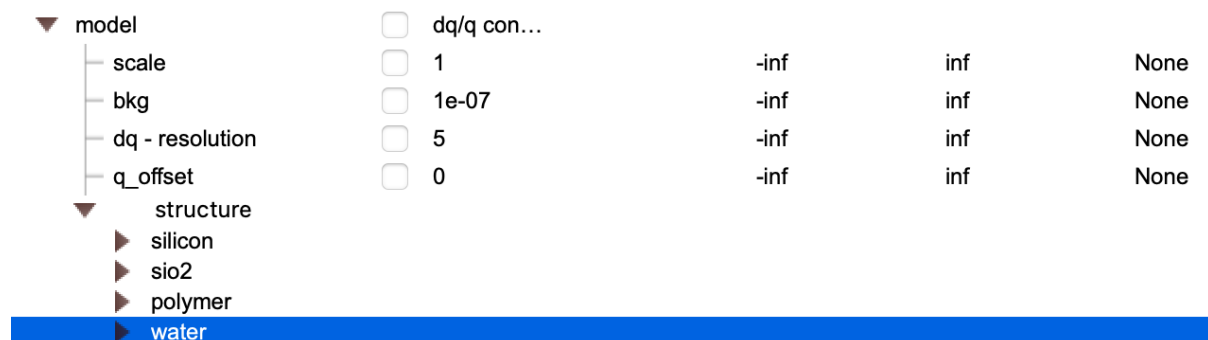



Figure 3

- Expand the ‘silicon’ and ‘water’ nodes, then change the real parts of the ‘fronting’ and ‘backing’ SLDs to 2.07 and 6.36 respectively. Note that the real and imaginary parts of the SLD are multiplied by 1e6 to make them easier to enter, no-one wants to type ‘e-6’ all the time.

^a The new Component is added after the Component you have selected. You can also change the order of Components by dragging them around.

- Using the  button on the reflectivity graph, use the mouse to zoom in on the critical edge area of the dataset. Does changing the SLD values for the fronting/backing mean that the critical edge is in the right place?
- Press **Ctrl-A** to autoscale the graph axes to the full scale of the data again.
- Expand the polymer node and change the real part of the *SLD* of the polymer slab to 1.0.
- Now change the *thick* of the polymer to 130 Å. With the polymer thickness parameter selected, use the slider at the bottom of the window so that the frequency of the Kiessig fringes in the model curve matches that of the data.^b As you vary the thickness look at the change in the SLD profile (“SLD” tab).
- Can you see the χ^2 value change for the e361r dataset as you vary the *thick* value?

What is the rough thickness of the film?

rough thickness = _____ Å

Using the mouse cursor, and the coordinates displayed at the bottom right of the reflectivity plot, find out the *difference* in Q values between two successive minima.

$\Delta Q =$ _____ Å⁻¹

Calculate $2\pi/\Delta Q$, how does this compare to value you estimated using the slider?

$2\pi/\Delta Q =$ _____ Å

- We notice that the background value seems way too low, change the *bkg* to 2e-5.
- Can you get the χ^2 value below 2500 by just changing the *bkg* and *thick* values with the slider?
- Try changing the roughness of some of the layers (in the range 0 to 10 Å) and investigate their effect on the model reflectivity and the SLD profile.
- Save the state of the experiment so far by choosing “File → Save Experiment As”, give it a filename “corefinement.mtft”, and save it somewhere accessible.

^b The slider varies the last selected/alterd parameter. You have to release the slider for the GUI to update.

Did you know?

In refnx the roughness of Component N refers to the interface between layers N-1 / N, and describes the 'standard deviation' of an error function. Layer 0 corresponds to the 'fronting' medium, Layer 1 would correspond to the SiO₂, etc.


For the model in Figure 3 the roughness of the 'sio2' Slab describes the interface between si/sio2. The roughness of the polymer Slab refers to the sio2/polymer interface.

REMEMBER: the roughness parameter has to make physical sense, and shouldn't really be more than half of the thickness of either of the adjoining slabs. Otherwise, the roughness of the layer would extend past either of the slabs. Check that this is the case before reporting a fit.

Example:

Layer	Thickness (Å)	Max Roughness (Å)
N-1	50	25 (but also depends on N-2)
N	100	25
N+1	200	50

Fitting a dataset

- Add the "e361r" dataset to the list of datasets to be fitted. You can do this by selecting the dataset in the data tree, then drag it to the list box on the bottom right hand side of the GUI (Figure 4). Alternatively you can select the dataset and press the  button.
- We have to decide which parameters we want to vary, and apply the bounds we want for each parameter. If you select a checkbox, then that parameter will be varied during the fit.
- Select the parameters we want to fit. Select the checkboxes for:
 - scale
 - bkg
 - sio2-thick
 - sio2-rough
 - polymer-thick
 - polymer-sld
 - polymer-rough
 - backing-sld
 - backing rough (upper bound should be 10 Å)

- It's necessary to give each of the fitted parameters physically reasonable bounds. The first step is to press the “Auto adjust limits” button, which gives a default lower bound of 0, and a default upper bound of twice the current parameter value. The roughness parameters should have upper bounds of 8 Å (make the backing roughness 10 Å).
- Given that the “e361r” was measured in D₂O, what do you think a reasonable upper bound might be for the SLD of the water? What about a reasonable lower bound? Why are we fitting the SLD of the D₂O at all?

Name	value	sigma	lb	ub	constraint
theoretical	<input type="checkbox"/> display	points: 1000	chi2: nan		
e361r	<input checked="" type="checkbox"/> display	points: 99	chi2: 2920.53		
model	<input checked="" type="checkbox"/> dq/q const. ...				
scale	<input checked="" type="checkbox"/> 1		0	2	None
bkg	<input checked="" type="checkbox"/> 1.46231e-05		0	2.92461e-05	None
dq - resolution	<input type="checkbox"/> 5		-inf	inf	None
q_offset	<input type="checkbox"/> 0		-inf	inf	None
structure					
si					
fronting - sld	<input type="checkbox"/> 2.07		-inf	inf	None
sio2					
1 - thick	<input checked="" type="checkbox"/> 15		0	30	None
1 - sld	<input type="checkbox"/> 3.47		-inf	inf	None
1 - isld	<input type="checkbox"/> 0		-inf	inf	None
1 - rough	<input checked="" type="checkbox"/> 3		0	6	None
1 - volfrac solv...	<input type="checkbox"/> 0		0	1	None
polymer					
thick	<input checked="" type="checkbox"/> 209.146		0	418.292	None
sld	<input checked="" type="checkbox"/> 1		0	2	None
isld	<input type="checkbox"/> 0		-inf	inf	None
rough	<input checked="" type="checkbox"/> 3		0	6	None
vlsolv	<input type="checkbox"/> 0		0	1	None
water					
backing - sld	<input checked="" type="checkbox"/> 6.36		0	12.72	None
backing - isld	<input type="checkbox"/> 0		-inf	inf	None
backing - rough	<input checked="" type="checkbox"/> 3		0	6	None
e365r	<input type="checkbox"/> display	points: 99	chi2: 96551.7		
e366r	<input type="checkbox"/> display	points: 99	chi2: 95247.9		

Figure 4 The e361 dataset is selected in the data tree, and the e361 dataset will be fitted (bottom right hand corner).

Data weighting

- If you have the “use errors” checkbox selected then the fit will use the measurement uncertainties supplied in the datafile (third column, standard deviation).
- You can select other weighting options via the “Fitting→Fit As” menu options. There you can select “lnY vs X”/“logY vs X”/“YX⁴ vs X”/“YX² vs X”.

Data weighting guide for neutrons

- A good default is logY vs X, **with** “use errors” selected. However, not using error bar weighting typically works for logY as well.
- Choosing other types of weighting options can highlight different aspects of the reflectivity curve. YX^4 is often used because the reflectivity from an infinitely sharp interface decays as $R \sim Q^{-4}$ above the critical edge. If you choose an option other than logY then use error bar weighting.
- **Do not fit as linY vs Q when the data is not weighted by the measurement uncertainties.**

Data weighting guide for X-rays

- **A good default is logY vs X, with “use errors” de-selected.** This is because the counting uncertainties are typically much smaller than systematic uncertainties. It’s hard to estimate those systematic uncertainties, so the y-error bars are typically way too small. Using counting uncertainties then causes the fit to be dominated by data near the critical edge.
- **It’s not advised to fit as linY vs X, with or without error bars.**

Instrumental Resolution Smearing

- If the data file contains the instrumental resolution smearing information (as a fourth column), then you should uncheck “dq/q const. smearing”, and refnx will use the information from the file. This is referred to as point-wise smearing.
- Otherwise NR instruments typically measure with a constant fractional dq/q resolution. You then need to select the “dq/q const. smearing” checkbox and enter a value for the “dq-resolution” parameter.
- **With the dq/q const. smearing checkbox selected, select the dq-resolution parameter, then use the slider to vary the resolution. What happens to the model reflectivity curve as the fractional dq/q resolution is varied from 0 to 20?**

-
- **Zoom in on the critical edge region (using the magnifying glass icon on the graph toolbar). What happens to the model curve around the critical edge as the fractional dq/q resolution is varied from 0 to 20?**
-


-
- Change the dq/q resolution back to 5. You can reset the graph view by pressing **Ctrl-A**.
 - **If you're fitting X-ray reflectometry data one typically sets the fractional dq/q resolution to 0. It's rare that an XRR instrument will provide the instrument resolution information.**

Further information on resolution smearing can be found here:

<https://refnx.readthedocs.io/en/latest/faq.html#how-is-instrumental-resolution-smearing-handled>

Fitting algorithm

The fitting algorithm can be chosen using the menu options "*Fitting → Algorithm*". Differential-evolution is a good default, and is frequently used^c for analysis of NR/XRR data. Each of the fitting algorithms should provide an estimate of the model parameters. However, the most robust estimate of parameter uncertainties is derived when using Markov Chain Monte Carlo sampling.

- Make sure the "differential-evolution" algorithm is selected.
- Start the fit by pressing the  button. You should see a progress dialogue appear as the optimisation runs. When it disappears the fit has finished.
- Have the fit parameters changed?
- Look at the fit uncertainties in the "sigma" column of the data tree. Do they make sense for the kind of data we've analysed?
- For a thick film with plenty of Kiessig fringes do you expect the uncertainty in film thickness to be:

LARGE / SMALL?
(circle one)
- In the console tab you can also see information on the fit. Historical information of previous fits is written there.
- Save the experiment again ("File → ...")

^c Wormington et al., Characterization of structures from X-ray scattering data using genetic algorithms, Philosophical Transactions of the Royal Society of London, 1999 , 357, 2827-2848

- Try altering the polymer film thickness using the slider, observe how χ^2 changes as you make small changes to the parameter value. Now go back to the optimal thickness value (either by loading the saved experiment or redoing the fit), and do the same for the different roughness parameters.

As a percentage of parameter value, is χ^2 more sensitive to the thickness or roughness?

Which of the roughness parameters is χ^2 most sensitive to?

- How sensitive is χ^2 to the value of the backing SLD?

Co-refining datasets

Now that we can fit an individual dataset we will fit the other two contrasts.

- In the data tree select the e365 and e366 datasets. With the datasets selected right-click to bring up the *contextual* menu (Figure 5).

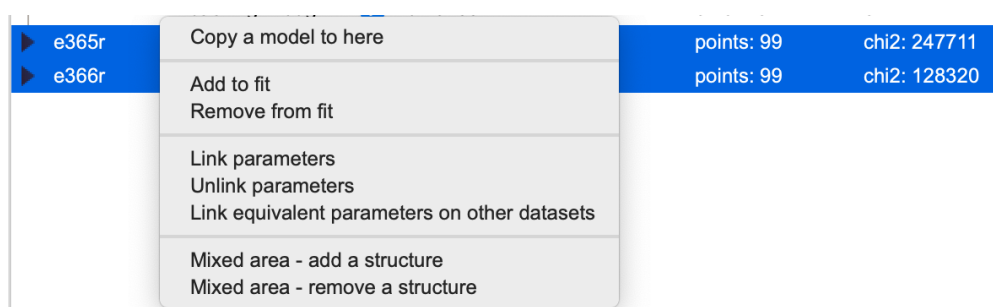


Figure 5 The contextual menu presented when one right-clicks on the data tree.

- Choose the “Copy a model to here” option.

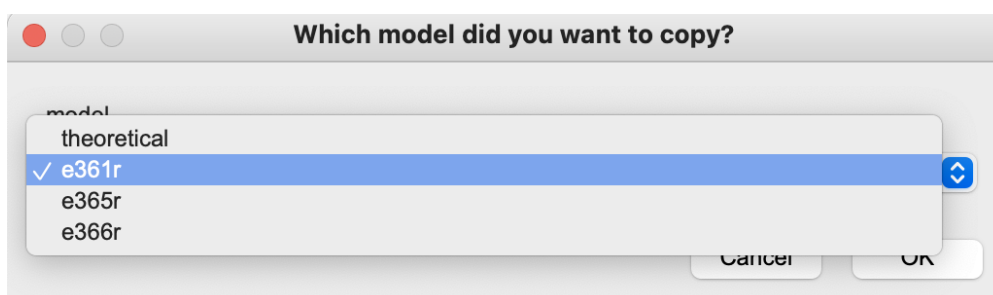


Figure 6

- In the dialogue that appears select the e361 dataset. This will copy the model setup from the e361 dataset to the e365 and e366 datasets. At this point no parameters are shared between any of the models.

- Make the e365 and e366 graphs visible by selecting the “display” checkbox on both the datasets in the data tree.
- Expand the model for e365 and e366, and change the real part of the backing SLD for e365 (HDmix) and e366 (H₂O) to 3.47 and -0.56 respectively. We don’t want to fit the backing SLD for e366, so deselect its checkbox. We do want to fit e365, but change it’s lower and upper bounds to [3.3, 3.5].

Dataset	Parameter	Value	Lower Bound	Upper Bound	Constraint
e365r	display	<input checked="" type="checkbox"/>			
	points	99			
	chi2	827.374			
	dq/q const. smearing?	<input checked="" type="checkbox"/>			
	scale	1.01356	0	2	None
	bkg	1.78879e-05	0	2.92461e-05	None
	dq - resolution	5.36403	-inf	inf	None
	q_offset	0	-inf	inf	None
	backing - sld	3.47	3.3	3.5	None
	backing - isld	0	-inf	inf	None
e366r	display	<input checked="" type="checkbox"/>			
	points	99			
	chi2	1581.64			
	dq/q const. smearing?	<input checked="" type="checkbox"/>			
	scale	1.01356	0	2	None
	bkg	1.78879e-05	0	2.92461e-05	None
	dq - resolution	5.36403	-inf	inf	None
	q_offset	0	-inf	inf	None
	backing - sld	-0.56	6.1	6.36	None
	backing - rough	6.28733	0	10	None

Figure 7

Making linkages between parameters

A sample measured in different solvent contrasts is often expected to have the same interfacial structure for each of the measurements. These datasets are frequently analysed at the same time, a process called *co-refinement* or *global fitting*. During co-refinement there will be parameters that are common across all the datasets. For example, in this tutorial we expect the polymer layer thickness to be the same for all three measurements. We therefore need to *link* the sets of common parameters. There are a few methods of doing so:

1. Select the polymer layer thickness parameter for e361. Press **Ctrl-3**, this is the keyboard shortcut for the menu item “*Model → Link equivalent parameters*”. It’s also possible to do this by the right-click contextual menu (Figure 5). A dialogue appears asking you to ‘select equivalent datasets to link’, select all the datasets (Figure 8). This process requires that all selected datasets have a similar interfacial model. If you look at the *constraint* column for the polymer thickness, for e365 and e366, you should see that it has a constraint applied to it.
2. It’s also possible to make linkages for any parameter combinations. To do this select all the parameters you want to link, then press **Ctrl-1**; this is the keyboard shortcut for the menu item “*Model → Link Parameters*”. Try this by linking the *backing-rough* parameter for all the datasets (roughness of the polymer/water interface), Figure 9. To select multiple parameters hold the **Ctrl** button down while mouse clicking on the

parameters you want to choose. The 'master' parameter will be the first parameter selected.

To unlink a parameter, select the parameters you wish to unlink and press **Ctrl-2** ("Model → Unlink Parameter").




Figure 8

backing - rough		<input checked="" type="checkbox"/>	6.28733		0	10	None
e365r		<input checked="" type="checkbox"/>	display	points: 99	chi2: 1812.14		
model		<input checked="" type="checkbox"/>	dq/q const. smearing?				
scale		<input checked="" type="checkbox"/>	1.01356		0	2	None
bkg		<input checked="" type="checkbox"/>	1.78879e-05		0	2.92461e-05	None
dq - resolution		<input type="checkbox"/>	5.36403		-inf	inf	None
q_offset		<input type="checkbox"/>	0		-inf	inf	None
structure							
si							
sio2							
polymer							
thick		<input type="checkbox"/>	222.007		0	418.292	Parameter(val...
sld		<input checked="" type="checkbox"/>	0.398808		0	2	None
isld		<input type="checkbox"/>	0		-inf	inf	None
rough		<input checked="" type="checkbox"/>	2.64826		0	6	None
vfsolv		<input type="checkbox"/>	0		0	1	None
water							
backing - sld		<input checked="" type="checkbox"/>	3.47		3.3	3.5	None
backing - isld		<input type="checkbox"/>	0		-inf	inf	None
backing - rough		<input type="checkbox"/>	6.28733		0	10	Parameter(val...
e366r		<input checked="" type="checkbox"/>	display	points: 99	chi2: 1542		
model		<input checked="" type="checkbox"/>	dq/q const. smearing?				
scale		<input checked="" type="checkbox"/>	1.01356		0	2	None
bkg		<input checked="" type="checkbox"/>	1.78879e-05		0	2.92461e-05	None
dq - resolution		<input type="checkbox"/>	5.36403		-inf	inf	None
q_offset		<input type="checkbox"/>	0		-inf	inf	None
structure							
si							
sio2							
polymer							
water							
backing - sld		<input type="checkbox"/>	-0.56		6.1	6.36	None
backing - isld		<input type="checkbox"/>	0		-inf	inf	None
backing - rough		<input type="checkbox"/>	6.28733		0	10	Parameter(val...

Figure 9 All the backing rough parameters are selected and linked. Notice that the polymer thickness for e365 is constrained.


- Try varying the polymer layer thickness, do you see all the model reflectivity curves changing simultaneously?
- Now link these parameters in separate sets across all the datasets. You can use either of the methods above.
 - sio2 thickness
 - sio2 roughness
 - polymer SLD
 - polymer roughness
 - polymer vfsolv
- We want to fit all these parameters so make sure the ‘master’ parameter has its checkbox selected.
- Make sure the lower and upper bounds for “polymer vfsolv” are [0, 1]. The solvent penetration is done by the backing medium. If you use *refnx* from a Python script you can make the solvent penetration any SLD you like (e.g. the fronting medium).

Example calculation: overall SLD of a layer

Let $\rho_{polymer} = 0.3$, $\rho_{solv} = 6.36$, $\phi_{solv} = \text{vfsolv} = 0.2$. The overall SLD of a layer is:

$$\rho = (1 - \phi_{solv})\rho_{polymer} + \phi_{solv}\rho_{solv}$$

$$\rho = 1.512 = 0.3 \times 0.8 + 0.2 \times 6.36$$

- Add the e365 and e366 datasets to the list of datasets to be fitted (Figure 4).
- Do the co-refinement by pressing . Hopefully your graphs look like Figure 10 and Figure 11.
- Save the experiment.




Figure 10




Figure 11

Exporting models/fits/graphs

You can export the reflectivity and SLD graphs by pressing the disk button in the toolbars of their respective graph windows.

You can export the fit curves and SLD profiles as files from the “*Model* → *Save fit curve*” and “*Model* → *Save SLD profile*” menu options.

You can export the fit coefficients and uncertainties as a CSV file using “*Model* → *Export Parameters*”.

Estimating Parameter Uncertainties using Markov Chain Monte Carlo (MCMC).

The most robust way of estimating parameter uncertainties is to derive the posterior probability distribution, $p(\theta|D, M)^d$.

- Select “Fitting → Algorithm → MCMC”.

Before you start an MCMC sampling run it’s a good idea to start from a situation where the model is reasonably fitted. This should be the case at this point in the tutorial.

- Start a fit.
- A dialogue appears, Figure 12, asking how many steps you’d like to perform. To start with we’re going to do a quick investigation, so we’re going to run around 500 steps. The different ways of initialising the 200 MCMC chains (walkers) are covar/jitter/prior:
 - covar – the walkers are distributed according to the covariance matrix around the current parameter location. This is a good option if you already have a good fit as the covariance matrix is related to the parameter uncertainties.
 - jitter – the walkers are perturbed with a small amount of noise. It may take longer for the walkers to distribute around their equilibrium positions.
 - prior – randomly distribute the walkers according to the priors (bounds) on each of the parameters.

The *thin* parameter controls the spacing between saved samples. For example, if thin=100 then 1 in every 100 samples are saved. This parameter is used to reduce autocorrelation (more on this later).

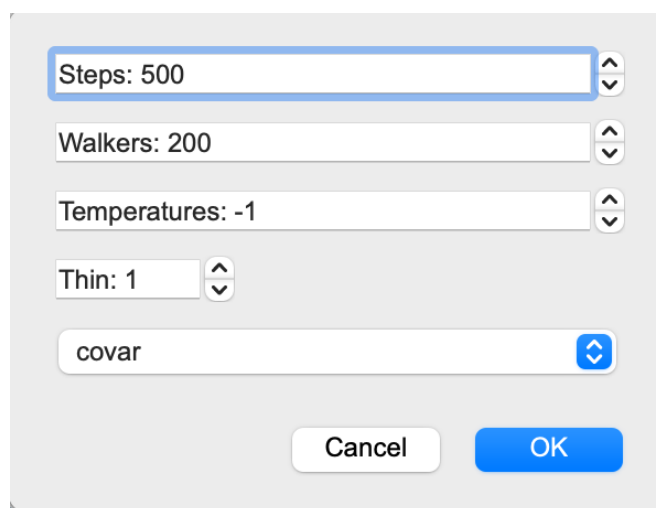
A dialog box for configuring MCMC sampling. It contains five input fields with up/down arrows: 'Steps: 500', 'Walkers: 200', 'Temperatures: -1', 'Thin: 1', and a dropdown menu set to 'covar'. At the bottom are 'Cancel' and 'OK' buttons.

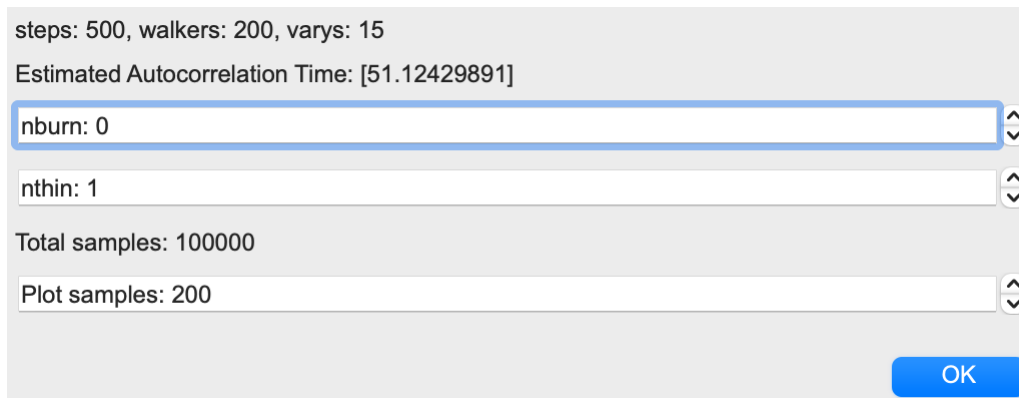
Figure 12

- When you click on “OK” a folder dialogue appears, this location is where the output of the MCMC sampling is saved. It’s a good idea to create a folder specifically for this purpose as there are a lot of files. Once you’ve chosen this directory the MCMC

^d <https://github.com/refnx/refnx/blob/master/paper/Nelson2019.pdf>

sampling starts, and a progress dialogue is presented. The sampling process takes ~100 sec on my 2015 quad-core iMac.

- Once the sampling finishes you'll be asked how you want to process the sampling run, Figure 13. Note the estimated autocorrelation time; this is the distance from any given step to when the next independent step is obtained. Leave *nthin* and *nburn* at their default values for the time being.



steps: 500, walkers: 200, varys: 15

Estimated Autocorrelation Time: [51.12429891]

nburn: 0

nthin: 1

Total samples: 100000

Plot samples: 200

OK

Figure 13

- Examine the MCMC files created during the sampling step, let's look at "steps-autocorrelation.png" to start with, Figure 14. It looks like zero autocorrelation is obtained after ~100 steps. However, this isn't the best autocorrelation graph, the curves drop significant below the X-axis. Normally to get a good estimate of the autocorrelation time one has to run for tens of times longer than the autocorrelation time. For example, if the autocorrelation time is 100, then acquiring 200 steps would only give ~2 samples contributing to this graph for a lag time of 100. This graph tells us that we should *thin* by a factor of 100, either during sampling or post processing.




Figure 14 The autocorrelation of each parameter in the MCMC chains.

- Figure 15 shows the variation of the backing SLD (*steps_param_8.png* for me) for dataset e361 as a function of step number. You'll notice that the parameter doesn't become distributed around its equilibrium position until step ~ 70 . You should examine all these graphs and figure out how long it takes for the parameters to reach their final distribution. This value will be the 'burn' time; so-called because this number of steps are simply discarded.




Figure 15

- Now repeat the MCMC sampling process by pressing the fit button. This time change the *thin* parameter to 100. Change the number of steps to 15 (a total of $15 * 100 = 1500$ steps will be run). This takes about 4.5 mins on my desktop. While you're waiting for the sampling to finish have a look at the other MCMC graphs from the previous sampling run.
This time around the estimated autocorrelation time is ~ 1 . When you process the data change *nburn* to 1 (we want to burn a total of 100 steps).
- The MCMC chain is saved in a file called *steps.chain*. You can reprocess this file at any time with "*Fitting \rightarrow Algorithm \rightarrow Process MCMC chain*". However, the chain must correspond to the setup currently contained in the GUI.
- Once the chain has been processed the parameter values and uncertainties are updated in the GUI. The parameter value is taken as the median value of the samples, with its uncertainty estimated as half the [15.9, 84.1] range. This uncertainty corresponds to the standard deviation of a normal distribution.
- The "steps_corner.png" image shows the covariance matrix for the system, Figure 16. If you look at the polymer roughness parameter (7th column from left) you can see that the distribution is truncated and is butting up against the upper limit of 8. Perhaps we should've set a higher upper bound.




Figure 16

- steps.png* and *steps_sld.png* show the reflectivity and SLD curves for the data, curves corresponding to the median of the MCMC samples, and plots for a certain number of the samples. This gives an impression of the spread of fits and spread of structures consistent with the data. See Figure 17 for the spread of SLD profiles. There you can see that the main uncertainty in the SLD profiles concerns the native SiO₂ layer. The slight dip when transitioning from Si to SiO₂ indicates that the Si/SiO₂ and SiO₂/polymer roughnesses might be too high compared to the SiO₂ thickness. Given that the fit isn't overly sensitive to the Si/SiO₂ and SiO₂/polymer roughnesses it might be reasonable to fix the roughness values at a reasonable value (3 Å?).




Figure 17 Spread of structural profiles for the three co-refined datasets.

Further work

If you've followed this tutorial successfully, and you still have time, then investigate:

- adding different types of Component, such as the LipidLeaflet.
- making a multilayer (hint: add a Stack Component).
- how you might analyse patchy layers (hint: the contextual menu, Figure 5, shows a "Mixed Area – add a structure" item).
- the same analysis, but performed in a *Jupyter* notebook. This is available in the tutorial repository as *corefinement/reflectometry_global.ipynb*. On IDAAS you will have to change the kernel: "Kernel → refnx".