

CRF example — written after class 10/20/16
(CS 585, UMass Amherst, Brendan O'Connor)

 finna bless us
y = [S] V V V

Tags: "V"erb and pr"O"noun (and [S]tart)

Let's use three feature templates:

Transition features:
for example
 $f_{VV}(x,y)$ = number of
V-V transitions in y

Word-tag observation
features: for example
 $f_{V,dog}(x,y)$ = number of
tokens that are word "dog"
under a Verb tag

"ends with s"-tag features:
 $f_{V-s}(x,y)$ = number of
tokens that end with -s and
are tagged as Verb

(Global features have to be COUNTS: the reason why is further below.)

For 3 word vocabulary and 2 tag types, that's J=14 total features.

Assume we have fixed model weights θ and would like to score the goodness of the above tag sequence.

Global feature vector $f(x,y)$ =	<table style="border-collapse: collapse; text-align: center;"> <tr> <td style="padding: 2px 5px;">f_{SV}</td><td style="padding: 2px 5px;">f_{SO}</td><td style="padding: 2px 5px;">$f_{V,V}$</td><td style="padding: 2px 5px;">$f_{V,O}$</td><td style="padding: 2px 5px;">$f_{O,V}$</td><td style="padding: 2px 5px;">$f_{O,O}$</td><td style="padding: 2px 5px;">$f_{V,finna}$</td><td style="padding: 2px 5px;">$f_{V,bless}$</td><td style="padding: 2px 5px;">$f_{V,us}$</td><td style="padding: 2px 5px;">$f_{O,finna}$</td><td style="padding: 2px 5px;">$f_{O,bless}$</td><td style="padding: 2px 5px;">$f_{O,us}$</td><td style="padding: 2px 5px;">$f_{V,-s}$</td><td style="padding: 2px 5px;">$f_{O,-s}$</td> </tr> <tr> <td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">2</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">2</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">0</td> </tr> </table>	f_{SV}	f_{SO}	$f_{V,V}$	$f_{V,O}$	$f_{O,V}$	$f_{O,O}$	$f_{V,finna}$	$f_{V,bless}$	$f_{V,us}$	$f_{O,finna}$	$f_{O,bless}$	$f_{O,us}$	$f_{V,-s}$	$f_{O,-s}$	1	0	2	0	0	0	1	0	1	1	0	0	2	0
f_{SV}	f_{SO}	$f_{V,V}$	$f_{V,O}$	$f_{O,V}$	$f_{O,O}$	$f_{V,finna}$	$f_{V,bless}$	$f_{V,us}$	$f_{O,finna}$	$f_{O,bless}$	$f_{O,us}$	$f_{V,-s}$	$f_{O,-s}$																
1	0	2	0	0	0	1	0	1	1	0	0	2	0																
Model parameters θ =	<table style="border-collapse: collapse; text-align: center;"> <tr> <td style="padding: 2px 5px;">θ_{SV}</td><td style="padding: 2px 5px;">θ_{SO}</td><td style="padding: 2px 5px;">$\theta_{V,V}$</td><td style="padding: 2px 5px;">$\theta_{V,O}$</td><td style="padding: 2px 5px;">$\theta_{O,V}$</td><td style="padding: 2px 5px;">$\theta_{O,O}$</td><td style="padding: 2px 5px;">$\theta_{V,finna}$</td><td style="padding: 2px 5px;">$\theta_{V,bless}$</td><td style="padding: 2px 5px;">$\theta_{V,us}$</td><td style="padding: 2px 5px;">$\theta_{O,finna}$</td><td style="padding: 2px 5px;">$\theta_{O,bless}$</td><td style="padding: 2px 5px;">$\theta_{O,us}$</td><td style="padding: 2px 5px;">$\theta_{V,-s}$</td><td style="padding: 2px 5px;">$\theta_{O,-s}$</td> </tr> <tr> <td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-0.2</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-0.8</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+0.1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+0.5</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+4.3</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-0.3</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-1.2</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-0.1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+0.1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+5.3</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-4.1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">-0.3</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+1.1</td><td style="border: 1px solid black; width: 20px; height: 20px; text-align: center;">+2.2</td> </tr> </table>	θ_{SV}	θ_{SO}	$\theta_{V,V}$	$\theta_{V,O}$	$\theta_{O,V}$	$\theta_{O,O}$	$\theta_{V,finna}$	$\theta_{V,bless}$	$\theta_{V,us}$	$\theta_{O,finna}$	$\theta_{O,bless}$	$\theta_{O,us}$	$\theta_{V,-s}$	$\theta_{O,-s}$	-0.2	-0.8	+0.1	+0.5	+4.3	-0.3	-1.2	-0.1	+0.1	+5.3	-4.1	-0.3	+1.1	+2.2
θ_{SV}	θ_{SO}	$\theta_{V,V}$	$\theta_{V,O}$	$\theta_{O,V}$	$\theta_{O,O}$	$\theta_{V,finna}$	$\theta_{V,bless}$	$\theta_{V,us}$	$\theta_{O,finna}$	$\theta_{O,bless}$	$\theta_{O,us}$	$\theta_{V,-s}$	$\theta_{O,-s}$																
-0.2	-0.8	+0.1	+0.5	+4.3	-0.3	-1.2	-0.1	+0.1	+5.3	-4.1	-0.3	+1.1	+2.2																

$$\begin{aligned} \text{Goodness score } G(y) &= \theta' f(x, y) = \sum_{j=1}^J \theta_j f_j(x, y) \\ &= -0.2 + 0 + 0.2 + 0 + 0 + 0 - 1.2 + 0 + 0.1 + 5.3 + 0 + 0 + 2.2 + 0 \end{aligned}$$

Global feature vector is from the sum of local feature vectors

$$f(x, y) = \sum_t f_t(y_{t-1}, y_t, x_t)$$

$f_t(y_{t-1}, y_t, x_t)$ = local feature vector including the transition between these two tags, and the observation of word at position t.

The local features are, for example:

$$f_{VV}(y_{prev}, y_{cur}, \text{curword}) = \{1 \text{ if } y_{prev}=V \text{ and } y_{cur}=V, \text{ else } 0\}$$

$$f_{V,dog}(y_{prev}, y_{cur}, \text{curword}) = \{1 \text{ if } y_{cur}=V \text{ and } \text{curword}=\text{"dog"}, \text{ else } 0\}$$

$$f_{V,-s}(y_{prev}, y_{cur}, \text{curword}) = \{1 \text{ if } y_{cur}=V \text{ and } \text{curword ends in "s"}, \text{ else } 0\}$$

And so on, repeated for different tags and words.

Example

.....trans. feats..... obs. feats.....

	f_{SV}	f_{SO}	f_{VW}	f_{VO}	f_{OV}	f_{OO}	$f_{V,finna}$	$f_{V,bless}$	$f_{V,us}$	$f_{O,finna}$	$f_{O,bless}$	$f_{O,us}$	$f_{V,-s}$	$f_{O,-s}$
f(START, V, finna)	1	0	0	0	0	0	1	0	0	0	0	0	0	0
+ f(V, V, bless)	0	0	1	0	0	0	0	1	0	0	0	0	1	0
+ f(V, V, us)	0	0	1	0	0	0	0	0	1	0	0	0	1	0
= f(x=finna bless us, y=V V V) =	1	0	2	0	0	0	1	0	1	1	0	0	2	0

Local feature decomposition implies that the scoring function decomposes, too.

$$G(y) = \theta' f(x, y) = \theta' \sum_t f_t(y_{t-1}, y_t, x_t) = \sum_t \theta' f_t(y_{t-1}, y_t, x_t)$$

$$= \theta' f(\text{START}, V, \text{finna}) + \theta' f(V, V, \text{bless}) + \theta' f(V, V, \text{us})$$

$$= \text{dotprod} \left(\begin{array}{c} \begin{array}{cccccccccccccccc} -0.2 & -0.8 & +0.1 & +0.5 & +4.3 & -0.3 & -1.2 & -0.1 & +0.1 & +5.3 & -4.1 & -0.3 & +1.1 & +2.2 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \\ + \\ \begin{array}{c} \begin{array}{cccccccccccccccc} -0.2 & -0.8 & +0.1 & +0.5 & +4.3 & -0.3 & -1.2 & -0.1 & +0.1 & +5.3 & -4.1 & -0.3 & +1.1 & +2.2 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \end{array} \\ + \\ \begin{array}{c} \begin{array}{cccccccccccccccc} -0.2 & -0.8 & +0.1 & +0.5 & +4.3 & -0.3 & -1.2 & -0.1 & +0.1 & +5.3 & -4.1 & -0.3 & +1.1 & +2.2 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \end{array} \end{array} \right)$$

In the assignment we're compiling these local feature scoring functions into A and B_t functions.

A(yprev, ycur) is the matrix of transition logprobs.

(If this was an HMM, A would be the table of log condprobs for possible transitions.)

$$A(y_{t-1}, y_t) = \sum_{j \in \text{transfeats}} \theta_j f_j(y_{t-1}, y_t, x_t)$$

B_t(ycur) is the log-prob weight for tag ycur, according to observation information at t.

(If this was an HMM, B_t would be the column of log condprobs for the word.)

$$B_t(y_t) = \sum_{j \in \text{obsfeats}} \theta_j f_j(\text{null}, y_t, x_t)$$

I'm writing "null" for the prev tag because the obs features don't depend on the prev tag: only the current one.

In HW4, note we separate the local feature function into two feature functions: one for the transition feature, and one for observation features. In lecture I tried to talk a single local feature function combining transition and observation features at the position just because that's how the Chen blogpost does it.

Why'd we remove the x's from the math notation? We could've kept them if we wanted to. We just want to set up the graph weights for the Viterbi algorithm. Viterbi doesn't need to know the words. It only has to know the log-prob weights for different tags at each timestep, and the log-prob weights for different tag transitions.

Sparse vectors: The feature vectors are big long vectors where most elements are 0. In practice, you should never actually allocate an array of length J in your computer's memory. Instead, use a "sparse representation" where you use a Python dict where you have key-value pairs only for features with a nonzero value.