

Initialize

■ Read in Statistical Add-in packages:

```
In[1]:= Off[General::spell1];  
Needs["ErrorBarPlots`"];  
Needs["MultivariateStatistics`"];
```

Review Discriminant functions (Lecture 11)

Let's build our geometric intuitions of what a simple perceptron unit does by viewing it from a more formal point of view. Perceptron learning is an example of nonparametric statistical learning, because it doesn't require knowledge of the underlying probability distributions generating the data (such distributions are characterized by a relatively small number of "parameters", such as the mean and variance of a Gaussian distribution). Of course, how well it does will depend on the generative structure of the data. Much of the material below is covered in Duda and Hart (1978).

Linear discriminant functions: Two category case

A discriminant function, $g(\mathbf{x})$ divides input space into two category regions depending on whether $g(\mathbf{x}) > 0$ or $g(\mathbf{x}) < 0$. (We've switched notation, $\mathbf{x} = \mathbf{f}$). The linear case corresponds to the simple perceptron unit we studied earlier:

$$g(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x} + w_0 \quad (1)$$

where \mathbf{w} is the weight vector and w_0 is the (scalar) threshold (sometimes called bias, although this "bias" has nothing to do with statistical "bias").

Discriminant functions can be generalized, for example to quadratic decision surfaces:

$$g(\mathbf{x}) = w_0 + \sum_{i=1} w_i x_i + \sum_{i=1} \sum_{j=1} w_{ij} x_i x_j \quad (2)$$

where $\mathbf{x} = \{x_1, x_2, x_3, \dots\}$. We've seen how $g(\mathbf{x})=0$ defines a decision surface which in the linear case is a hyperplane.

Suppose \mathbf{x}_1 and \mathbf{x}_2 are vectors, with endpoints sitting on the hyperplane, then their difference is a vector lying in the hyperplane

$$\begin{aligned} \mathbf{w} \cdot \mathbf{x}_1 + w_0 &= \mathbf{w} \cdot \mathbf{x}_2 + w_0 \\ \mathbf{w} \cdot (\mathbf{x}_1 - \mathbf{x}_2) &= 0 \end{aligned} \quad (3)$$

so the weight vector \mathbf{w} is normal to any vector lying in the hyperplane. Thus \mathbf{w} determines how the plane is oriented. The normal vector \mathbf{w} points into the region for which $g(\mathbf{x})>0$, and $-\mathbf{w}$ points into the region for which $g(\mathbf{x})<0$.

Let \mathbf{x} be a point on the hyperplane. If we project \mathbf{x} onto the normalized weight vector $\mathbf{x} \cdot \mathbf{w} / |\mathbf{w}|$, we have the normal distance of the hyperplane from the origin equal to:

$$\mathbf{w} \cdot \mathbf{x} / |\mathbf{w}| = -w_0 / |\mathbf{w}| \quad (4)$$

Thus, the threshold determines the position of the hyperplane.

One can also show that the normal distance of \mathbf{x} to the hyperplane is given by:

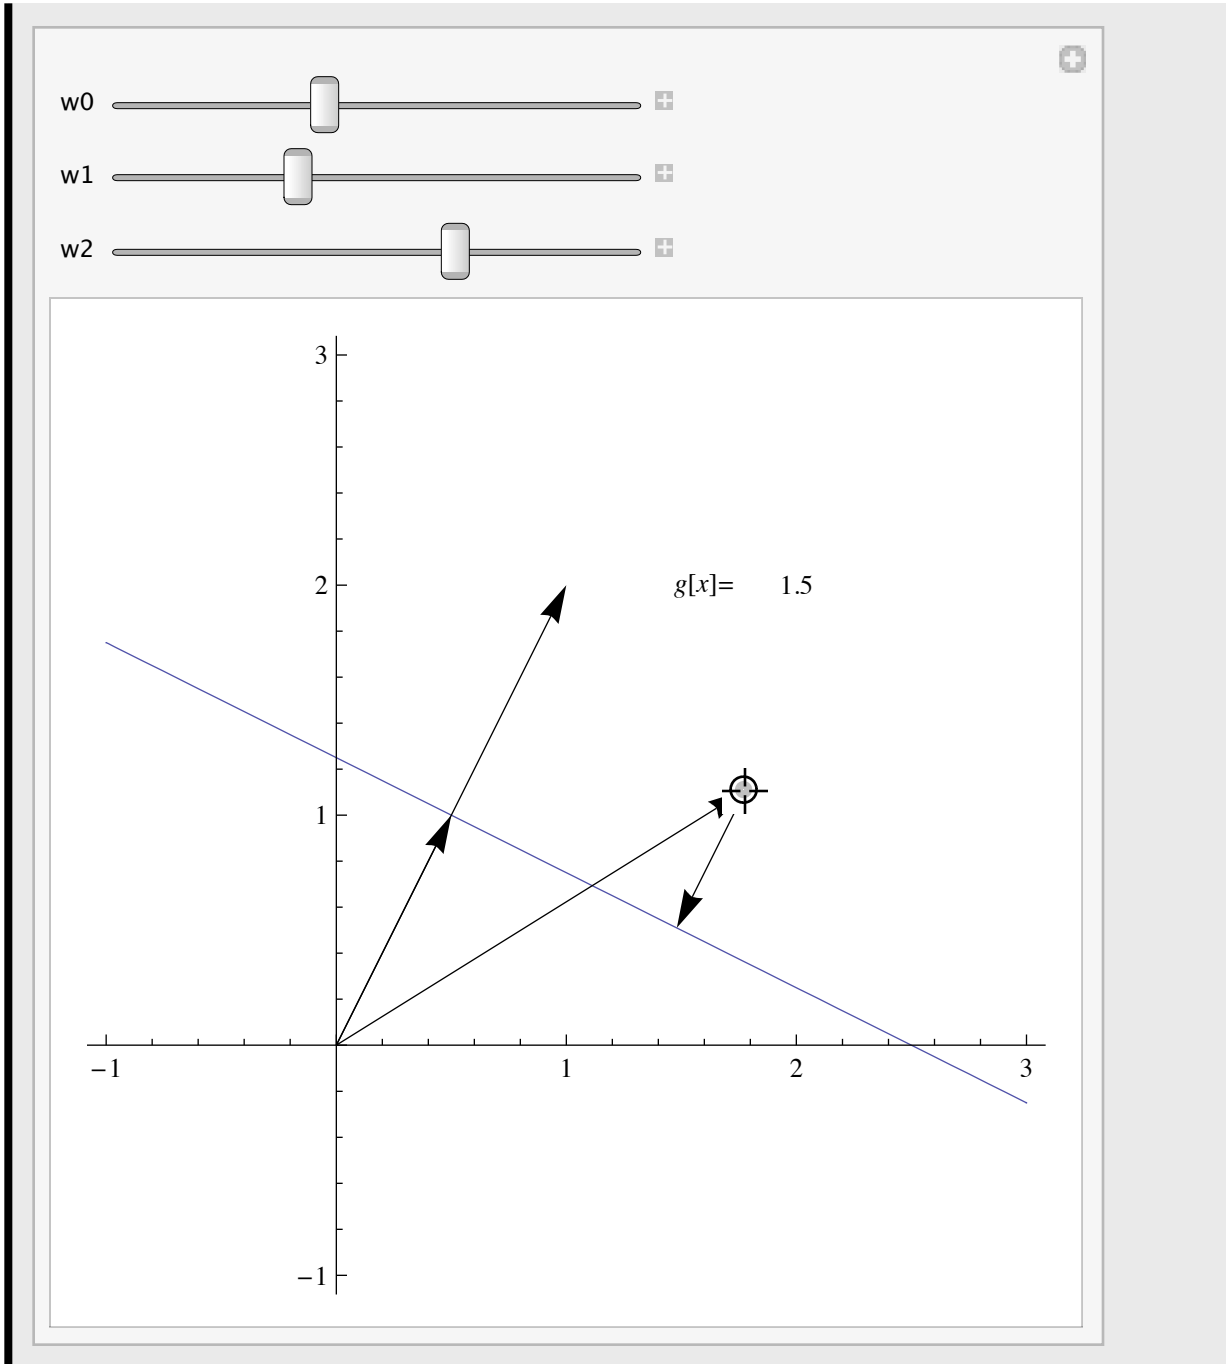
$$g(\mathbf{x}) / |\mathbf{w}| \quad (5)$$

So we've seen that: 1) discriminant function divides the input space by a hyperplane decision surface; 2) The orientation of the surface is determined by the weight vector \mathbf{w} ; 3) the location is determined by the threshold w_0 ; 4) the discriminant function gives a measure of how far an input vector is from the hyperplane.

The figure summarizes the basic properties of the linear discriminant.

```
In[675]:= Manipulate[
  (* x0={2,1};*)
  w = {w1, w2};
  wn = w / Norm[w];
  g[x_] := {w1, w2} . x + w0;
  gg = Plot[Tooltip[x2 /. Solve[{w1, w2} . {x1, x2} + w0 == 0, x2],
    "discriminant"], {x1, -1, 3}];
  ggg = Graphics[g[Dynamic[MousePosition["Graphics"]]]];
  Show[{gg, Graphics[Inset["g[x]=", {1.6, 2}]],
    Graphics[Inset[ToString[g[x0]], {2, 2}]],
    Graphics[{Tooltip[Arrow[{0, 0}, w], "w"],
      Tooltip[Arrow[{0, 0}, (-w0 / Norm[w]) * wn], "-w0 / |w|"],
      Tooltip[Arrow[{0, 0}, x0], "x"],
      Tooltip[Arrow[{x0, x0 - wn * g[x0] / Norm[w]}, "g(x) / |w|"]]}],
    PlotRange -> {{-1, 3}, {-1, 3}}, AxesOrigin -> {0, 0}, Axes -> True,
    AspectRatio -> 1, {{w0, -2.5}, -6, 3}, {{w1, 1}, 0, 3}, {{w2, 2}, 0, 3},
    {{x0, {2, 1}}, Locator}]
```

Out[675]=



Task-dependent Dimensionality reduction

Fisher's linear "discriminant"

The idea is that the original input space may be impractically huge, but if we can find a subspace (hyperplane) that preserves the distinctions between categories as well as possible, we can make our decisions in smaller space. We will derive the Fisher linear "discriminant".

This is closely related to the psychology idea of finding "distinctive" features. E.g. consider bird identification. If I want to discriminate cardinals from other birds in my backyard, I can make use of the fact that (male) cardinals may be the only birds that are red. So even tho' the image of a bird can have lots of dimensions, if I project the image on to the "red" axis, I can do fairly well with just one number. How about male vs. female human faces?

■ Generative model: two nearby gaussian classes

Define two bivariate base distributions

```
In[4]:= (ar = {{1, 0.99}, {0.99, 1}};
ndista = MultinormalDistribution[{0, -1}, ar];)
(br = {{1, .9}, {.9, 2}};
ndistb = MultinormalDistribution[{0, 1}, br];)
```

Find the expression for the probability distribution function of ndista

```
pdf = PDF[ndista, {x1, x2}]
```

$$1.12822 e^{\frac{1}{2}(-x_1(50.2513 x_1 - 49.7487(x_2 + 1)) - (x_2 + 1)(50.2513(x_2 + 1) - 49.7487 x_1))}$$

Use Mean[] and CovarianceMatrix[ndista] to verify the population mean and the covariance matrix of ndistb

```
In[6]:= Mean[ndistb]
```

```
Out[6]= {0, 1}
```

```
In[7]:= Covariance[ndista]
```


```
Out[7]=  $\begin{pmatrix} 1 & 0.99 \\ 0.99 & 1 \end{pmatrix}$ 
```

Try different covariant matrices. Should they be symmetric? Constraints on the determinant of ar, br?

Make a contour plot of the PDF ndista


```
In[664]:= pdfa = PDF[ndista, {x1, x2}];
ContourPlot[pdfa, {x1, -3, 3}, {x2, -3, 3}, PlotPoints -> 64,
PlotRange -> All, ImageSize -> Small]
```

Out[665]=



```
In[10]:= nsamples = 500;  
a = Table[Random[ndista], {nsamples}];  
ga = ListPlot[a, PlotRange → {{-8, 8}, {-8, 8}}, AspectRatio → 1,  
PlotStyle → Hue[0.2`]];  
b = Table[Random[ndistb], {nsamples}];  
gb = ListPlot[b, PlotRange → {{-8, 8}, {-8, 8}}, AspectRatio → 1,  
PlotStyle → Hue[0.6`]];  
Show[ga, gb]
```

Out[15]=



Use `Mean[]` to find the *sample mean* of `b`. What is the *sample covariance* of `b`?

```
In[16]:= Mean[b]
```

```
Out[16]:= {-0.0164449, 1.00993}
```

```
In[17]:= Covariance[b]
```

```
Out[17]=  $\begin{pmatrix} 1.06974 & 0.998697 \\ 0.998697 & 2.09141 \end{pmatrix}$ 
```

■ Try out different projections of the data by varying the slope (m) of a projection line

```
Clear[x, y, n1, n2];  
{x, y} . {n1, n2}  
Map[#1 * {n1, n2} &, {x, y} . {n1, n2}]
```

We'll use the Map[] function to calculate the projection of a data point (x,y) onto unit normal (n1, n2) to produce a vector in the direction of the unit vector.

```
Out[672]= {n1 x + n2 y}
```

```
Out[673]= (n1 (n1 x + n2 y) n2 (n1 x + n2 y))
```

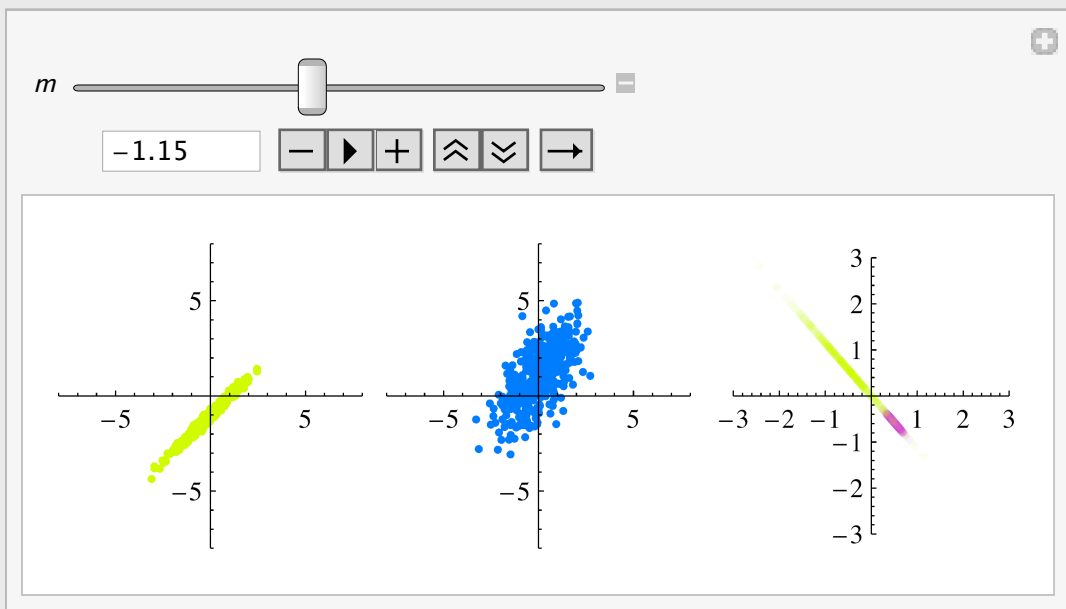


```

In[28]:= Manipulate[
  wnvec = {1, m} / Sqrt[1 + m^2];
  aproj = (#1 wnvec &) /@ (a.wnvec);
  bproj = (#1 wnvec &) /@ (b.wnvec);
  gbproj = ListPlot[{aproj, bproj}, AspectRatio -> 1,
    PlotStyle -> {{PointSize[.03], Hue[0.8], Opacity[.05]},
      {PointSize[.03], Hue[0.2], Opacity[.05]}}],
    PlotRange -> {{-3, 3}, {-3, 3}}];
  GraphicsRow[{ga, gb, gbproj}], {{m, 2/3}, -10, 10}]

```

Out[28]=



By trial and error, find a value of m that separates the classes well along the projection line

```

In[30]:= wnvec = {1, -1.15} / Sqrt[1 + 1.15^2]

```

```

Out[30]:= {0.656179, -0.754606}

```

Calculate the "signal-to-noise" ratio along the projection line: difference between the means divided by the square root of the product of the standard deviations along the line

```

Mean[aproj]

```

```

{0.46178, -0.307853}

```

Theory for simple 2-class case

(see Duda and Hart for general case)

A measure of the separation between the projections is the difference between the means along a projection line in the direction w :

$$| w \cdot (m_a - m_b) |$$

and

$$m_a = \frac{1}{N} \sum_{i=1}^N x_i, \text{ summed over the } N \text{ } x \text{'s from class a}$$

$$m_b = \frac{1}{M} \sum x_i, \text{ summed over the } M \text{ } x \text{'s from class b}$$
(6)

Suppose w (**wnvec**) is the unknown unit vector perpendicular to the discriminant line.

In our case above, the vector difference between the means is:

Mean [a] - Mean [b]

{-0.0146813, -2.03368}

and the difference between the means projected onto a discriminant line is:

wnvec . (Mean [a] - Mean [b])

1.11587

To improve separation, we can't just scale w , because the noise scales too.

We'd like the difference between the means to be large relative to the variation for each class. We can define a measure of the scatter for the projected samples in say class a ($a==1$), by:

$$\sum_{y \in \text{class a}} (y - \tilde{m}_a)^2$$
(7)

where \tilde{m}_a is the sample mean of the points from class a projected onto line $y=w \cdot x$

Or in terms of the Mathematic example:

Apply [Plus, aproj - wnvec . Mean [a]] ;

The total scatter S is defined by the sum of the scatters for both classes (a and b).

$$S = \sum_{y \in \text{class a}} (y - \tilde{m}_a)^2 + \sum_{y \in \text{class b}} (y - \tilde{m}_b)^2$$

If we divide the above number by the total number of points, we have an estimate of the variance of the combined data along the projected axis.

We now have the basic ingredients behind intuition for the Fisher linear discriminant. We'd like to find that \mathbf{w} for which J :

$$J(\mathbf{w}) = \frac{|\tilde{\mathbf{m}}_a - \tilde{\mathbf{m}}_b|^2}{S} = \frac{|\mathbf{w} \cdot (\mathbf{m}_a - \mathbf{m}_b)|^2}{S}$$

is biggest. We want to maximize the difference between the projected class means, while minimizing the dispersion of the data on the projected line.

One can show that $S = \mathbf{w}^T \cdot \mathbf{S}_W \cdot \mathbf{w}$, where

\mathbf{S}_W is measure of within-class variation called the *within-class scatter matrix*:

$$\mathbf{S}_W = \sum_{i=1}^2 \sum_{\mathbf{x} \in \text{Class } i} (\mathbf{x} - \mathbf{m}_i) (\mathbf{x} - \mathbf{m}_i)^T \quad (8)$$

For the numerator, a measure of between class variation is the *between-class scatter matrix*:

$$\mathbf{S}_B = (\mathbf{m}_1 - \mathbf{m}_2) \cdot (\mathbf{m}_1 - \mathbf{m}_2)^T \quad (9)$$

and the difference between the projected means can be show to be:

$$|\tilde{\mathbf{m}}_a - \tilde{\mathbf{m}}_b|^2 = \mathbf{w}^T \cdot \mathbf{S}_B \cdot \mathbf{w}$$

Find \mathbf{w} (corresponding to slope) to maximize the cost function

$$J(\mathbf{w}) = \frac{\mathbf{w}^T \cdot \mathbf{S}_B \cdot \mathbf{w}}{\mathbf{w}^T \cdot \mathbf{S}_W \cdot \mathbf{w}} \quad (10)$$

Answer:

$$\mathbf{w} = \mathbf{S}_W^{-1} \cdot (\mathbf{m}_a - \mathbf{m}_b) \quad (11)$$

■ Demo: Finding Fisher's linear discriminant


```
normalize[x_] := x / Sqrt[x.x];
```

```
ma = Mean[a];
mb = Mean[b];
```

```
Sa = Sum[Outer[Times, a[[i]] - ma, a[[i]] - ma], {i, 1, nsamples}];
Sb = Sum[Outer[Times, b[[i]] - mb, b[[i]] - mb], {i, 1, nsamples}];
Sw = Sa + Sb;
wldf = normalize[Inverse[Sw].(ma - mb)]
```

```
{0.674727, -0.738067}
```


```
aproj = (#1 wldf &) /@ (a.wldf);  
gaproj = ListPlot[aproj, AspectRatio -> 1, PlotStyle -> Hue[0.3`]];  
bproj = (#1 wldf &) /@ (b.wldf);  
gbproj = ListPlot[bproj, AspectRatio -> 1, PlotStyle -> Hue[0.7`]];  
Show[ga, gb, gaproj, gbproj]
```



We started off with a 2-dimensional input problem and turned it into a 1-D problem. For the n-dimensional case, see Duda and Hart.

■ Compare with the principal component axes

```
c = Join[a, b];  
ListPlot[c, PlotRange → {{-8, 8}, {-8, 8}}]
```



```
g1 = ListPlot[c, PlotRange → {{-8, 8}, {-8, 8}}, AspectRatio → 1,  
  DisplayFunction → Identity];
```

```
auto = Covariance[c]  
eigvalues = Eigenvalues[auto]  
eigauto = Eigenvectors[auto]
```


```
(1.01937  0.943791 )  
(0.943791 2.51208 )
```

```
{2.96897, 0.562481}
```

```
(-0.435724 -0.90008 )  
(-0.90008  0.435724 )
```

```
gPCA = Plot[{eigauto[[1,2]]/eigauto[[1,1]] x,  
            eigauto[[2,2]]/eigauto[[2,1]] x},  
            {x,-4,4}, AspectRatio->1,  
            DisplayFunction->Identity,  
            PlotStyle->{RGBColor[.2,0,1]}];
```

```
Show[g1, gPCA, DisplayFunction -> $DisplayFunction]
```



How does the principal component (biggest variance) compare with the Fisher discriminant line?

Kalman filter

■ Blackboard notes

1D example

```
In[693]:= Clear[y, k, c, w]
```


```
In[694]:= sw = 1.0;  
ndistw = NormalDistribution[0, sw];  
w := RandomReal[ndistw];  
  
sv0 = 2.0;  
ndistv = NormalDistribution[0, sv0];  
v := RandomReal[ndistv];
```

■ Generate synthetic data

```
In[700]:= iter = 20;  
c = 1;  
y0 = 0.0;
```

```
In[703]:= y[k_] := k + c + w
ys = NestList[y, y0, iter - 1];
x = ys + RandomReal[ndistv, iter];
gxys = ListPlot[{x, ys}, Joined -> {False, True},
  PlotRange -> {-5, 1.2 * iter}]
```

Out[706]=



■ Initialize arrays

```
In[821]:= st = Table[0, {iter}];
yh = Table[0, {iter}];
K = Table[0, {iter}];
stp = Table[0, {iter}];
yhp = Table[0, {iter}];
```

■ Initial estimates

```
yhp[[1]] = 5;
stp[[1]] = 1;

t = 1;

sv = 2.0 * sv0;
```

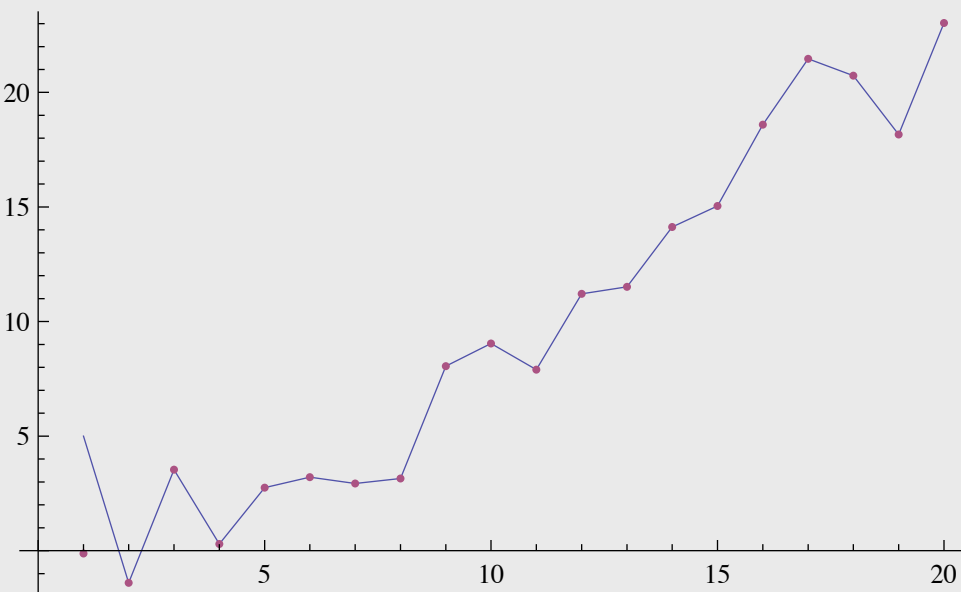

■ Time update--predict the state and variance ahead

```
In[1006]:= yh[[t + 1]] = yhp[[t]] + c;  
st[[t + 1]] = stp[[t]] + sw^2;
```

■ Measurement update--correction

```
In[1008]:= K[[t + 1]] = st[[t + 1]] / (st[[t + 1]] + sv^2);  
  
yhp[[t + 1]] = yh[[t + 1]] + K[[t + 1]] * (x[[t + 1]] - yh[[t + 1]]);  
  
stp[[t + 1]] = (1 - K[[t + 1]]) * st[[t + 1]];  
  
t++;  
t = Min[t, iter - 1];  
ListPlot[{yhp, x}, Joined → {True, False, True}]
```

Out[1013]=




Try $sv = .001*sv0$;

■ Show the Kalman filter estimate together with the true path

In[1014]:=

```
Show[ListPlot[ys, Joined → True],
      ErrorListPlot[Table[{yhp[[i]], Sqrt[stp[[i]]]}, {i, 1, iter}]]]
```

Out[1014]=



2D tracking example

```
sw0 = 10.0;
Q = DiagonalMatrix[{sw0, sw0, .2 * sw0, .2 * sw0}];
ndistw = MultinormalDistribution[{0, 0, 0, 0}, Q];
w := RandomReal[ndistw];

sv0 = 50.0;
R0 = {{sv0, 0.0}, {0.0, sv0}};
ndistv = MultinormalDistribution[{0, 0}, R0];
v := RandomReal[ndistv];
```


■ Generate synthetic data

```
iter = 100;  
y0 = {.0, 0.0, .1, .1};
```

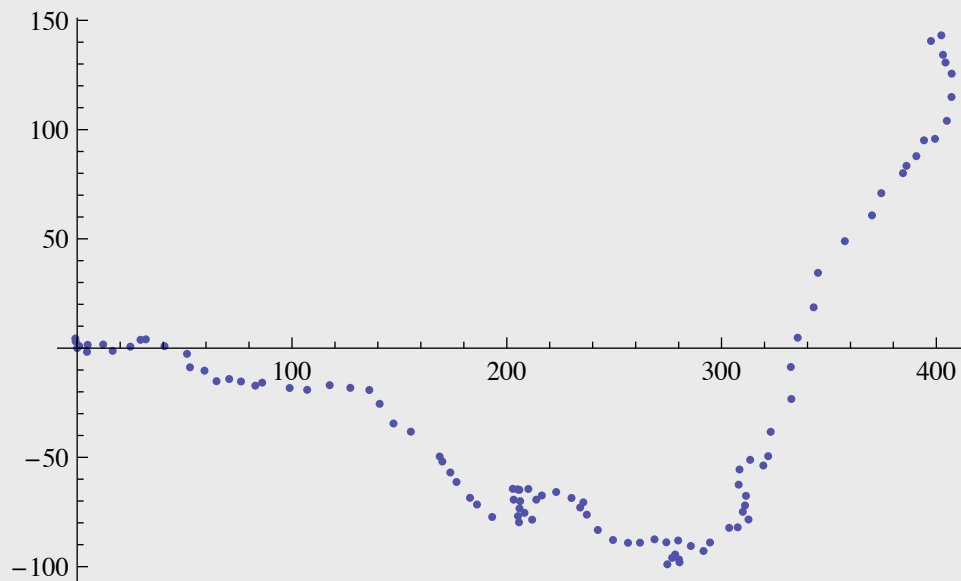
```
A = {{1, 0, 1, 0}, {0, 1, 0, 1}, {0, 0, 1, 0}, {0, 0, 0, 1}};  
H = {{1, 0, 0, 0}, {0, 1, 0, 0}};
```

```
y[k_] := A.k + w  
ys = NestList[y, y0, iter - 1];  
x = H.# & /@ ys + RandomReal[ndistv, iter];
```

```
ListPlot[x]
```



```
ListPlot[H.# & /@ ys]
```




■ Initialize arrays

```
yh = Table[{0, 0, 0, 0}, {iter}];
K = Table[{{0, 0}, {0, 0}, {0, 0}, {0, 0}}, {iter}];
P = Table[DiagonalMatrix[{0, 0, 0, 0}], {iter}];
Pm = P;
yhp = yh;
```

■ Initial estimates

```
R = 100 R0; (* Change measurement noise assumption. If R is small,
then the data is trusted--the dots get joined. If the data is
believed to be noisy,
the state estimate is smoothed *)
```

```
Dynamic[
  gxys = ListPlot[{x, H.# & /@ys, H.# & /@yhp}, Joined -> {False, True, True}]]
```



```
For[t = 1, t < iter, t++,
  (*Time update--predict the state and error covariance *)
  yh[[t + 1]] = A.yhp[[t]];
  Pm[[t + 1]] = A.P[[t]].Transpose[A] + Q;

  (*Measurement update--correction*)
  K[[t + 1]] = Pm[[t + 1]].HT.Inverse[H.Pm[[t + 1]].HT + R];

  P[[t + 1]] = (IdentityMatrix[4] - K[[t + 1]].H).Pm[[t + 1]];

  yhp[[t + 1]] = yh[[t + 1]] + K[[t + 1]].(x[[t + 1]] - H.yh[[t + 1]]);
  (*gxys=ListPlot[{x,H.#&/@ys,H.#&/@yhp},Joined->{False,True,True}];*)
  Pause[.05];
];
```

References

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