

Introduction to Neural Networks
U. Minn. Psy 5038
Daniel Kersten
Belief Propagation

Initialize

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

Last time

Neural population codes representing probability distributions


Generative modeling: Multivariate gaussian, mixtures

- Drawing samples
- Mixtures of gaussians
- Will use mixture distributions in the next lecture on EM application to segmentation

Introduction to Optimal inference and task


Introduction to Bayesian learning

Interpolation using smoothness: Gradient descent



For simplicity, we'll assume 1-D as in the lecture on sculpting the energy function. In anticipation of formulating the problem in terms of a graph that represents conditional probability dependence, we represent *observable* depth cues by y^* ,

and the true ("hidden") depth estimates by y .



(Figure from Weiss (1999)).

First-order smoothness

Earlier we saw that under specific assumptions, biologically plausible neural updating can be seen decrease the value of an energy (or "cost") function. One can also start off with an assumed cost function, determined say by a set of constraints, and use gradient descent to derive an update rule that minimizes the cost. (See supplementary material in Lecture 20). However, such an update rule will not necessarily resemble a biologically plausible neural mechanism.

For an interpolation problem, we can write the energy or cost function by:

$$J(Y) = \sum_k w_k (y_k - y_k^*)^2 + \lambda \sum_i (y_i - y_{i+1})^2$$

where $w_k = \text{xs}[[k]]$ is an "indicator function", and $y_k^* = d$, are the data values. The indicator function is 1 if there is data available, and zero otherwise. .

Gradient descent gives the following local update rule:

$$y_k \leftarrow y_k + \eta_k \left(\lambda \left(\frac{y_{k-1} + y_{k+1}}{2} - y_k \right) + w_k (y_k^* - y_k) \right)$$

λ is a free parameter that controls the degree of smoothness, i.e. smoothness at the expense of fidelity to the data.

Gauss-Seidel: $\eta[k_]:=1/(\lambda+\text{xs}[[k]])$

Successive over-relaxation (SOR): $\eta2[k_]:=1.9/(\lambda+\text{xs}[[k]])$;

A simulation: Straight line with random missing data points

■ Make the data

Consider the problem of interpolating a set of points with missing data, marked by an indicator function with the following notation:

$$w_k = \text{xs}[[k]], y^* = \text{data}, y = f.$$


We'll assume the true model is that $f=y=j$, where $j=1$ to size . `data` is a function of the sampling process on $f=j$

```
In[3]:= size = 32; xs = Table[0, {i, 1, size}]; xs[[1]] = 1; xs[[size]] = 1;
data = Table[N[j] xs[[j]], {j, 1, size}];
g3 = ListPlot[Table[N[j], {j, 1, size}], Joined -> True,
  PlotStyle -> {RGBColor[0, 0.5, 0]}];
g2 = ListPlot[data, Joined -> False,
  PlotStyle -> {Opacity[0.35], RGBColor[0.75, 0., 0], PointSize[Large]}];
```

The green line shows the a straight line connecting the data points. The red dots on the abscissa mark the points where data are missing.

```
In[4]:= Show[{g2, g3}, ImageSize -> Medium]
```

Out[4]=



Let's set up two matrices, **Tm** and **Sm** such that the gradient of the energy is equal to:

$$\mathbf{Tm} \cdot \mathbf{f} - \mathbf{Sm} \cdot \mathbf{f}.$$

Sm will be our filter to exclude non-data points. **Tm** will express the "smoothness" constraint.

```
In[5]:= Sm = DiagonalMatrix[xs];
Tm = Table[0, {i, 1, size}, {j, 1, size}];
For[i=1, i<=size, i++, Tm[[i, i]] = 2];
Tm[[1, 1]] = 1; Tm[[size, size]] = 1; (*Adjust for the boundaries*)
For[i=1, i<size, i++, Tm[[i+1, i]] = -1];
For[i=1, i<size, i++, Tm[[i, i+1]] = -1];
```

Check the update rule code for small size=10:

```
In[11]:= Clear[f, d, λ]
(λ * Tm.Array[f, size] - Sm.((Array[d, size]) - Array[f, size])) //
  MatrixForm;
```

■ Run gradient descent

```
In[13]:= Clear[Tf, f1];
dt = 1; λ = 2;
Tf[f1_] := f1 - dt*(1/(λ+xs))*(Tm.f1 - λ*Sm.(data-f1));
```

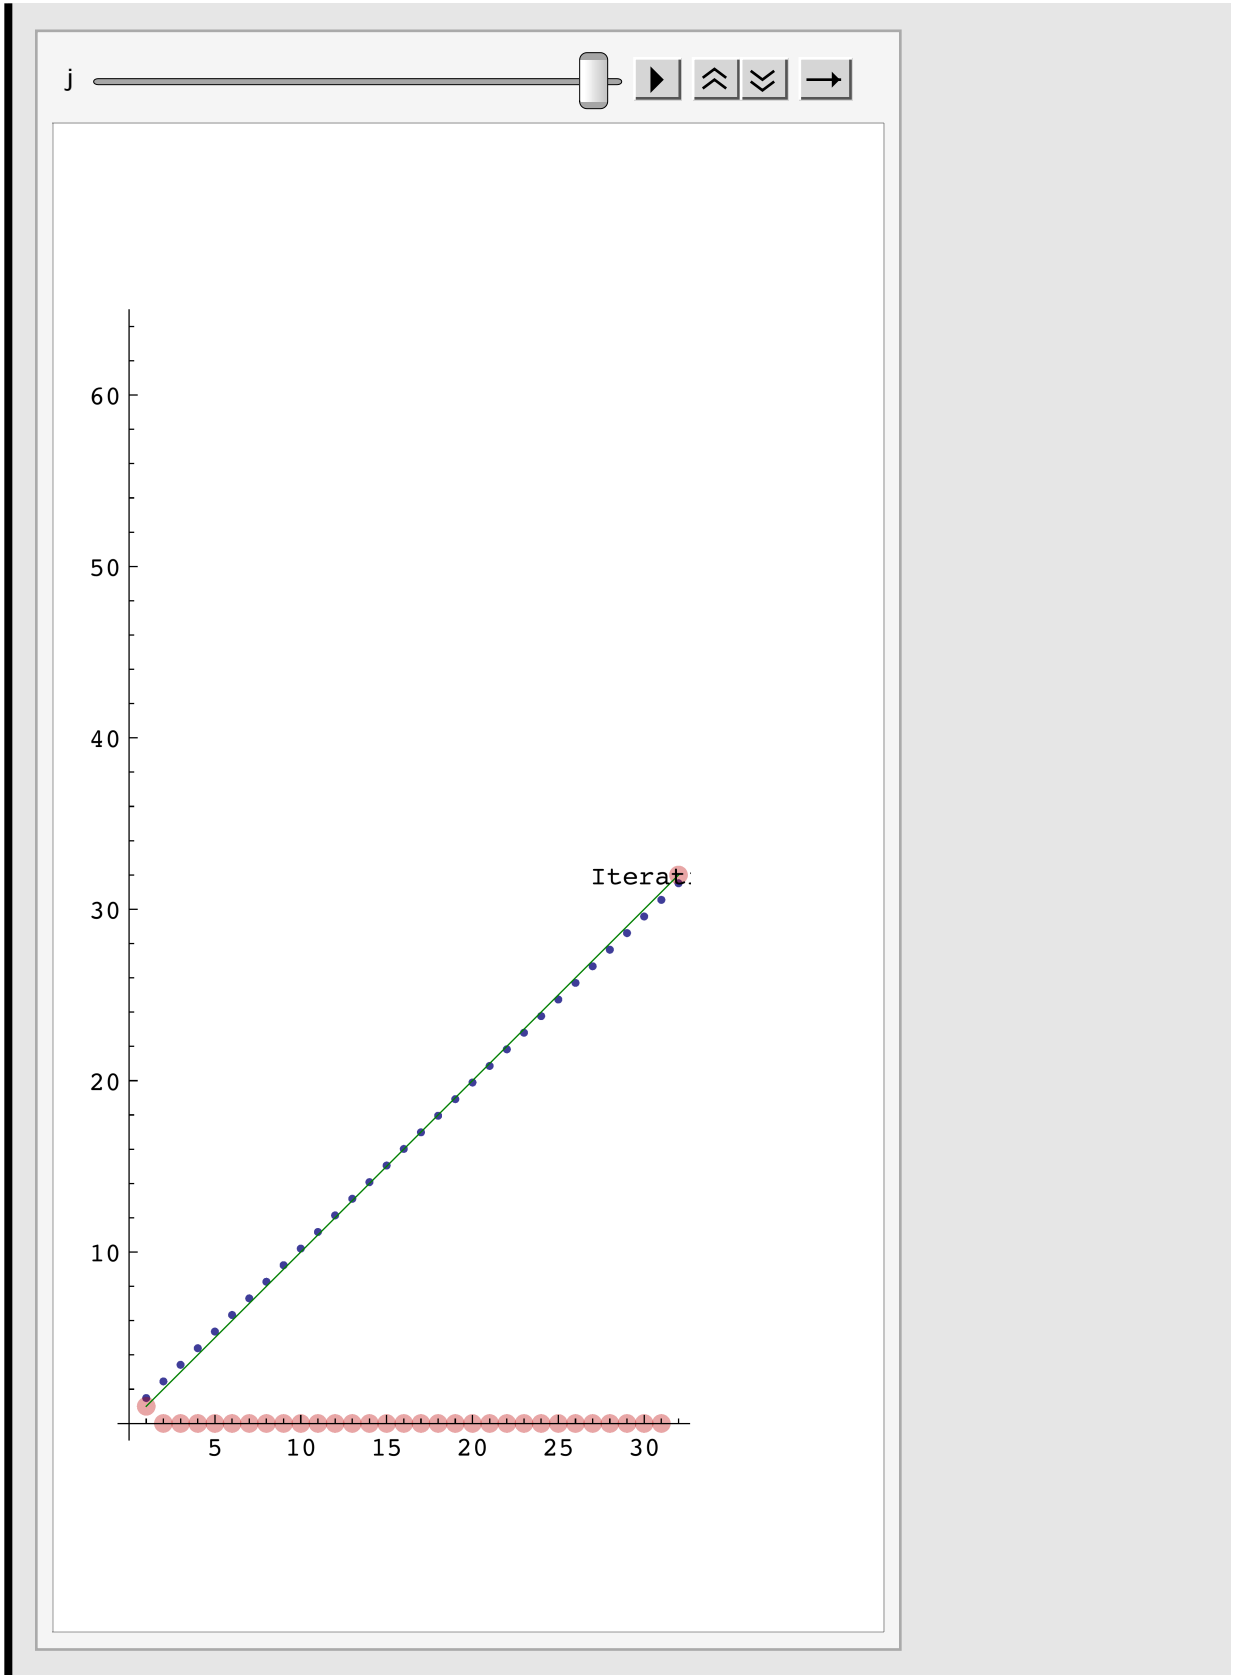
We will initialize the state vector to zero, and then run the network for **iter** iterations:

```
In[46]:= f0 = Table[0, {i, 1, size}];
result = f0;
f = f0;
iter = 25;
```

Now plot the interpolated function.

```
Animate[
  result = Nest[Tf, f, iter];
  f = result;
  g1 = ListPlot[result, Joined → False, AspectRatio → Automatic,
    PlotRange → {{0, size}, {-1, size + 1}}, ImageSize → Medium];
  Show[
    {g1, g2, g3,
      Graphics[{{Text["Iteration=" <> ToString[iter], {
        size/2, size/2}]}]}],
    PlotRange → {-1, size + 1}], {j, 1, 10}]
```

Out[45]=




Try starting with $f =$ random values. Try various numbers of iterations.

Try different sampling functions $xs[[i]]$.

Belief Propagation

Same interpolation problem, but now using belief propagation

Example is taken from Yair Weiss.(Weiss, 1999)



Probabilistic generative model

$$\text{data}[[i]] = y^*[i] = \text{xs}[[i]] y[[i]] + \text{dnoise}, \quad \text{dnoise} \sim N[0, \sigma_D] \quad (1)$$

$$y[[i+1]] = y[[i]] + \text{znoise}, \quad \text{znoise} \sim N[0, \sigma_R] \quad (2)$$

The first term is the "data formation" model, i.e. how the data is directly influenced by the interaction of the underlying influences or causes: y with the sampling and noise. The second term reflects our prior assumptions about the smoothness of y , i.e. nearby y 's are correlated, and in fact identical except for some added noise. So with no noise the prior reflects the assumption that lines are horizontal--all y 's are the same.

Some theory

We'd like to know the distribution of the random variables at each node i , conditioned on all the data: I.e. we want the posterior

$$p(Y_i=u \mid \text{all the data})$$

If we could find this, we'd be able to: 1) say what the most probable value of the y value is, and 2) give a measure of confidence

Let $p(Y_i=u \mid \text{all the data})$ be normally distributed: $\text{NormalDistribution}[\mu_i, \sigma_i]$.

Consider the i th unit. The posterior $p(Y_i=u \mid \text{all the data}) =$

$$p(Y_i=u \mid \text{all the data}) \propto p(Y_i=u \mid \text{data before } i) p(\text{data at } i \mid Y_i=u) p(Y_i=u \mid \text{data after } i) \quad (3)$$

Suppose that $p(Y_i=u \mid \text{data before } i)$ is also gaussian:

$$p(Y_i=u \mid \text{data before } i) = \alpha[u] \sim \text{NormalDistribution}[\mu\alpha, \sigma\alpha]$$

and so is probability conditioned on the data after i :

$$p(Y_i=u \mid \text{data after } i) = \beta[u] \sim \text{NormalDistribution}[\mu\beta, \sigma\beta]$$

And the noise model for the data:

$$p(\text{data at } i \mid Y_i=u) = L[u] \sim \text{NormalDistribution}[\mathbf{y}_p, \sigma_D]$$

$$\mathbf{y}_p = \text{data}[[i]]$$

So in terms of these functions, the posterior probability of the i th unit taking on the value u can be expressed as proportional to a product of the three factors:

$$p(Y_i=u \mid \text{all the data}) \propto \alpha[u] * L[u] * \beta[u] \quad (4)$$

```
αdist = NormalDistribution[μα, σα];
```

```
α[u] = PDF[αdist, u];
```

```
Ddist = NormalDistribution[yp, σD];
```

```
L[u] = PDF[Ddist, u];
```

```
βdist = NormalDistribution[μβ, σβ];
```

```
β[u] = PDF[βdist, u];
```

```
α[u] * L[u] * β[u]
```

$$e^{-\frac{(u-\mu\alpha)^2}{2\sigma\alpha^2} - \frac{(u-\mu\beta)^2}{2\sigma\beta^2} - \frac{(u-\mathbf{y}_p)^2}{2\sigma_D^2}}$$

$$\frac{1}{2\sqrt{2}\pi^{3/2}\sigma\alpha\sigma\beta\sigma_D}$$

This just another gaussian distribution on $Y_i=u$. What is its mean and variance? Finding the root enables us to complete the square to see what the numerator looks like. In particular, what the mode (=mean for gaussian) is.

$$\text{Solve}\left[-D\left[-\frac{(u - \mu\alpha)^2}{2\sigma\alpha^2} - \frac{(u - \mu\beta)^2}{2\sigma\beta^2} - \frac{(u - y_p)^2}{2\sigma_D^2}, u\right], u\right] == 0, u$$

$$\left\{ \left\{ u \rightarrow \frac{\frac{\mu\alpha}{\sigma\alpha^2} + \frac{\mu\beta}{\sigma\beta^2} + \frac{y_p}{\sigma_D^2}}{\frac{1}{\sigma\alpha^2} + \frac{1}{\sigma\beta^2} + \frac{1}{\sigma_D^2}} \right\} \right\}$$

The update rule for the variance is:

$$\sigma^2 \rightarrow \frac{1}{\sigma\alpha^2} + \frac{1}{\sigma\beta^2} + \frac{1}{\sigma_D^2}$$

How do we get $\mu\alpha, \mu\beta, \sigma\alpha, \sigma\beta$?

We express the probability of the i th unit taking on the value u in terms of the values of the neighbor before, conditioning on what is known (the observed measurements), and marginalizing over what isn't (the previous "hidden" node value, v , at the $i-1$ th location).

We have three terms to worry about that depend on nodes in the neighborhood preceding i :

$$\alpha[u] = \int_{-\infty}^{\infty} \alpha_p[v] * S[u] * L[v] dv \propto \int_{-\infty}^{\infty} e^{-\frac{(v-y_p)^2}{2\sigma_D^2} - \frac{(u-v)^2}{2\sigma_R^2} - \frac{(v-\mu\alpha_p)^2}{2\sigma\alpha_p^2}} dv \quad (5)$$

$\alpha_p = \alpha_{i-1}$. $S[u]$ is our smoothing term, or transition probability: $S[u] = p(u | v)$. $L[v]$ is the likelihood of the previous data node, given its hidden node value, v .

```
Rdist = NormalDistribution[v,  $\sigma_R$ ];
S[u] = PDF[Rdist, u];

avdist = NormalDistribution[ $\mu\alpha_p, \sigma\alpha_p$ ];
 $\alpha_p[v] = PDF[avdist, v];$ 

Lp[v] = PDF[Ddist, v];
```

Integrate[$\alpha_p[v] * S[u] * Lp[v]$, { v , -Infinity, Infinity}]

$$\frac{1}{2 \sqrt{2} \pi^{3/2} \sigma_D \sigma_R \sigma \alpha_p} \text{If} \left[\text{Re} \left[\frac{1}{\sigma_D^2} + \frac{1}{\sigma_R^2} + \frac{1}{\sigma \alpha_p^2} \right] > 0, \right.$$

$$\frac{e^{-\frac{(u - \mu \alpha_p)^2 \sigma_D^2 + \mu \alpha_p^2 \sigma_R^2 + u^2 \sigma \alpha_p^2 + y_p^2 (\sigma_R^2 + \sigma \alpha_p^2) - 2 y_p (\mu \alpha_p \sigma_R^2 + u \sigma \alpha_p^2)}{2 (\sigma_R^2 \sigma \alpha_p^2 + \sigma_D^2 (\sigma_R^2 + \sigma \alpha_p^2))}}}{\sqrt{2 \pi} \sqrt{\frac{1}{\sigma_D^2} + \frac{1}{\sigma_R^2} + \frac{1}{\sigma \alpha_p^2}}}, \text{Integrate} \left[\right.$$

$$\left. e^{\frac{1}{2} \left(-\frac{(v - y_p)^2}{\sigma_D^2} - \frac{(u - v)^2}{\sigma_R^2} - \frac{(v - \mu \alpha_p)^2}{\sigma \alpha_p^2} \right)} \right], \{v, -\infty, \infty\}, \text{Assumptions} \rightarrow \text{Re} \left[\frac{1}{\sigma_D^2} + \frac{1}{\sigma_R^2} + \frac{1}{\sigma \alpha_p^2} \right] \leq 0 \left. \right]$$

■ Some uninspired *Mathematica* manipulations

To find an expression for the mode of the above calculated expression for $\alpha[u]$

$$\text{D} \left[-\frac{(u - \mu \alpha_p)^2 \sigma_D^2 + \mu \alpha_p^2 \sigma_R^2 + u^2 \sigma \alpha_p^2 + y_p^2 (\sigma_R^2 + \sigma \alpha_p^2) - 2 y_p (\mu \alpha_p \sigma_R^2 + u \sigma \alpha_p^2)}{2 (\sigma_R^2 \sigma \alpha_p^2 + \sigma_D^2 (\sigma_R^2 + \sigma \alpha_p^2))}, u \right]$$

$$\frac{2(u - \mu \alpha_p) \sigma_D^2 + 2u \sigma \alpha_p^2 - 2y_p \sigma \alpha_p^2}{2((\sigma_R^2 + \sigma \alpha_p^2) \sigma_D^2 + \sigma_R^2 \sigma \alpha_p^2)}$$

Solve[-% == 0, u]

$$\left\{ \left\{ u \rightarrow \frac{\mu \alpha_p \sigma_D^2 + y_p \sigma \alpha_p^2}{\sigma_D^2 + \sigma \alpha_p^2} \right\} \right\}$$

$$\text{Simplify} \left[\left(\frac{\mu \alpha_p \sigma_D^2}{\sigma_R^2 \sigma \alpha_p^2 + \sigma_D^2 (\sigma_R^2 + \sigma \alpha_p^2)} + \frac{y_p \sigma \alpha_p^2}{\sigma_R^2 \sigma \alpha_p^2 + \sigma_D^2 (\sigma_R^2 + \sigma \alpha_p^2)} \right) / (\sigma_D^2 * \sigma \alpha_p^2) \right]$$

$$\frac{\mu \alpha_p \sigma_D^2 + y_p \sigma \alpha_p^2}{\sigma \alpha_p^2 (\sigma_R^2 + \sigma \alpha_p^2) \sigma_D^2 + \sigma_R^2 \sigma \alpha_p^4}$$

$$\text{Simplify} \left[\left(\frac{\sigma_D^2}{\sigma_R^2 \sigma \alpha_p^2 + \sigma_D^2 (\sigma_R^2 + \sigma \alpha_p^2)} + \frac{\sigma \alpha_p^2}{\sigma_R^2 \sigma \alpha_p^2 + \sigma_D^2 (\sigma_R^2 + \sigma \alpha_p^2)} \right) / (\sigma_D^2 + \sigma \alpha_p^2) \right]$$

$$\frac{\sigma_D^2 + \sigma \alpha_p^2}{\sigma \alpha_p^2 (\sigma_R^2 + \sigma \alpha_p^2) \sigma_D^4 + \sigma_R^2 \sigma \alpha_p^4 \sigma_D^2}$$

$$\left(\frac{\mu \alpha_p \sigma_D^2 + y_p \sigma \alpha_p^2}{\sigma_D^2 \sigma_R^2 \sigma \alpha_p^4 + \sigma_D^4 \sigma \alpha_p^2 (\sigma_R^2 + \sigma \alpha_p^2)} \right) / \left(\frac{\sigma_D^2 + \sigma \alpha_p^2}{\sigma_D^2 \sigma_R^2 \sigma \alpha_p^4 + \sigma_D^4 \sigma \alpha_p^2 (\sigma_R^2 + \sigma \alpha_p^2)} \right)$$

$$\frac{\mu \alpha_p \sigma_D^2 + y_p \sigma \alpha_p^2}{\sigma_D^2 + \sigma \alpha_p^2}$$

So we now have rule that tells us how to update the $\alpha(u)=p(y_i=uldata \text{ before } i)$, in terms of the mean and variance parameters of the previous node:

$$\mu \alpha \leftarrow \frac{\mu \alpha_p \sigma_D^2 + y_p \sigma \alpha_p^2}{\sigma_D^2 + \sigma \alpha_p^2} = \frac{\frac{\mu \alpha_p \sigma_D^2}{\sigma \alpha_p^2 \sigma_D^2} + \frac{y_p \sigma \alpha_p^2}{\sigma \alpha_p^2 \sigma_D^2}}{\frac{\sigma_D^2}{\sigma \alpha_p^2 \sigma_D^2} + \frac{\sigma \alpha_p^2}{\sigma \alpha_p^2 \sigma_D^2}} = \frac{\frac{\mu \alpha_p}{\sigma \alpha_p^2} + \frac{y_p}{\sigma_D^2}}{\frac{1}{\sigma \alpha_p^2} + \frac{1}{\sigma \alpha_p^2}}$$

The update rule for the variance is:

$$\sigma \alpha^2 \leftarrow \sigma_R^2 + \frac{1}{\frac{1}{\sigma_D^2} + \frac{1}{\sigma \alpha_p^2}}$$

A similar derivation gives us the rules for $\mu \beta, \sigma \beta^2$

$$\mu \beta \leftarrow \frac{\frac{\mu \beta_a + y_a}{\sigma \beta_a^2 + \sigma_D^2}}{\frac{1}{\sigma \beta_a^2} + \frac{1}{\sigma \alpha_a^2}}$$

$$\sigma \beta^2 \leftarrow \sigma_R^2 + \frac{1}{\frac{1}{\sigma_D^2} + \frac{1}{\sigma \beta_a^2}}$$

Where the subscript index p (for "previous", i.e. unit $i-1$) is replaced by a (for "after", i.e. unit $i+1$).

Recall that sometimes we have data and sometimes we don't. So replace:

$$y_p \rightarrow \mathbf{xs}[i-1] \mathbf{data}[i-1] = w_{i-1} y_{i-1}^*$$

(6)

And similarly for y_a .

■ Summary of update rules

The ratio, $\left(\frac{\sigma_D}{\sigma_R}\right)^2$ plays the role of λ above. If $\sigma_D^2 \gg \sigma_R^2$, there is greater smoothing. If $\sigma_D^2 \ll \sigma_R^2$, there is more fidelity to the data. (Recall $y^* \rightarrow \text{data}.w_k \rightarrow \text{xs}[[k]]$). But now we have a principled way of assigning the relative amount of smoothing.


We'll follow Weiss, and also make a (hopefully not too confusing) notation change to avoid the square superscripts for $\sigma_D^2 \rightarrow \sigma_D, \sigma_R^2 \rightarrow \sigma_R$.

$$\mu_i \leftarrow \frac{\frac{w_i}{\sigma_D} Y_i^* + \frac{1}{\sigma_i^\alpha} \mu_i^\alpha + \frac{1}{\sigma_i^\beta} \mu_i^\beta}{\frac{w_i}{\sigma_D} + \frac{1}{\sigma_i^\alpha} + \frac{1}{\sigma_i^\beta}}$$

$$\sigma_i \leftarrow \frac{1}{\frac{w_i}{\sigma_D} + \frac{1}{\sigma_i^\alpha} + \frac{1}{\sigma_i^\beta}}$$

$$\mu_i^\alpha \leftarrow \frac{\frac{1}{\sigma_{i-1}^\alpha} \mu_{i-1}^\alpha + \frac{w_{i-1}}{\sigma_D} Y_{i-1}^*}{\frac{1}{\sigma_{i-1}^\alpha} + \frac{w_{i-1}}{\sigma_D}}$$

$$\sigma_i^\alpha \leftarrow \sigma_R + \left(\frac{1}{\sigma_{i-1}^\alpha} + \frac{w_{i-1}}{\sigma_D} \right)^{-1}$$



A simulation: Belief propagation for interpolation with missing data

■ Initialization


```
In[50]:= size = 32;
           $\mu_0 = 1$ ;
           $\mu_\alpha = 1$ ;  $\sigma_\alpha = 100\,000$ ; (*large uncertainty *)
           $\mu_\beta = 1$ ;  $\sigma_\beta = 100\,000$ ; (*large*)
           $\sigma_R = 4.0$ ;  $\sigma_D = 1.0$ ;
           $\mu = \text{Table}[\mu_0, \{i, 1, \text{size}\}]$ ;
           $\sigma = \text{Table}[\sigma_\alpha, \{i, 1, \text{size}\}]$ ;
           $\mu_\alpha = \text{Table}[\mu_0, \{i, 1, \text{size}\}]$ ;
           $\sigma_\alpha = \text{Table}[\sigma_\alpha, \{i, 1, \text{size}\}]$ ;
           $\mu_\beta = \text{Table}[\mu_0, \{i, 1, \text{size}\}]$ ;
           $\sigma_\beta = \text{Table}[\sigma_\beta, \{i, 1, \text{size}\}]$ ;
          iter = 0;
          i = 1;
          j = size;
```

The code below implements the above iterative equations, taking care near the boundaries. The plot shows the estimates of $y_i = \mu$, and the error bars show $\pm\sigma_i$.

■ Belief Propagation Routine: Execute this cell "manually" for each iteration

```
In[63]:= yfit = Table[{0, 0}, {i1, 1, size}];
g1b = ErrorListPlot[{yfit}];
Dynamic[
  Show[
    {g1b, g2, g3,
      Graphics[{Text["Iteration=" <> ToString[iter], { $\frac{\text{size}}$ , size}]}]}],
    PlotRange → {-50, 50}, Axes → {False, True}]]
```

Out[65]=



Execute the next cell to run 31 iterations. The display is slowed down so that you can see the progression of the updates in the above graph.

In[66]:=

```

Do [
  Pause [.5];

  
$$\mu[[i]] = \frac{\frac{xs[[i]] data[[i]]}{\sigma D} + \frac{\mu\alpha[[i]]}{\sigma\alpha[[i]]} + \frac{1. \mu\beta[[i]]}{\sigma\beta[[i]]}}{\frac{xs[[i]]}{\sigma D} + \frac{1}{\sigma\alpha[[i]]} + \frac{1}{\sigma\beta[[i]]}};$$


  
$$\sigma[[i]] = \frac{1.}{\frac{xs[[i]]}{\sigma D} + \frac{1}{\sigma\alpha[[i]]} + \frac{1}{\sigma\beta[[i]]}};$$


  
$$\mu[[j]] = \frac{\frac{xs[[j]] data[[j]]}{\sigma D} + \frac{\mu\alpha[[j]]}{\sigma\alpha[[j]]} + \frac{1. \mu\beta[[j]]}{\sigma\beta[[j]]}}{\frac{xs[[j]]}{\sigma D} + \frac{1}{\sigma\alpha[[j]]} + \frac{1}{\sigma\beta[[j]]}};$$


  
$$\sigma[[j]] = \frac{1.}{\frac{xs[[j]]}{\sigma D} + \frac{1}{\sigma\alpha[[j]]} + \frac{1}{\sigma\beta[[j]]}};$$


  nextj = j - 1;

  
$$\mu\alpha[[nextj]] = \frac{\frac{xs[[j]] data[[j]]}{\sigma D} + \frac{1. \mu\alpha[[j]]}{\sigma\alpha[[j]]}}{\frac{xs[[j]]}{\sigma D} + \frac{1}{\sigma\alpha[[j]]}};$$


  
$$\sigma\alpha[[nextj]] = \sigma R + \frac{1.}{\frac{xs[[j]]}{\sigma D} + \frac{1}{\sigma\alpha[[j]}};$$


  nexti = i + 1;

  
$$\mu\beta[[nexti]] = \frac{\frac{xs[[i]] data[[i]]}{\sigma D} + \frac{1. \mu\beta[[i]]}{\sigma\beta[[i]]}}{\frac{xs[[i]]}{\sigma D} + \frac{1}{\sigma\beta[[i]]}};$$


  
$$\sigma\beta[[nexti]] = \sigma R + \frac{1.}{\frac{xs[[i]]}{\sigma D} + \frac{1}{\sigma\beta[[i]}};$$


  j--;
  i++;
  iter++;
  yfit = Table[{μ[[i1]], σ[[i1]]}, {i1, 1, size}];
  glb = ErrorListPlot[{yfit}];
  , {size - 1}];

```

Exercises

Run the descent algorithm using successive over-relaxation (SOR): $\eta_2[k_] := 1.9 / (\lambda + xs[[k]])$.

How does convergence compare with Gauss-Seidel?

Run Belief Propagation using: $\sigma_R = 1.0$; $\sigma_D = 4.0$; How does fidelity to the data compare with the original case

($\sigma_R = 4.0$; $\sigma_D = 1.0$).

BP with missing sine wave data

■ Generate sine wave with missing data

```
In[67]:= size = 64; xs = Table[RandomInteger[1], {i, 1, size}];
data = Table[N[Sin[2 π j / 20] xs[[j]]], {j, 1, size}];
g3b = ListPlot[Table[N[Sin[2 π j / 20]], {j, 1, size}], Joined → True,
PlotStyle → {RGBColor[0, 0.5, 0]}];
g2b = ListPlot[data, Joined → False, PlotStyle → {RGBColor[0.75, 0., 0]}];
```


■ Initialize

```
In[68]:= μ0 = 1;
μα = 1; σα = 100 000; (*large uncertainty *)
μβ = 1; σβ = 100 000; (*large*)
σR = .5; σD = .1;
μ = Table[μ0, {i, 1, size}];
σ = Table[σα, {i, 1, size}];
μα = Table[μ0, {i, 1, size}];
σα = Table[σα, {i, 1, size}];
μβ = Table[μ0, {i, 1, size}];
σβ = Table[σβ, {i, 1, size}];
iter = 0;
i = 1;
j = size;
```



```
In[80]:= yfit = Table[{0, 0}, {i1, 1, size}];  
g1bb = ErrorListPlot[{yfit}];  
Dynamic[Show[{g1bb, g2b, g3b}, PlotRange → {-2, 2}, Axes → {False, True}]]
```

Out[82]=



■ SINE WAVE DEMO: Belief Propagation Routine

```

Do [
  Pause [0.2];

  
$$\mu[i] = \frac{\frac{xs[i] \text{ data}[i]}{\sigma_D} + \frac{\mu\alpha[i]}{\sigma\alpha[i]} + \frac{1 \cdot \mu\beta[i]}{\sigma\beta[i]}}{\frac{xs[i]}{\sigma_D} + \frac{1}{\sigma\alpha[i]} + \frac{1}{\sigma\beta[i]}};$$


  
$$\sigma[i] = \frac{1.}{\frac{xs[i]}{\sigma_D} + \frac{1}{\sigma\alpha[i]} + \frac{1}{\sigma\beta[i]}};$$


  
$$\mu[j] = \frac{\frac{xs[j] \text{ data}[j]}{\sigma_D} + \frac{\mu\alpha[j]}{\sigma\alpha[j]} + \frac{1 \cdot \mu\beta[j]}{\sigma\beta[j]}}{\frac{xs[j]}{\sigma_D} + \frac{1}{\sigma\alpha[j]} + \frac{1}{\sigma\beta[j]}};$$


  
$$\sigma[j] = \frac{1.}{\frac{xs[j]}{\sigma_D} + \frac{1}{\sigma\alpha[j]} + \frac{1}{\sigma\beta[j]}};$$


  nextj = j - 1;

  
$$\mu\alpha[\text{nextj}] = \frac{\frac{xs[j] \text{ data}[j]}{\sigma_D} + \frac{1 \cdot \mu\alpha[j]}{\sigma\alpha[j]}}{\frac{xs[j]}{\sigma_D} + \frac{1}{\sigma\alpha[j]}};$$


  
$$\sigma\alpha[\text{nextj}] = \sigma_R + \frac{1.}{\frac{xs[j]}{\sigma_D} + \frac{1}{\sigma\alpha[j]}};$$


  nexti = i + 1;

  
$$\mu\beta[\text{nexti}] = \frac{\frac{xs[i] \text{ data}[i]}{\sigma_D} + \frac{1 \cdot \mu\beta[i]}{\sigma\beta[i]}}{\frac{xs[i]}{\sigma_D} + \frac{1}{\sigma\beta[i]}};$$


  
$$\sigma\beta[\text{nexti}] = \sigma_R + \frac{1.}{\frac{xs[i]}{\sigma_D} + \frac{1}{\sigma\beta[i]}};$$


  j--;
  i++;
  iter++;
  yfit = Table[{μ[i1], σ[i1]}, {i1, 1, size}];
  glbb = ErrorListPlot[{yfit}];
  , {size - 1}]

```

References

- Applebaum, D. (1996). Probability and Information . Cambridge, UK: Cambridge University Press.
- Frey, B. J. (1998). *Graphical Models for Machine Learning and Digital Communication*. Cambridge, Massachusetts: MIT Press.
- Jepson, A., & Black, M. J. (1993). *Mixture models for optical flow computation*. Paper presented at the Proc. IEEE Conf. Comput. Vision Pattern Recog., New York.
- Kersten, D. and P.W. Schrater (2000), *Pattern Inference Theory: A Probabilistic Approach to Vision*, in *Perception and the Physical World*, R. Mausfeld and D. Heyer, Editors. , John Wiley & Sons, Ltd.: Chichester. (pdf)
- Kersten, D., & Madarasmı, S. (1995). The Visual Perception of Surfaces, their Properties, and Relationships. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 19, 373-389.
- Madarasmı, S., Kersten, D., & Pong, T.-C. (1993). The computation of stereo disparity for transparent and for opaque surfaces. In C. L. Giles & S. J. Hanson & J. D. Cowan (Eds.), *Advances in Neural Information Processing Systems 5*. San Mateo, CA: Morgan Kaufmann Publishers.
- Pearl, Judea. (1997) Probabilistic Reasoning in Intelligent Systems : Networks of Plausible Inference. (amazon.com link)
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge, UK: Cambridge University Press.
- Weiss Y. (1999) Bayesian Belief Propagation for Image Understanding submitted to SCTV 1999. (gzipped postscript 297K)
- Weiss, Y. (1997). *Smoothness in Layers: Motion segmentation using nonparametric mixture estimation*. Paper presented at the Proceedings of IEEE conference on Computer Vision and Pattern Recognition.
- Yuille, A., Coughlan J., Kersten D.(1998) (pdf)

For notes on Graphical Models, see:<http://www.cs.berkeley.edu/~murphyk/Bayes/bayes.html>