

## Initialize

### ■ Read in Statistical Add-in packages:

```
Off[General::"spell1"];  
<< "MultivariateStatistics`";  
<< LinearRegression`;  
<< "ErrorBarPlots`";
```

---

## Statistical learning, model selection & the bias/variance dilemma

Earlier we summarized optimal rules for minimizing risk, assuming that we know the distributions of the generative model.

But what if we don't? Consider the regression problem, fitting data that may be a complex function of the input.

The problem in general is how to choose the function that both remembers the relationship between  $\mathbf{x}$  and  $\mathbf{y}$ , and generalizes with new values of  $\mathbf{x}$ . At first one might think that it should be as general as possible to allow all kinds of maps.

For example, if one is fitting a curve, you might wish to use a very high-order polynomial, or a back-prop network with lots of hidden units. There is a drawback, however, to the flexibility afforded by extra degrees of freedom in fitting the data. We can get drastically different fits for different sets of data that are randomly drawn from the same underlying process. The fact that we get different fit parameters (e.g. slope of a regression line) each time means that although we may exactly fit the data each time, we introduce variation between the average fit (over all data sets) and the fits over the ensemble of data sets. We could get around this problem with a huge amount of data, but the problem is that the amount of required data can grow exponentially with the order of the fit--an example of the so-called "curse of dimensionality".

On the other hand, if the function is restrictive, (e.g. straight lines through the origin), then we will get similar fits for different data sets, because all we have to adjust is one parameter--the slope. The problem here, is that the fit is only good if the underlying process is in fact a straight line through the origin. If it isn't a straight line for instance, there will be a fixed error or **bias** that will never go away, no matter how much data we collect. Statisticians refer to the trade-off between simple but biased fits, and complex but data-dependent variation, as the *bias/variance* dilemma.

To sum up, lots of parameter flexibility (or lots of hidden units) has the benefit of fitting anything, but at the cost of sensitivity to variability in the data set--there is *variance* introduced by the fits found over multiple training sets (e.g. of a small fixed size).

A fit with very few parameters is not as sensitive to the inevitable variability in the training set, but can give large constant errors or *bias* if the data do not match the underlying model.

There is no general way of getting around this problem, and neural networks are no exception. We generalize linear regression to non-linear fits using error back-propagation. Because back-propagation models can have lots of hidden layers with many units and weights, they form a class of very flexible approximators and can fit almost any function. But these models can show high variability in their fits from one data set to the next, even when the data comes from the same underlying process. Lots of hidden units can mean low bias, but at a high cost in variance.

## Demonstration of bias/variance for regression

Suppose we have an unknown, underlying generative model given by:  $p(y|x)p(x)$ . From this we obtain a set of samples  $\{x_i, y_i\}, i=1 \dots N$ . We posit some estimator to fit the data:  $y_i \sim f(x_i)$ . In general, there will be some cost assigned to errors in  $f$ 's ability to predict the  $y$ 's. E.g. the expected value of the squared difference between the fits and the true expected value of  $y$ , call it  $\hat{y}$ . We can write this cost as the sum of two terms:

$$E[(f - \hat{y})^2] = (E[f] - \hat{y})^2 + E[(f - E[f])^2]$$

The first term on the right is the bias (squared), and the second term the variance. The bias is the "constant error" which tells us how far off we'll be no matter how much data we have. The second term, the "variance", tells us how much variation we have in the ensemble of fits  $f$  about the average of all the fit,  $E[f]$ .

Let's represent the expectation of  $f$ ,  $E[f]$ , by  $\tilde{f}$ :

$$E[(f - \hat{y})^2] = (\tilde{f} - \hat{y})^2 + E[(f - \tilde{f})^2]$$

A more general formulation of the bias/variance trade-off, for a risk function  $R()$  is:

$$\int R(\alpha_D; x) P(\alpha_D) d\alpha = \int (y - \hat{y}(x))^2 P(y|x) dy + \int (\tilde{f}(x) - f(x, \alpha))^2 P(\alpha|x) d\alpha$$

Use the following link for the notes:

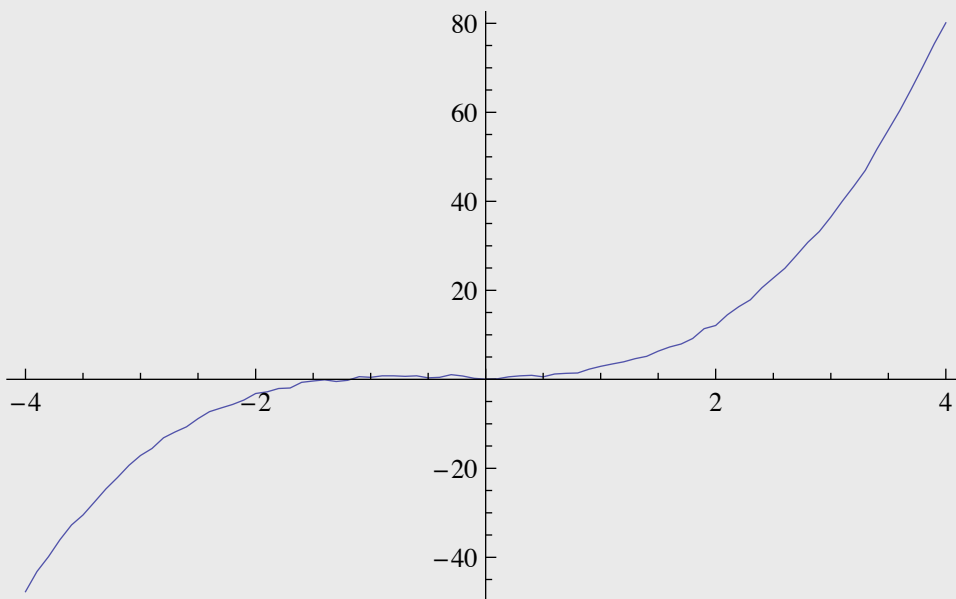
[http://gandalf.psych.umn.edu/~kersten/kersten-lab/courses/Psy5038WF2003/MathematicaNotebooks/Lect\\_27\\_Bias-Variance/biasvarianceNotes.pdf](http://gandalf.psych.umn.edu/~kersten/kersten-lab/courses/Psy5038WF2003/MathematicaNotebooks/Lect_27_Bias-Variance/biasvarianceNotes.pdf)

Let's demonstrate the effects of the bias/variance trade-off by estimating the values of  $\tilde{f}$  and  $\hat{y}$ , given a generative model.

### ■ *Mathematica's regression package*

Go to Help, and find the Linear Regression package. Look up **Regress**. We are going to use Regress as our learning model. We could have used our error-back prop network, or other learning algorithms that produce a set of fit parameters. The principles would be the same.

```
ff[x_, α_] := α.{1, x, x2, x3} + RandomReal[];
α = {0, 0, 1, 1};
xd = Table[{x, ff[x, α]}, {x, -4, 4, 0.1}];
ListPlot[xd, AxesOrigin → {0, 0}, Joined → True]
Regress[xd, {1, x, x2, x3}, x][[1]]
```



	Estimate	SE	TStat	PValue
1	0.501449	0.0490169	10.2301	$4.44089 \times 10^{-16}$
ParameterTable → $x$	0.0436236	0.034955	1.24799	0.215816
$x^2$	1.01042	0.00668392	151.172	0.
$x^3$	0.99746	0.00325677	306.273	0.

## Learning from one data set

### ■ True model


```
ff[x_, α_] := α.{1, x, x^2, x^3};
```

### ■ Generative data process: true plus some noise

```
noise = 15;  
ffn[x_, α_] := ff[x, α] + 1.5 * RandomReal[{-noise, noise}];
```

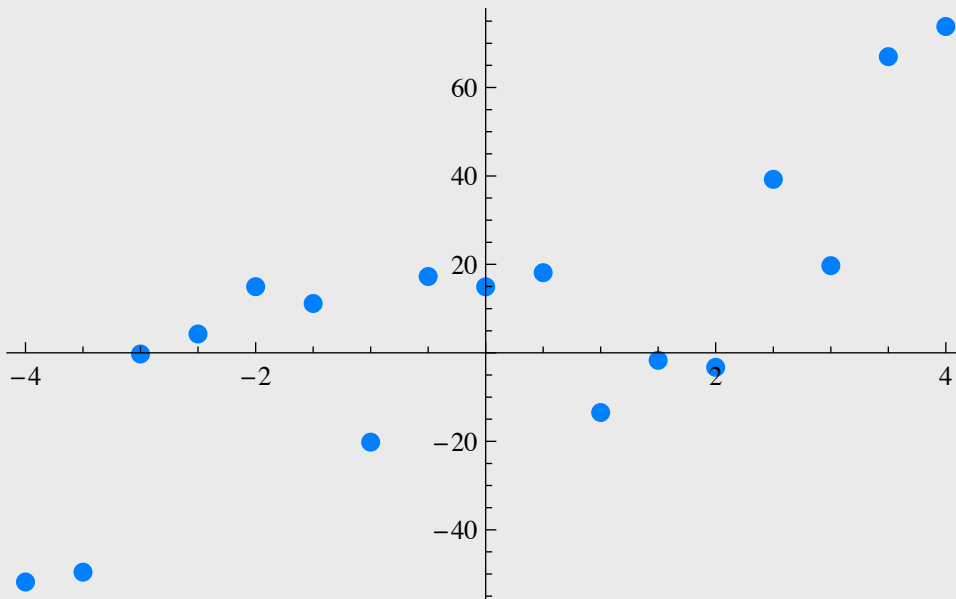
### ■ Choose domain and true model parameters. Calculate a set of samples from true model ffp, evaluated at xp.

```
xp = Table[x, {x, -4, 4, 0.5}];  
α = {0, 0, 1, 1};  
ffp = (ff[#1, α] &) /@ xp;  
gffp = ListPlot[Transpose[{xp, ffp}],  
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True]
```



- Run one experiment to collect y values for data process:

```
y = (ffn[#1,  $\alpha$ ] &) /@ xp;
gy = ListPlot[Transpose[{xp, y}], PlotStyle -> {PointSize[0.02], Hue[0.6]}]
```

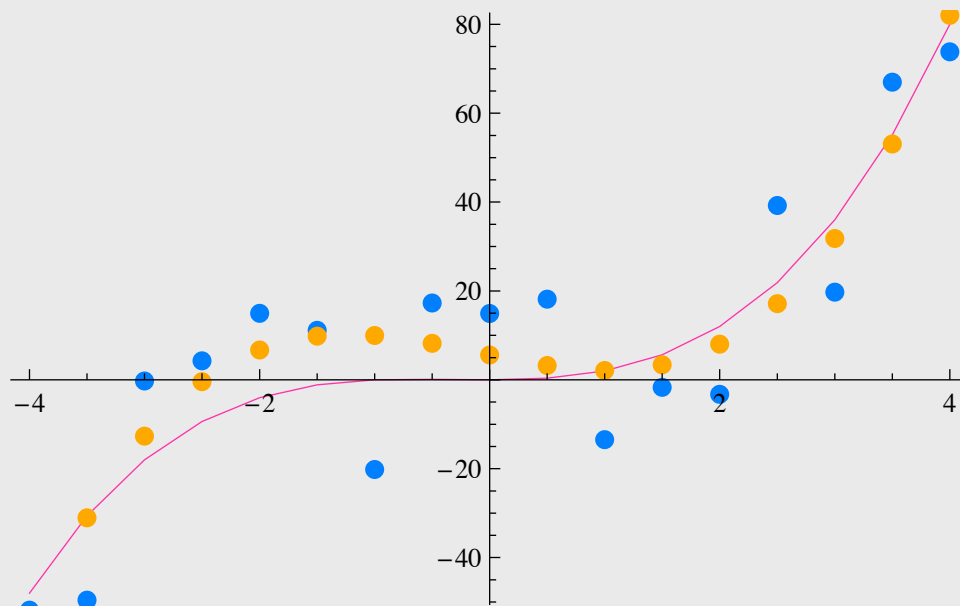


- Estimate model parameters using polynomial regression. Return model parameters, and predicted responses, fsquiggle

```
 $\alpha$ D = Regress[Transpose[{xp, y}], {1, x, x2, x3}, x,
  RegressionReport -> BestFitParameters][[1, 2]]
fsquiggle =
  Regress[Transpose[{xp, y}], {1, x, x2, x3}, x,
  RegressionReport -> PredictedResponse][[1, 2]];
gfsquiggle = ListPlot[Transpose[{xp, fsquiggle}],
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];
```

```
{5.59987, -5.34946, 0.441812, 1.41805}
```

```
Show[gffp, gy, gfsquiggle]
```



Repeat the above, and notice how the model parameters and the fit changes. Try changing the basis functions used for fitting.

## Comparing learning from multiple data sets

We'd like some idea of how learning generalizes depending on model complexity

### ■ The "right" model

First, assume by some incredibly lucky guess, we've chosen the right model  $\{x^2, x^3\}$ , and want to find the parameters.


```
y = ffn[#1, α] & /@ xp;
fsquiggle =
  Regress[Transpose[{xp, y}], {x^2, x^3}, x,
    RegressionReport → PredictedResponse][[1, 2]];

```

```

gffp = ListPlot[Transpose[{xp, ffp}],
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True];
gsmooth = Plot[Fit[Transpose[{xp, y}], {x^2, x^3}, x] /. x -> x2,
  {x2, -4, 4}];
gfsquiggle = ListPlot[Transpose[{xp, fsquiggle}],
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];
Show[gffp, gy, gfsquiggle, gsmooth]

```



### ■ A simple, but wrong model

But now suppose that we are trying to fit the data with an inappropriate model. In particular, suppose that it is weak, say a linear model:

```


y = ffn[#1, α] & /@ xp;
fsquiggle =
  Regress[Transpose[{xp, y}], {1, x}, x,
  RegressionReport -> PredictedResponse][[1, 2]];

```

```

gffp = ListPlot[Transpose[{xp, ffp}],
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True];
gy = ListPlot[Transpose[{xp, y}], PlotStyle -> {PointSize[0.02], Hue[0.6]}];
gsmooth = Plot[Fit[Transpose[{xp, y}], {1, x}, x] /. x -> x^2, {x2, -4, 4}];
gfsquiggle = ListPlot[Transpose[{xp, fsquiggle}],
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];
Show[gffp, gy, gfsquiggle, gsmooth]

```



The bias is the squared difference between the average y values (blue) and the model fits (orange). The variance is the squared difference between the predicted responses (blue) and the true (red line). So we can see that the bias is high. The variance (represented by the square root of the variance, black points in the graph) is not zero. But how does it compare with a model with lots of parameters?

### ■ A complex model, with too many parameters

Now let's try over-fitting. (Analogous to having lots of hidden units and/or layers in a non-linear feedforward network).

```

y = ffn[#1, α] & /@ xp;
fsquiggle =
  Regress[Transpose[{xp, y}],
    {1, x, x^2, x^3, x^4, x^5, x^6, x^7, x^8, x^9, x^10, x^11, x^12},
    x, RegressionReport -> PredictedResponse][[1, 2]];


```



```

gffp = ListPlot[Transpose[{xp, ffp}],
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True];
gsmooth =
  Plot[
    Fit[Transpose[{xp, y}], {1, x, x^2, x^3, x^4, x^5, x^6, x^7,
      x^8, x^9, x^10, x^11, x^12}, x] /. x -> x2, {x2, -4, 4}];
gy = ListPlot[Transpose[{xp, y}], PlotStyle -> {PointSize[0.02], Hue[0.6]}];
gfsquiggle = ListPlot[Transpose[{xp, fsquiggle}],
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];
Show[gffp, gy, gfsquiggle, gsmooth]

```



Note how the bias (discrepancy between the orange and blue) is lower than with too few parameters. But we have higher variance than with the "right model" family.

## Comparing learning from multiple data sets

To get a quantitative idea of how learning generalizes, we are going to run a bunch of experiments to get a bunch of data sets. From these we get multiple fits. Thus we can get estimates of the average value of  $y$ , and the variance of the model predictions (with respect to the true model) over multiple fits. And calculate estimates of the second (variance) and third (bias) terms below.

$$\int R(\alpha_D; x) P(\alpha_D) d\alpha = \int (y - \hat{y}(x))^2 P(y|x) dy + \int (\tilde{f}(x) - f(x, \alpha))^2 P(\alpha_D) d\alpha + \dots$$

## ■ The right model

First, assume by some incredibly lucky guess, we've chosen the right model  $\{x^2, x^3\}$ , and want to find the parameters.

```

iter = 500;
ysum = Table[0, {i, 1, Length[xp]}];
fsquigglesum = Table[0, {i, 1, Length[xp]}];
y = Table[0, {i, 1, iter}];
fsquiggle = Table[0, {i, 1, iter}];
var = Table[0, {i, 1, Length[xp]}];
For[i = 1, i ≤ iter, i++,
  y[[i]] = ffn[#1, α] & /@ xp;
  fsquiggle[[i]] =
    Regress[Transpose[{xp, y[[i]]}], {x^2, x^3}, x,
      RegressionReport → PredictedResponse][[1, 2]];
  ysum = ysum + y[[i]];
  fsquigglesum = fsquigglesum + fsquiggle[[i]];
];

yhat = ysum / iter;
fsquigglemean = fsquigglesum / iter;

For[i = 1, i ≤ iter, i++,
  var = var + (fsquiggle[[i]] - fsquigglemean)^2;
];

stdright = Sqrt[var / iter]
biasright = Abs[(fsquigglemean - yhat)]

```

```

{9.67287, 6.69123, 4.56277, 3.4552, 3.36697, 3.80177, 4.28539, 4.61684,
 4.7319, 4.61765, 4.29274, 3.83196, 3.4557, 3.64267, 4.84061, 7.03628, 10.082}


```

```

{0.365582, 0.336445, 0.122321, 0.82054, 0.442705, 0.165885, 0.397393, 0.3387, 0.0811485,
 0.734664, 0.0786217, 0.607147, 1.20684, 0.265881, 0.214032, 0.599569, 0.33919}

```

```
gffp = ListPlot[Transpose[{xp, ffp}],  
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True];  
gy = ListPlot[Transpose[{xp, yhat}],  
  PlotStyle -> {PointSize[0.02], Hue[0.6]}];  
gfsquiggle = ListPlot[Transpose[{xp, fsquigglemean}],  
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];  
Show[gffp, gy, gfsquiggle]
```



### ■ A simple, but wrong model

But now suppose that we are trying to fit the data with an inappropriate model. In particular, suppose that it is weak, say a linear model:

```

iter = 500;
ysum = Table[0, {i, 1, Length[xp]}];
y = Table[0, {i, 1, iter}];
fsquiggle = Table[0, {i, 1, iter}];
fsquigglesum = Table[0, {i, 1, Length[xp]}];
var = Table[0, {i, 1, Length[xp]}];
For[i = 1, i ≤ iter, i++,
  y[[i]] = ffN[#1, α] & /@ xp;
  fsquiggle[[i]] =
    Regress[Transpose[{xp, y[[i]]}], {1, x}, x,
      RegressionReport → PredictedResponse][[1, 2]];
  ysum = ysum + y[[i]];
  fsquigglesum = fsquigglesum + fsquiggle[[i]];
];
yhat = ysum / iter;
fsquigglemean = fsquigglesum / iter;

For[i = 1, i ≤ iter, i++,
  var = var + (fsquiggle[[i]] - fsquigglemean) ^ 2;
];

stdsimple = Sqrt[var / iter]
biassimple = Abs[(fsquigglemean - yhat)]

```

```

{6.23651, 5.68331, 5.15491, 4.65975, 4.2096, 3.82037, 3.5124, 3.30845,
 3.2283, 3.28102, 3.46056, 3.74873, 4.12283, 4.56175, 5.04862, 5.57088, 6.11948}


```

```

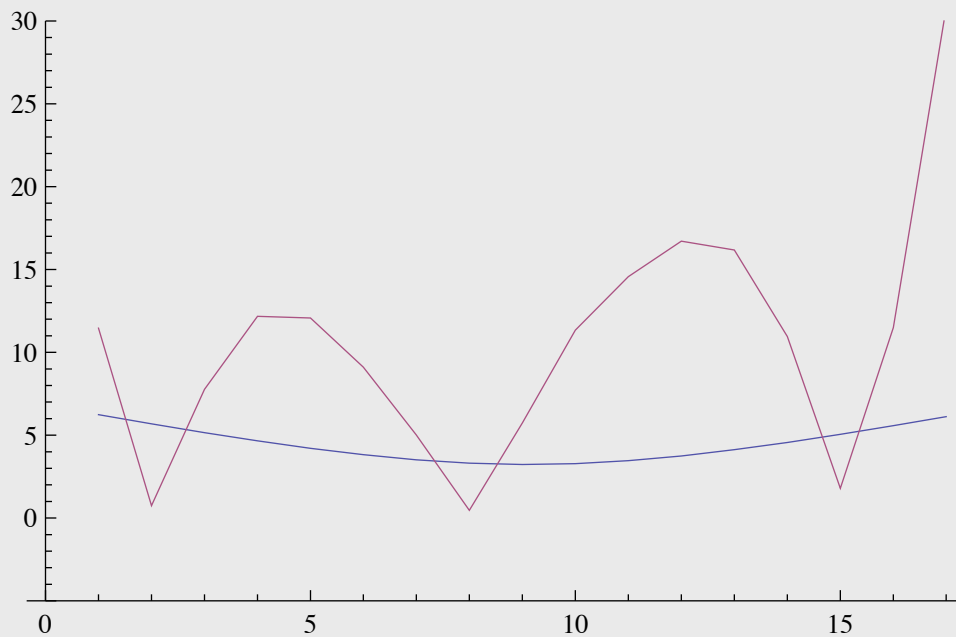
{11.4579, 0.750538, 7.77006, 12.1698, 12.0739, 9.10867, 5.0109, 0.462797,
 5.73876, 11.3372, 14.5717, 16.7097, 16.1792, 10.9587, 1.79657, 11.4887, 30.8399}

```

```
gffp = ListPlot[Transpose[{xp, ffp}],  
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True];  
gy = ListPlot[Transpose[{xp, yhat}],  
  PlotStyle -> {PointSize[0.02], Hue[0.6]}];  
gfsquiggle = ListPlot[Transpose[{xp, fsquigglemean}],  
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];  
Show[gffp, gy, gfsquiggle]
```



```
ListPlot[{stdsimple, biassimple}, Joined -> {True, True},
PlotRange -> {-5, 30}, AxesOrigin -> {0, -5}]
```



The bias is the squared difference between the average y values (blue) and the model fits (orange). The variance is the squared difference between the predicted responses (blue) and the true (red line). So we can see that the bias is high. The variance (represented by the square root of the variance, black points in the graph) is not zero. But how does it compare with a model with lots of parameters?

### ■ A complex model, with too many parameters

Now let's try over-fitting. (Analogous to having lots of hidden units and/or layers in a non-linear feedforward network).

```

iter = 500;
ysum = Table[0, {i, 1, Length[xp]}];
y = Table[0, {i, 1, iter}];
fsquiggle = Table[0, {i, 1, iter}];
fsquigglesum = Table[0, {i, 1, Length[xp]}];
var = Table[0, {i, 1, Length[xp]}];
For[i = 1, i ≤ iter, i++,
    y[[i]] = ffN[#1, α] & /@ xp;
    fsquiggle[[i]] =
    Regress[Transpose[{xp, y[[i]]}],
      {1, x, x^2, x^3, x^4, x^5, x^6, x^7, x^8, x^9, x^10, x^11, x^12},
      x, RegressionReport → PredictedResponse][[1, 2]];
    ysum = ysum + y[[i]];
    fsquigglesum = fsquigglesum + fsquiggle[[i]];
];
yhat = ysum / iter;
fsquigglemean = fsquigglesum / iter;

For[i = 1, i ≤ iter, i++,
    var = var + (fsquiggle[[i]] - fsquigglemean)^2;
];

stdcomplex = Sqrt[var / iter]
biascomplex = Abs[(fsquigglemean - yhat)]

```

```

{13.1502, 13.3392, 13.0614, 11.7048, 10.2144, 9.62229, 10.1483, 9.02625,
 9.34861, 9.21209, 10.0991, 10.6289, 10.7709, 12.4631, 13.1423, 12.9494, 12.8427}


```

```

{0.00348965, 0.0369638, 0.171234, 0.447006, 0.696128, 0.60354, 0.189878, 0.0282221, 0.232563,
 0.500542, 0.264381, 0.243366, 0.480492, 0.354719, 0.144665, 0.0323508, 0.00312415}

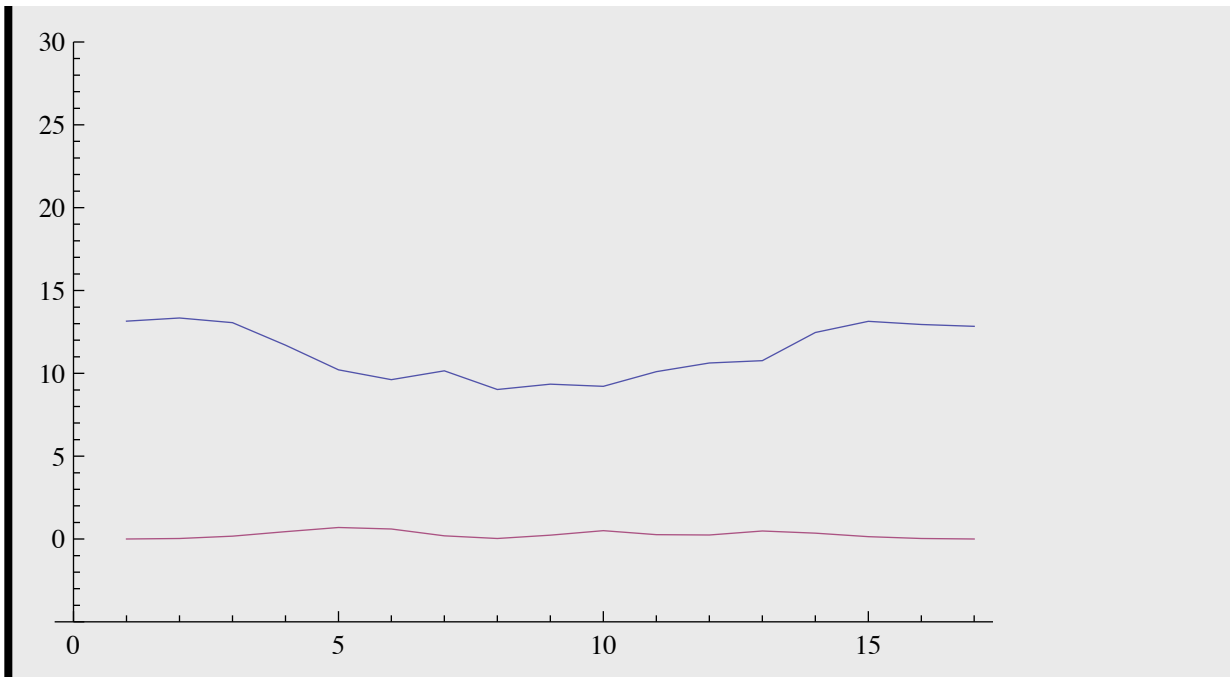
```

```
gffp = ListPlot[Transpose[{xp, ffp}],  
  PlotStyle -> {PointSize[0.02], Hue[0.9]}, Joined -> True];  
gy = ListPlot[Transpose[{xp, yhat}],  
  PlotStyle -> {PointSize[0.02], Hue[0.6]}];  
gfsquiggle = ListPlot[Transpose[{xp, fsquigglemean}],  
  PlotStyle -> {PointSize[0.02], Hue[0.1]}];  
Show[gffp, gy, gfsquiggle]
```





```
ListPlot[{stdcomplex, biascomplex}, Joined -> {True}, PlotRange -> {-5, 30},
AxesOrigin -> {0, -5}]
```



Note how the bias (discrepancy between the orange and blue) is lower than with too few parameters. But we have higher variance than with the "right model" family.

## Model selection

### Bayesian model selection

MacKay (1992)

### Cross-validation

#### Beyond linear separation, beyond sigmoidal kernels

Density estimation vs. classification

Map the data (through some non-linear mapping, e.g. polynomial) to a higher-dimensional space to find the optimal hyperplane separating the data. But what is "optimal". Require good generalization, small VC dimension. Construct the hyperplane on a small number of support vectors, then the generalization ability will be high.

Main source: Vapnik (1995)

<http://svm.first.gmd.de/>

Demo links:

<http://svm.dcs.rhbnc.ac.uk/pagesnew/GPat.shtml>

---

## References

Duda, R. O., & Hart, P. E. (1973). Pattern classification and scene analysis. New York.: John Wiley & Sons.

Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern classification* (2nd ed.). New York: Wiley.

(Amazon.com)

Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1), 1-58.

Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13(2), 1-9.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian Inference. *Annual Review of Psychology*, 55.

MacKay, D. J. C. (1992). Bayesian interpolation. *Neural Computation*, 4(3), 415-447.

Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge, UK: Cambridge University Press.

Vapnik, V. N. (1995). *The nature of statistical learning*. New York: Springer-Verlag.

<http://neuron.eng.wayne.edu/software.html>

© 1998, 2001, 2003, 2005, 2007 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.