## CSE 427

## Markov Models and Hidden Markov Models

## How Proteins "Read" DNA

E.g.:


Helix-Turn-Helix Motif
Leucine Zipper Motif

## Down in the Groove

Different patterns of hydrophobic methyls, potential H bonds, etc. at edges of different base pairs. They're accessible, esp. in major groove


## DNA Methylation

CpG - 2 adjacent nts, same strand (not Watson-Crick pair; " $p$ " mnemonic for the phosphodiester bond of the DNA backbone)

C of CpG is often (70-80\%) methylated in mammals i.e., CH3 group added (both strands)

cytosine

Why? Generally silences transcription. X-inactivation, imprinting, repression of mobile elements, some cancers, aging, and developmental differentiation

How? DNA methyltransferases convert hemi- to fullymethylated

Major exception: promoters of housekeeping genes

## Same

 PairingMethyl-C alters major groove profile, not base-pairing


## ${ }^{6}$ CpG Islands"

Methyl-C mutates to T relatively easily
Net: CpG is less common than expected genome-wide:

cytosine $\mathrm{f}(\mathrm{CpG})<\mathrm{f}(\mathrm{C})^{* f}(\mathrm{G})$
BUT in promoter (\& other) regions, CpG remain unmethylated, so CpG $\rightarrow$ TpG less likely there: makes "CpG Islands"; often mark gene-rich regions

thymine

## QDASAS

CpG Islands
More CpG than elsewhere
More C \& G than elsewhere, too
Typical length: few 100 to few 1000 bp
Questions
Is a short sequence (say, 200 bp ) a CpG island or not? Given long sequence (say, IO-IOOkb), find CpG islands?

## Markov \& Hidden Markov Models

## References:

Eddy, "What is a hidden Markov model?" Nature Biotechnology, 22, \#I0 (2004) I3I5-6.

Durbin, Eddy, Krogh and Mitchison, "Biological Sequence Analysis", Cambridge, 1998

Rabiner, "A Tutorial on Hidden Markov Models and Selected Application in Speech Recognition," Proceedings of the IEEE, v 77 \#2,Feb 1989, 257-286

## Independence

A key issue: Previous models we've talked about assume independence of nucleotides in different positions - definitely unrealistic.

## Markov Chains

A sequence $x_{1}, x_{2}, \ldots$ of random variables is a $k$-th order Markov chain if, for all $i, i^{\text {th }}$ value is independent of all but the previous $k$ values:

$$
P\left(x_{i} \mid x_{1}, x_{2}, \ldots, x_{i-1}\right)=P\left(x_{i} \mid x_{i-k}, x_{i-k+1}, \ldots, x_{i-1}\right)
$$

Example I: Uniform random ACGT Example 2: Weight matrix model
Example 3: ACGT, but $\downarrow \operatorname{Pr}(G$ following C)

## A Markov Model (Ist order)



States: A,C,G,T
Emissions: corresponding letter
Transitions: $a_{s t}=P\left(x_{i}=t \mid x_{i-1}=s\right) \longleftarrow$ st order

## A Markov Model (Ist order)



States: A,C,G,T
Emissions: corresponding letter
Transitions: $a_{s t}=P\left(x_{i}=t \mid x_{i-1}=s\right)$
Begin/End states

## $\operatorname{Pr}$ of emitting sequence $x$

$$
\begin{aligned}
x & =x_{1} x_{2} \ldots x_{n} \\
P(x) & =P\left(x_{1}, x_{2}, \ldots, x_{n}\right) \sum_{\text {laws of probability }} \\
& =P\left(x_{1}\right) \cdot P\left(x_{2} \mid x_{1}\right) \cdots P\left(x_{n} \mid x_{n-1}, \ldots, x_{1}\right) \\
& =P\left(x_{1}\right) \cdot P\left(x_{2} \mid x_{1}\right) \cdots P\left(x_{n} \mid x_{n-1}\right) \sum_{\text {if }}^{1 s^{t}} \text { order } \\
& =P\left(x_{1}\right) \prod_{i=1}^{n-1} a_{x_{i}, x_{i+1}} \\
& =\prod_{i=0}^{n-1} a_{x_{i}, x_{i+1}} \quad \text { (with Begin state) }
\end{aligned}
$$

## Training

Max likelihood estimates for transition probabilities are just the frequencies of transitions when emitting the training sequences
E.g., from 48 CpG islands in 60 kbp :

| + | A | c | G | T | - | A | C | G | T |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| A | 0.180 | 0.274 | 0.426 | 0.120 | A | 0.300 | 0.205 | 0.285 | 0.210 |
| C | 0.171 | 0.368 | 0.274 | 0.188 | C | 0.322 | 0.298 | 0.078 | 0.302 |
| G | 0.161 | 0.339 | $\overline{0.375}$ | 0.125 | G | 0.248 | 0.246 | 0.298 | 0.208 |
| T | 0.079 | 0.355 | 0.384 | 0.182 | T | 0.177 | 0.239 | 0.292 | 0.292 |

## Discrimination/Classification

Log likelihood ratio of CpG model vs background model

## CpG Island Scores



Figure 3.2 The histogram of the length-normalised scores for all the sequences. CpG islands are shown with dark grey and non-CpG with light grey.

# What does a 2nd order Markov Model look like? 

## 3rd order?

## Questions

QI: Given a short sequence, is it more likely from feature model or background model? Above
Q2: Given a long sequence, where are the features in it (if any)

Approach I: score 100 bp (e.g.) windows
Pro: simple
Con: arbitrary, fixed length, inflexible
Approach 2: combine +/- models.

## Combined Model



Emphasis is "Which (hidden) state?" not "Which model?"

## Hidden Markov Models (HMMs)

States: Paths:
Transitions:
Emissions:
Observed data: emission sequence Hidden data:
$1,2,3, \ldots$
$e_{k}(b)=P\left(x_{i}=b \mid \pi_{i}=k\right)$
state/transition sequence
sequences of states $\pi=\left(\pi_{1}, \pi_{2}, \ldots\right)$
$a_{k, l}=P\left(\pi_{i}=l \mid \pi_{i-1}=k\right)$

## The Occasionally Dishonest Casino

1 fair die, 1 "loaded" die, occasionally swapped


Rolls
315116246446644245311321631164152133625144543631656626566666 Die FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLL Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLL

Rolls 651166453132651245636664631636663162326455236266666625151631 Die LLLLLLFFFFFFFFFFFFLLLLLLLLLLLLLLLLFFFLLLLLLLLLLLLLLFFFFFFFFF Viterbi LLLLLLFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLFFFFFFFF

Rolls 222555441666566563564324364131513465146353411126414626253356 Die FFFFFFFFLLLLLLLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLL Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFL

Rolls 366163666466232534413661661163252562462255265252266435353336 Die LLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF Viterbi LLLLLLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF

Rolls 233121625364414432335163243633665562466662632666612355245242 Die FFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLLLLFFFFFFFFFFF Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLFFFFFFFFFFF

Figure 3.5 The numbers show 300 rolls of a die as described in the example. Below is shown which die was actually used for that roll (F for fair and $L$ for loaded). Under that the prediction by the Viterbi algorithm is shown.

## Inferring hidden stuff

Joint probability of a given path $\pi \&$ emission sequence $x$ :

$$
P(x, \pi)=a_{0, \pi_{1}} \prod_{i=1}^{n} e_{\pi_{i}}\left(x_{i}\right) \cdot a_{\pi_{i}, \pi_{i+1}}
$$

But $\pi$ is hidden; what to do? Some alternatives:
Most probable single path

$$
\pi^{*}=\arg \max _{\pi} P(x, \pi)
$$

Sequence of most probable states

$$
\hat{\pi}_{i}=\arg \max _{k} P\left(\pi_{i}=k \mid x\right)
$$

## The Viterbi Algorithm: The most probable path

Viterbi finds: $\quad \pi^{*}=\arg \max _{\pi} P(x, \pi)$
Possibly there are $10^{99}$ paths of prob $10^{-99}$
More commonly, one path (+ slight variants) dominate others.
(If not, other approaches may be preferable.)
Key problem: exponentially many paths $\pi$

## Unrolling an HMM



Conceptually, sometimes convenient
Note exponentially many paths

## Viterbi

$v_{l}(i)=$ probability of the most probable path emitting $x_{1}, x_{2}, \ldots, x_{i}$ and ending in state $l$

Initialize:


## Viterbi Traceback

Above finds probability of best path
To find the path itself, trace backward to the state $k$ attaining the max at each stage

Rolls
315116246446644245311321631164152133625144543631656626566666 Die FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLL Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLL

Rolls 651166453132651245636664631636663162326455236266666625151631 Die LLLLLLFFFFFFFFFFFFLLLLLLLLLLLLLLLLFFFLLLLLLLLLLLLLLFFFFFFFFF Viterbi LLLLLLFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLFFFFFFFF

Rolls 222555441666566563564324364131513465146353411126414626253356 Die FFFFFFFFLLLLLLLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLL Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFL

Rolls 366163666466232534413661661163252562462255265252266435353336 Die LLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF Viterbi LLLLLLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF

Rolls 233121625364414432335163243633665562466662632666612355245242 Die FFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLLLLFFFFFFFFFFF Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLFFFFFFFFFFF

Figure 3.5 The numbers show 300 rolls of a die as described in the example. Below is shown which die was actually used for that roll (F for fair and $L$ for loaded). Under that the prediction by the Viterbi algorithm is shown.

## Is Viterbi "best"?

Viterbi finds $\pi^{*}=\arg \max _{\pi} P(x, \pi)$


Most probable (Viterbi) path goes through 5, but most probable state at 2 nd step is 6 (I.e., Viterbi is not the only interesting answer.)

## An HMM (unrolled)



Emissions/sequence positions $\longrightarrow$

## Viterbi: best path to each state



## The Forward Algorithm

For each state/time, want total probability of all paths leading to it, with given
emissions


## The Backward Algorithm

Similar: for each state/time, want total probability of all paths from it, with given emissions, conditional on that state.


## In state $k$ at step $i$ ?

$$
\begin{aligned}
& P\left(x, \pi_{i}=k\right) \\
& \quad=P\left(x_{1}, \ldots, x_{i}, \pi_{i}=k\right) \cdot P\left(x_{i+1}, \ldots, x_{n} \mid x_{1}, \ldots, x_{i}, \pi_{i}=k\right) \\
& \quad=P\left(x_{1}, \ldots, x_{i}, \pi_{i}=k\right) \cdot P\left(x_{i+1}, \ldots, x_{n} \mid \pi_{i}=k\right) \\
& \quad=f_{k}(i) \cdot b_{k}(i) \\
& P\left(\pi_{i}=k \mid x\right)=\frac{P\left(x, \pi_{i}=k\right)}{P(x)}=\frac{f_{k}(i) \cdot b_{k}(i)}{P(x)}
\end{aligned}
$$

## Posterior Decoding, I

Alternative 1: what's the most likely state at step i?

$$
\hat{\pi}_{i}=\arg \max _{k} P\left(\pi_{i}=k \mid x\right)
$$

Note: the sequence of most likely states $\neq$ the most likely sequence of states. May not even be legal!


## The Occasionally Dishonest Casino

1 fair die, 1 "loaded" die, occasionally swapped


Rolls
315116246446644245311321631164152133625144543631656626566666 Die FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLL Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLL

Rolls 651166453132651245636664631636663162326455236266666625151631 Die LLLLLLFFFFFFFFFFFFLLLLLLLLLLLLLLLLFFFLLLLLLLLLLLLLLFFFFFFFFF Viterbi LLLLLLFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLLFFFFFFFF

Rolls 222555441666566563564324364131513465146353411126414626253356 Die FFFFFFFFLLLLLLLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLL Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFL

Rolls 366163666466232534413661661163252562462255265252266435353336 Die LLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF Viterbi LLLLLLLLLLLLFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF

Rolls 233121625364414432335163243633665562466662632666612355245242 Die FFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLLLLFFFFFFFFFFF Viterbi FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFLLLLLLLLLLLLLLLLLLLFFFFFFFFFFF

Figure 3.5 The numbers show 300 rolls of a die as described in the example. Below is shown which die was actually used for that roll (F for fair and $L$ for loaded). Under that the prediction by the Viterbi algorithm is shown.

## Posterior Decoding



Figure 3.6 The posterior probability of being in the state corresponding to the fair die in the casino example. The x axis shows the number of the roll. The shaded areas show when the roll was generated by the loaded die.

## Posterior Decoding, II

Alternative 1: what's most likely state at step $i$ ?

$$
\hat{\pi}_{i}=\arg \max _{k} P\left(\pi_{i}=k \mid x\right)
$$

Alternative 2: given some function $g(k)$ on states, what's its expectation. E.g., what's probability of " + " model in CpG HMM ( $g(k)=1$ iff $k$ is " + " state)?

$$
G(i \mid x)=\sum_{k} P\left(\pi_{i}=k \mid x\right) \cdot g(k)
$$

## CpG Islands again

Data: 41 human sequences, totaling 60 kbp , including 48 CpG islands of about Ikbp each

Viterbi:
Found 46 of 48
plus I2I "false positives"
Posterior Decoding:
same 2 false negatives
plus 236 false positives

Post-process:
46/48
67 false pos

46/48
83 false pos
Post-process: merge within
500; discard < 500

## Training

Given model topology \& training sequences, learn transition and emission probabilities

If $\pi$ known, then MLE is just frequency observed in training data

$$
\begin{aligned}
& a_{k, l}=\frac{\text { count of } k \rightarrow l \text { transitions }}{\text { count of } k \rightarrow \text { anywhere transitions }} \\
& e_{k}(b)=\ldots
\end{aligned}
$$

If $\pi$ hidden, then use EM : given $\pi$, estimate $\theta$; given $\theta$ estimate $\pi . \quad\}^{2}$ ways

## Viterbi Training given $\pi$, estimate $\theta$; given $\theta$ estimate $\pi$

Make initial estimates of parameters $\theta$
Find Viterbi path $\pi$ for each training sequence
Count transitions/emissions on those paths, getting new $\theta$
Repeat
Not rigorously optimizing desired likelihood, but still useful \& commonly used.
(Arguably good if you're doing Viterbi decoding.)

## Baum-Welch Training

 given $\theta$, estimate $\pi$ ensemble; then re-estimate $\theta$$$
\begin{aligned}
& P\left(\pi_{i}=k, \pi_{i+1}=l \mid x, \theta\right) \\
& \quad=\frac{f_{k}(i \mid \theta) a_{k, l} e_{l}\left(x_{i+1}\right) b_{l}(i+1 \mid \theta)}{P(x \mid \theta)}
\end{aligned}
$$

Estimated \# of $k \rightarrow l$ transitions $\hat{A}_{k, l}$

$$
=\sum_{\text {training seqs } x^{j}} \sum_{i} P\left(\pi_{i}=k, \pi_{i+1}=l \mid x^{j}, \theta\right)
$$

New estimate $\hat{a}_{k, l}=\frac{\hat{A}_{k, l}}{\sum_{l} \hat{A}_{k, l}}$
Emissions: similar

## True Model



Log-odds per roll
True model 0.101 bits 300-roll est. 0.097 bits 30k-roll est. 0.100 Bits (NB: overfitting)


B-W Learned Model


## HMM Summary

Viterbi - best single path
Forward - Sum over all paths
Backward - similar
Baum-Welch - Training via EM and forward/backward (aka the forward/backward algorithm)

Viterbi training - also "EM", but Viterbi-based

## HMMs in Action: Pfam

Proteins fall into families, both across \& within species

Ex: Globins, GPCRs, Zinc Fingers, Leucine zippers,...
Identifying family very useful: suggests function, etc.

So, search \& alignment are both important
One very successful approach: profile HMMs

Helix HBA_HUMAN HBB_HUMAN MYG_PHYCA GLB3_CHITP GLB5_PETMA LGB2__LUPLU GLB1_GLYDI Consensus

AAAAAAAAAAAAAAAA BBBBBBBBBBBBBBBBCCCCCCCCCCC ---------VLSPADKTNVKAAWGKVGA--HAGEYGAEALERMFLSFPTTKTYFPHF --------VHLTPEEKSAVTALWGKV----NVDEVGGEALGRLLVVYPWTQRFFESF ----------VLSEGEWQLVLHVWAKVEA--DVAGHGQDILIRLFKSHPETLEKFDRF ----------LSADQISTVQASFDKVKG------DPVGILYAVFKADPSIMAKFTQF PIVDTGSVAPLSAAEKTKIRSAWAPVYS--TYETSGVDILVKFFTSTPAAQEFFPKF --------GALTESQAALVKSSWEEFNA--NI PKHTHRFFILVLEIAPAAKDLFS-F --------GLSAAQRQVIAATWKDIAGADNGAGVGKDCLIKFLSAHPQMAAVFG-F Ls... v a W kv . . g . L. . f . P . F F

Helix
HBA_HUMAN HBB_HUMAN MYG_PHYCA GLB3_CHITP GLB5_PETMA LGB2_LUPLU GLB1_GLYDI Consensus

Helix
HBA_HUMAN HBB__HUMAN MYG_PHYCA GLB3_CHITP GLB5_PETMA LGB2_LUPLU GLB1_GLYDI Consensus

DDDDDDDEEEEEEEEEEEEEEEEEEEEE
FFFFFFFFFFFF -DLS-----HGSAQVKGHGKKVADALTNAVAHV---D--DMPNALSALSDLHAHKL-GDLSTPDAVMGNPKVKAHGKKVLGAFSDGLAHL---D--NLKGTFATLSELHCDKL-KHLKTEAEMKASEDLKKHGVTVLTALGAILKK----K-GHHEAELKPLAQSHATKH-AG-KDLESIKGTAPFETHANRIVGFFSKIIGEL--P---NIEADVNTFVASHKPRG-K̈GLTTADQLKKSADVRWHAERIINAVNDAVASM--DDTEKMSMKLRDLSGKHAKSF-LK-GTSEVPQNNPELQAHAGKVFKLVYEAAIQLQVTGVVVTDATLKNLGSVHVSKG-SG----AS---DPGVAALGAKVLAQIGVAVSHL--GDEGKMVAQMKAVGVRHKGYGN
t .. . v..Hg kv. a a...l d . a 1.1 H
FFGGGGGGGGGGGGGGGGGGG
HНННннHHHHHHHHHHHHHHHHHHHH -RVDPVNFKLLSHCLLVTLAAHLPAEFTPAVHASLDKFLASVSTVLTSKYR--------HVDPENFRLLGNVLVCVLAHHFGKEFTPPVQAAYQKVVAGVANALAHKYH-------KIPIKYLEFISEAIIHVLHSRHPGDFGADAQGAMNKALELFRKDIAAKYKELGYQG --VTHDQLNNFRAGFVSYMKAHT--DFA-GAEAAWGATLDTFFGMIFSKM

--VADAHFPVVKEAILKTIKEVVGAKWSEELNSAWTIAYDELAIVIKKEMNDAA---KHIKAQYFEPLGASLLSAMEHRIGGKMNAAAKDAWAAAYADISGALISGLQS----v. f 1 . . . .. f . aa. k. . 1 sky

# Alignment of 7 globins. A-H mark 8 alpha helices. Consensus line: upper case $=6 / 7$, lower $=4 / 7$, dot $=3 / 7$. Could we have a profile (aka weight matrix) $\mathrm{w} /$ indels? 

## Profile Hmm Structure



Figure 5.2 The transition structure of a profile HMM.
Mj : Match states ( 20 emission probabilities)
l : $\quad$ Insert states (Background emission probabilities)
$\mathrm{D}_{\mathrm{j}}$ : Delete states (silent - no emission)

## Silent States

Example: chain of states, can skip some


Problem: many parameters.
A solution: chain of "silent" states; fewer parameters (but less detailed control)


Algorithms: basically the same.

## Using Profile HMM's

## Search

Forward or Viterbi
Scoring
Log likelihood (length adjusted)
Log odds vs background
Z scores from either


Alignment
Viterbi

## Likelihood vs Odds Scores




Figure 5.5 To the left the length-normalized LL score is shown as a function of sequence length. The right plot shows the same for the log-odds score.

## Z-Scores



Figure 5.6 The Z-score calculated from the LL scores (left) and the log-odds (right).

## Pfam Model Building

Hand-curated "seed" multiple alignments
Train profile HMM from seed alignment
Hand-chosen score threshold(s)
Automatic classification/alignment of all other protein sequences

7973 families in Rfam 18.0, 8/2005 (covers $\sim 75 \%$ of proteins)

## Model-building refinements

Pseudocounts (count $=0$ common when training with 20 aa's)

$$
e_{i}(a)=\frac{C_{i, a}+A \cdot q_{a}}{\sum_{a} C_{i, a}+A}, \quad A \sim 20, q_{a}=\text { background }
$$ ( $\sim 50$ training sequences)

Pseudocount "mixtures", e.g. separate pseudocount vectors for various contexts (hydrophobic regions, buried regions,...)
(~10-20 training sequences)

## More refinements

Weighting: may need to down weight highly similar sequences to reflect phylogenetic or sampling biases, etc.

Match/insert assignment: Simple threshold, e.g. "> $50 \%$ gap $\Rightarrow$ insert", may be suboptimal. Can use forward-algorithm-like dynamic programming to compute max a posteriori assignment.

## Numerical Issues

Products of many probabilities $\rightarrow 0$
For Viterbi: just add logs
For forward/backward: also work with logs, but you need sums of products, so need "log-of-sum-of-product-of-exp-of-logs", e.g., by table/interpolation

Keep high precision and perhaps scale factor
Working with log-odds also helps.

## Model structure

Define it as well as you can.
In principle, you can allow all transitions and hope to learn their probabilities from data, but it usually works poorly - too many local optima

## Duration Modeling

Self-loop duration: geometric pn(I-p)
min, then geometric

"negative binomial"


More general: possible (but slower)

