

Fitting Data with *Mathematica* for a Nonlinear Model

In a number of experiments this semester, you will be asked to fit experimental data to a theoretical curve where a nonlinear fit is needed. This notebook walks through an example of calculating chi-squared in such a case where LinearModelFit won't work.

For this notebook, we will look at fitting a simple oscillating function to some data.

```
In[45]:= Clear["Global`*"]; (* Clear out any stored definitions. *)
DateString[]
```

```
Out[46]=
Mon 16 Feb 2026 10:19:56
```

Entering Data

Importing from a file

```
In[47]:= SetDirectory[NotebookDirectory[]]
```

```
Out[47]=
/Users/doughera/notes/www/public_html/courses/phys238-2026/notes/Ch08
```

```
In[48]:= FilePrint["nonlinear-modelfit-1-data.txt", 5] (* Print out the first 5 lines *)
```

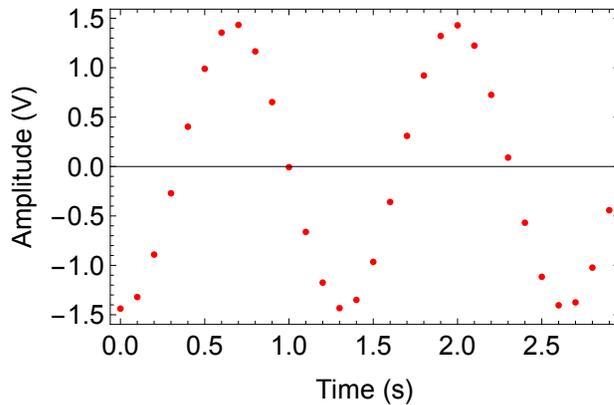
```
0.00  -1.438
0.10  -1.321
0.20  -0.891
0.30  -0.271
0.40   0.403
```

This data file is just two columns of numbers separated by spaces. Mathematica can import this as a simple table.

```
In[49]:= data = Import["nonlinear-modelfit-1-data.txt", "Table"];
```

```
In[50]:= dataplot = ListPlot[data, PlotRange → All,
  PlotStyle → Red, Frame → True, LabelStyle → Larger,
  FrameLabel → {"Time (s)", "Amplitude (V)"}, ImageSize → Scaled[0.5]]
```

Out[50]=



Fitting a Model to the Data

For this exercise, we will try to fit this to a cosine function, allowing for the possibility that the average value might not be zero, due to some offset in the sensor. This function has 4 free parameters: A , ω , ϕ , and $xoff$.

```
In[51]:= x[A_, ω_, φ_, t_, xoff_] := A Cos[ω t + φ] + xoff
```

```
In[52]:= Clear[chisq]
```

```
chisq[data_, A_, ω_, φ_, xoff_] :=
  Sum[(data[[i, 2]] - x[A, ω, φ, data[[i, 1]], xoff])2, {i, 1, Length[data]}] /
  (Length[data] - 4)
```

For convenience of our `chisq` visualization, pick Amplitude and offset values that seem plausible. Look at varying ω and ϕ .

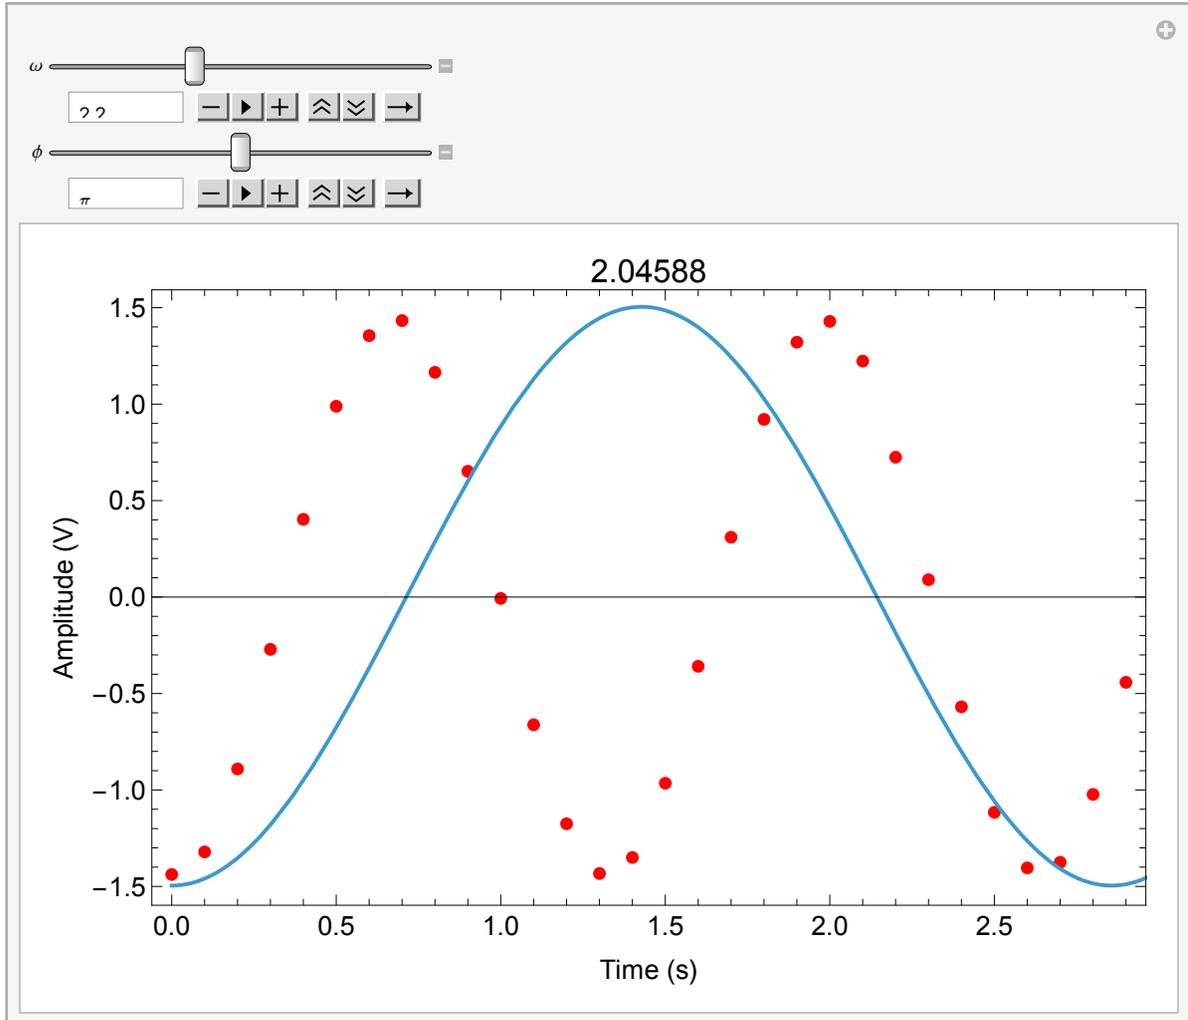
```
In[54]:= myxoff = 0.0044; myA = 1.5;
```

```

In[56]:= Manipulate[
  Show[{
    dataplot,
    Plot[x[myA,  $\omega$ ,  $\phi$ , t, myxoff], {t, 0, 3}]
  }, ImageSize  $\rightarrow$  Large, PlotLabel  $\rightarrow$  chisq[data, myA,  $\omega$ ,  $\phi$ , myxoff]],
  {{ $\omega$ , 0.65}, 0, 6, 0.05, Appearance  $\rightarrow$  "Open"},
  {{ $\phi$ ,  $\pi$ }, 0,  $2\pi$ , 0.1, Appearance  $\rightarrow$  "Open"}
]

```

Out[56]=



Minimizing χ^2

One-dimensional minimization

Picking $A = 1.5$ and a phase of π , look at how χ^2 varies with ω .

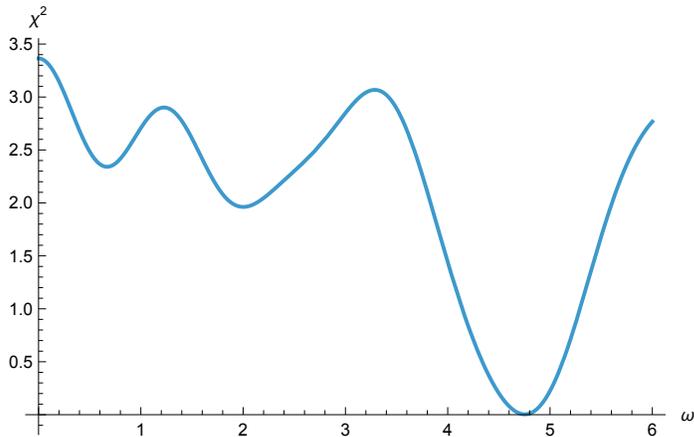
```

In[57]:= myA = 1.5; my $\phi$  =  $\pi$ ;

```

```
In[58]:= Plot[chisq[data, myA,  $\omega$ , my $\phi$ , myxoff], { $\omega$ , 0, 6}, AxesLabel -> {" $\omega$ ", " $\chi^2$ "}]
```

```
Out[58]=
```



There clearly is a global minimum near 4.8, but there are also local minima near 0.7 and 2.0. Without help, *Mathematica* might get trapped in a local minimum.

```
In[59]:= FindMinimum[chisq[data, myA,  $\omega$ , my $\phi$ , myxoff],  $\omega$ ]
(* Mathematica tries starting at  $\omega = 1$  *)
```

```
Out[59]=
```

```
{2.34117, { $\omega \rightarrow 0.670019$ }}
```

```
In[60]:= FindMinimum[chisq[data, myA,  $\omega$ , my $\phi$ , myxoff], { $\omega$ , 3}] (* Tell it to start at 3 *)
```

```
Out[60]=
```

```
{1.96213, { $\omega \rightarrow 2.00135$ }}
```

This finds the second minimum. An even better initial guess will find the global minimum.

```
In[61]:= FindMinimum[chisq[data, myA,  $\omega$ , my $\phi$ , myxoff], { $\omega$ , 4}] (* Tell it to start near 4 *)
```

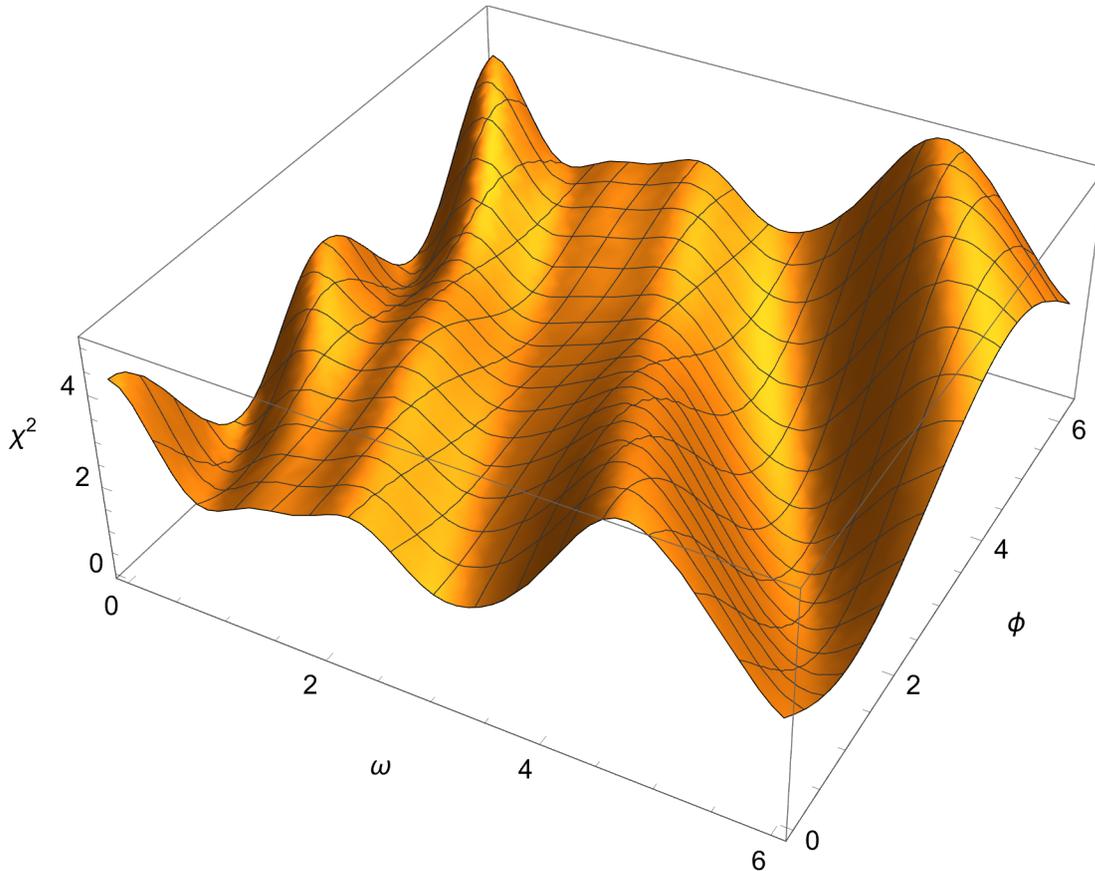
```
Out[61]=
```

```
{0.0036504, { $\omega \rightarrow 4.75364$ }}
```

Two-dimensional minimization

```
In[64]:= Plot3D[chisq[data, myA,  $\omega$ ,  $\phi$ , myxoff], { $\omega$ , 0, 6}, { $\phi$ , 0,  $2\pi$ }, PlotRange  $\rightarrow$  All,  
AxesLabel  $\rightarrow$  {" $\omega$ ", " $\phi$ ", " $\chi^2$ "}, LabelStyle  $\rightarrow$  Larger, ImageSize  $\rightarrow$  Large]
```

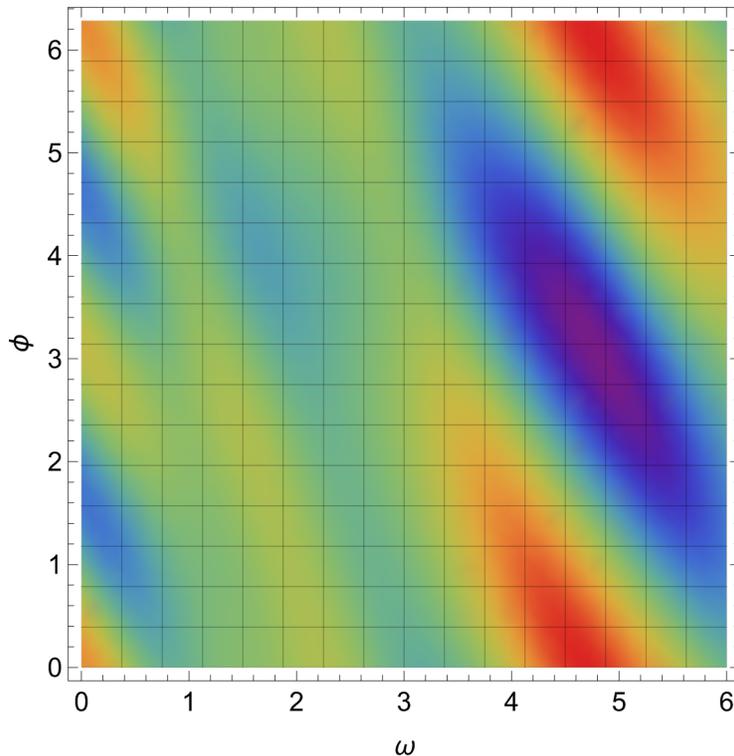
Out[64]=



Sometimes it helps to visualize this sort of thing as a 2d density plot. The minimum near $\omega = 4.75$ and $\phi = \pi$ is evident.

```
In[66]:= DensityPlot[chisq[data, myA,  $\omega$ ,  $\phi$ , myxoff], { $\omega$ , 0, 6}, { $\phi$ , 0, 2  $\pi$ },
  PlotRange  $\rightarrow$  All, FrameLabel  $\rightarrow$  {" $\omega$ ", " $\phi$ "}, LabelStyle  $\rightarrow$  Larger,
  Mesh  $\rightarrow$  True, ColorFunction  $\rightarrow$  "Rainbow", ImageSize  $\rightarrow$  Scaled[0.6]]
```

Out[66]=



Here again is a failed search. Given our guess for amplitude and offset, find the best values for ω and ϕ :

```
In[68]:= result = FindMinimum[chisq[data, myA,  $\omega$ ,  $\phi$ , myxoff], { $\omega$ ,  $\phi$ }]
```

Out[68]=

```
{1.18902, { $\omega$   $\rightarrow$  0.0753507,  $\phi$   $\rightarrow$  1.5484}}
```

Mathematica chooses initial starting values of '1'. This gets caught in the local minimum on the left hand side. However, if you give it a good starting point for at least one of the parameters, *Mathematica* finds the minimum correctly. Here we will let it find all of the parameters, but give a good starting hint. (This gives a negative amplitude. More on that below.)

```
In[69]:= result = FindMinimum[chisq[data, A,  $\omega$ ,  $\phi$ , off], {A, { $\omega$ , 4},  $\phi$ , off}]
```

Out[69]=

```
{0.0000598132, {A  $\rightarrow$  -1.43493,  $\omega$   $\rightarrow$  4.7841,  $\phi$   $\rightarrow$  -0.0628613, off  $\rightarrow$  0.00441981}}
```

NonlinearModelFit

`NonlinearModelFit[]` performs this minimization automatically. Though it has certain algorithms to attempt to avoid local minima, you still sometimes need to give it hints.

```
In[70]:= fit = NonlinearModelFit[data, x[A,  $\omega$ ,  $\phi$ , t, xoff], {A,  $\omega$ ,  $\phi$ , xoff}, t]
```

⋯ NonlinearModelFit: Failed to converge to the requested accuracy or precision within 100 iterations.

```
Out[70]=
```

```
FittedModel [ -879. + 879. Cos [0.0479 - 0.0355 t ] ]
```

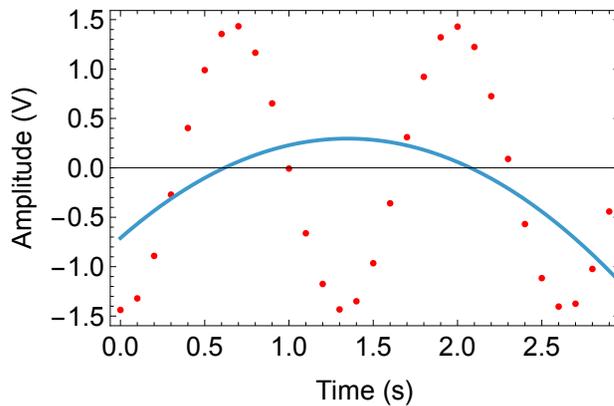
```
In[71]:= fit["BestFitParameters"]
```

```
Out[71]=
```

```
{A → 879.285,  $\omega$  → 0.0355321,  $\phi$  → -0.0478572, xoff → -878.989}
```

```
In[72]:= fitplot = Plot[fit[t], {t, 0, 3}, PlotRange → All];
Show[dataplot, fitplot]
```

```
Out[73]=
```



```
In[74]:= Sqrt[fit["EstimatedVariance"]] (* The typical error is about 1 volt. *)
```

```
Out[74]=
```

```
1.01463
```

This fit is quite poor. Give it a hint for ω , and it does much better.

```
In[75]:= fit = NonlinearModelFit[data, x[A,  $\omega$ ,  $\phi$ , t, xoff], {A, { $\omega$ , 4},  $\phi$ , xoff}, t];
fit["BestFitParameters"]
```

```
Out[76]=
```

```
{A → -1.43493,  $\omega$  → 4.7841,  $\phi$  → -0.0628613, xoff → 0.00441981}
```

```
In[77]:= Sqrt[fit["EstimatedVariance"]] (* The typical error is now about 8 mV. *)
```

```
Out[77]=
```

```
0.0077339
```

However, note the negative amplitude and near zero phase. If all you want is the frequency, then that is fine. However, if you want a positive amplitude, then you also have to give a hint for the phase:

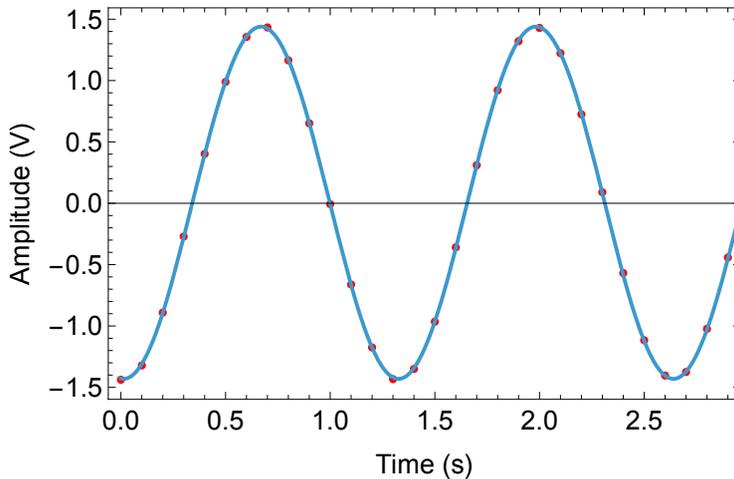
```
In[80]:= fit = NonlinearModelFit[data, x[A,  $\omega$ ,  $\phi$ , t, xoff], {A, { $\omega$ , 4}, { $\phi$ ,  $\pi$ }, xoff}, t];
fit["BestFitParameters"]
```

```
Out[81]=
```

```
{A → 1.43493,  $\omega$  → 4.7841,  $\phi$  → 3.07873, xoff → 0.00441981}
```

```
In[82]:= fitplot = Plot[fit[t], {t, 0, 3}];
Show[{dataplot, fitplot}, ImageSize -> Scaled[0.6]]
```

Out[83]=



Fit Results and Diagnostics

NonlinearModelFit returns a variety of measures for how "good" the fit is. The online help has more details, but four particularly useful ones are shown here.

Best Fit Parameters

```
In[87]:= fit["BestFitParameters"]
```

Out[87]=

```
{A -> 1.43493, ω -> 4.7841, φ -> 3.07873, xoff -> 0.00441981}
```

The best fit parameters are given as a set of replacement rules. You can assign a value to a variable with the `./` notation. For example, to extract the phase you could do

```
In[88]:= phase = φ /. fit["BestFitParameters"]
```

Out[88]=

```
3.07873
```

Similarly, you can calculate the period by using the replacement rule for ω :

```
In[89]:= T = (2 π / ω) /. fit["BestFitParameters"]
```

Out[89]=

```
1.31335
```

Parameter Confidence Interval

This item gives both the uncertainty and the 95% confidence-level interval.

```
In[90]:= fit["ParameterConfidenceIntervalTable"]
```

```
Out[90]=
```

	Estimate	Standard Error	Confidence Interval
A	1.43493	0.00199071	{1.43084, 1.43902 }
ω	4.7841	0.00169737	{4.78062, 4.78759 }
ϕ	3.07873	0.00289789	{3.07277, 3.08469 }
xoff	0.00441981	0.00144711	{0.00144523 , 0.00739438 }

The table entries are available in the raw form, and you can extract the ones you want with the Part `[[]]` notation. For example, the ω values are in the 2nd row, 1st and 2nd columns

```
In[91]:= fit["ParameterConfidenceIntervalTableEntries"]
```

```
Out[91]=
```

```
{ {1.43493, 0.00199071, {1.43084, 1.43902}},
  {4.7841, 0.00169737, {4.78062, 4.78759}}, {3.07873, 0.00289789, {3.07277, 3.08469}},
  {0.00441981, 0.00144711, {0.00144523, 0.00739438}} }
```

```
In[92]:= {best $\omega$ ,  $\delta\omega$ } = fit["ParameterConfidenceIntervalTableEntries"][[2, {1, 2}]]
```

```
Out[92]=
```

```
{4.7841, 0.00169737}
```

Correlation Matrix

The "CorrelationMatrix" tells you about the correlations among the fit parameters. Low values mean the parameters are poorly correlated. This is good -- it means they are mostly independent. On the other hand, high values mean the parameters are highly correlated. This is often a sign of trouble in the proposed fit. For example, if you are trying to fit the parameters A and B in the function

$$f[x] = A \frac{x}{B}$$

then the fitted parameters for A and B will be highly correlated. You could double A and also double B, and still get the same answer.

```
In[95]:= fit["CorrelationMatrix"] // MatrixForm
```

```
Out[95]//MatrixForm=
```

$$\begin{pmatrix} 1. & 0.0313374 & -0.000416606 & 0.125407 \\ 0.0313374 & 1. & -0.872054 & 0.16148 \\ -0.000416606 & -0.872054 & 1. & -0.17947 \\ 0.125407 & 0.16148 & -0.17947 & 1. \end{pmatrix}$$

For example, the 3rd element on the first row (0.000416606) means that changes in the first parameter (A) are quite weakly correlated with changes in the third parameter (ϕ). The 1's along the diagonals mean each parameter is perfectly correlated with itself, as expected. Checking the correlation matrix is a good way to check whether your proposed fitting function really has the independent parameters you thought it did.

Here, the only tricky one is the -0.87 correlation between ω and the phase ϕ . Because there are only two full oscillations in the data, you can (partially) compensate for increasing ϕ by decreasing ω . This shows up clearly in the density plot of χ^2 , where I have fixed A and xoff to approximately the right values, but allowed ω and ϕ to vary. The deep purple diagonal valley at the center shows how you can

increase ω but decrease ϕ without changing χ^2 much. (You can also fiddle with the Manipulate box above to see this.)

RMS Error

The RMS error is the typical amount the fit misses a data point. Here it is about 0.0077, which is about 0.5% of the amplitude. Depending on the specific sensor used, this is probably a plausible uncertainty.

```
In[99]:= RMSE = Sqrt[fit["EstimatedVariance"]]
Out[99]= 0.0077339
```

If you want to compare that to the amplitude, and express it as a percentage, you can write:

```
In[100]:= (100 * Sqrt[fit["EstimatedVariance"]]) / (A /. fit["BestFitParameters"])
Out[100]= 0.538974
```

The resulting fit object has many other “Properties” you can query. Here is a list of all of them:

```
In[101]:= fit["Properties"]
Out[101]= {AdjustedRSquared, AIC, AICc, ANOVA, BestFitAround, BestFitDataAround,
  BestFit, BestFitParameters, BIC, CorrelationMatrix, CovarianceMatrix,
  CurvatureConfidenceRegion, Data, Weights, EstimatedVariance, FitCurvature,
  FitResiduals, Function, TabularFunction, HatDiagonal, MaxIntrinsicCurvature,
  MaxParameterEffectsCurvature, MeanPredictions, MeanPredictionBands,
  ParameterEstimates, ParameterConfidenceRegion, ParameterBias, PredictedResponse,
  Properties, Response, RSquared, SingleDeletionVariances, SinglePredictions,
  SinglePredictionBands, StandardizedResiduals, StudentizedResiduals}
```

Summary of Usage

Import Data

```
In[102]:= SetDirectory[NotebookDirectory[]];
In[103]:= data = Import["nonlinear-modelfit-1-data.txt", "Table"];
  dataplot = ListPlot[data, PlotRange → All,
    PlotStyle → Red, Frame → True, LabelStyle → Larger,
    FrameLabel → {"Time (s)", "Amplitude (V)"}, ImageSize → Scaled[0.5]];
```

Construct Model

```
In[105]:=
x[A_, ω_, φ_, t_, xoff_] := A Cos[ω t + φ] + xoff
```

Run NonlinearModel Fit, giving hints for parameters as needed

```
In[110]:=
fit = NonlinearModelFit[data, x[A, ω, φ, t, xoff],
  {A, {ω, 4}, {φ, π}, xoff}, t]; (* Here we give initial guesses for ω and φ *)
fit["ParameterConfidenceIntervalTable"]
```

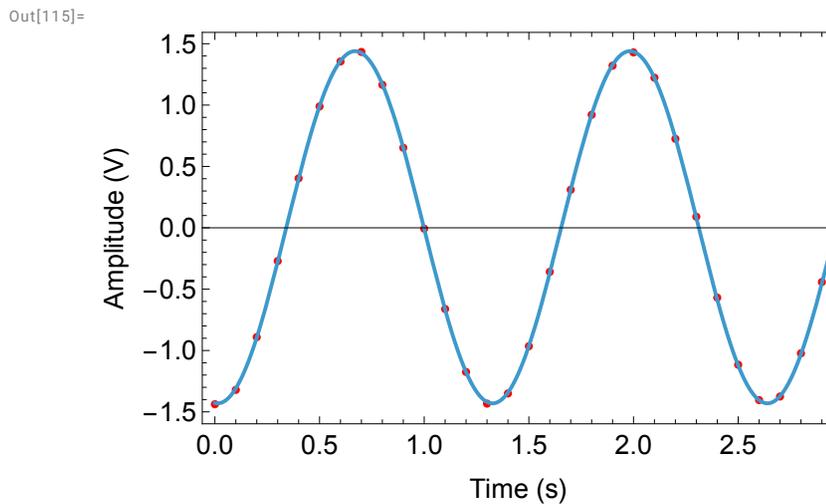
```
Out[111]=
```

	Estimate	Standard Error	Confidence Interval
A	1.43493	0.00199071	{1.43084, 1.43902 }
ω	4.7841	0.00169737	{4.78062, 4.78759 }
ϕ	3.07873	0.00289789	{3.07277, 3.08469 }
xoff	0.00441981	0.00144711	{0.00144523 , 0.00739438 }

```
In[113]:=
RMSE = Sqrt[fit["EstimatedVariance"]] (* Find typical error. *)
```

```
Out[113]=
0.0077339
```

```
In[114]:=
fitplot = Plot[fit[t], {t, 0, 3}];
Show[{dataplot, fitplot}, ImageSize → Scaled[0.6]]
```



Extracting values (with uncertainty) from the fit

```
In[118]:=
phase = φ /. fit["BestFitParameters"]
```

```
Out[118]=
3.07873
```

In[117]:=

```
{best $\omega$ ,  $\delta\omega$ } = fit["ParameterConfidenceIntervalTableEntries"][[2, {1, 2}]]
```

Out[117]=

```
{4.7841, 0.00169737}
```