CS440/ECE448 Lecture 22: Hidden Markov Models

Mark Hasegawa-Johnson, 3/2023

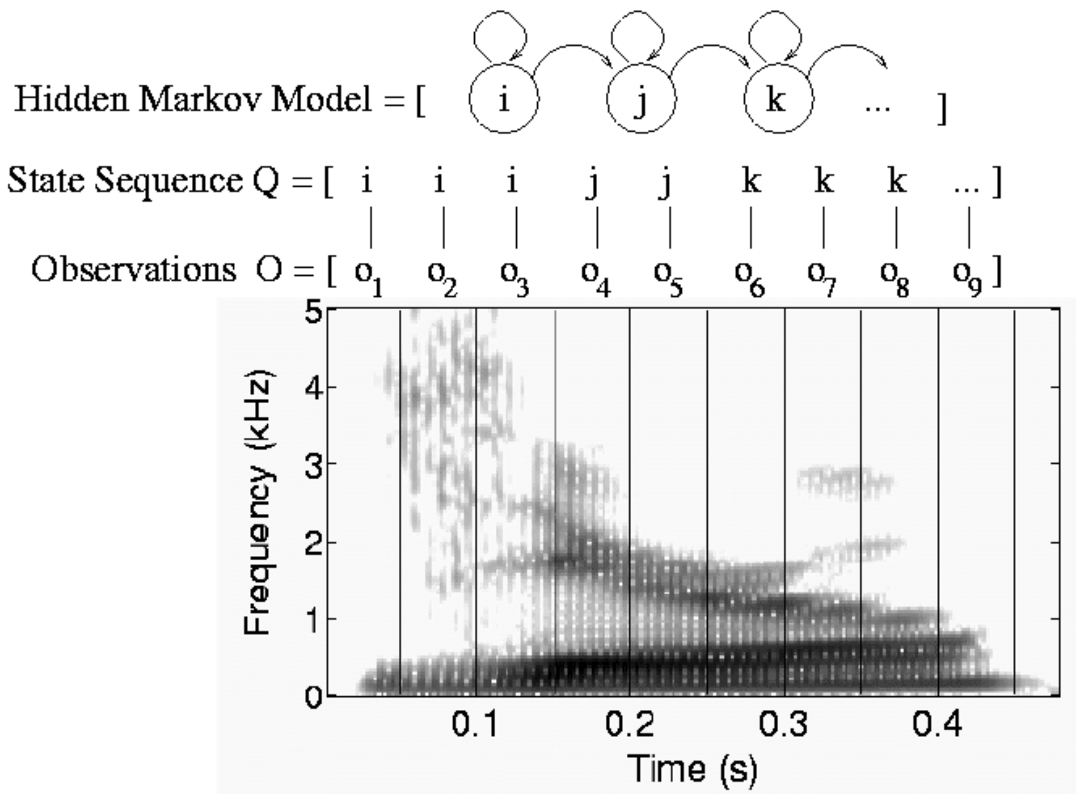

Outline

- Review: Bayesian classifer, Bayesian networks
- HMM: Probabilistic reasoning over time
- Viterbi algorithm

Review: Bayesian Classifier

- Class label $Y = y$, drawn from some set of labels
- Observation $X = x$, drawn from some set of features
- Bayesian classifier: choose the class label, y, that minimizes your probability of making a mistake:

$$\hat{y} = \underset{y}{\operatorname{argmax}}\, P(Y = y|X = x)$$

Bayesian network: A better way to represent knowledge

A Bayesian network is a graph in which:

- Each variable is a node.
- An arrow between two nodes means that the child depends on the parent.
- If the child has no direct dependence on the parent, then there is no arrow.

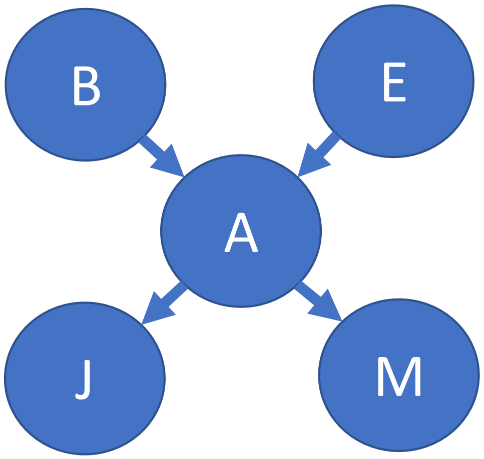

Outline

- Review: Bayesian classifer, Bayesian networks
- HMM: Probabilistic reasoning over time
- Viterbi algorithm

Example: Speech Recognition

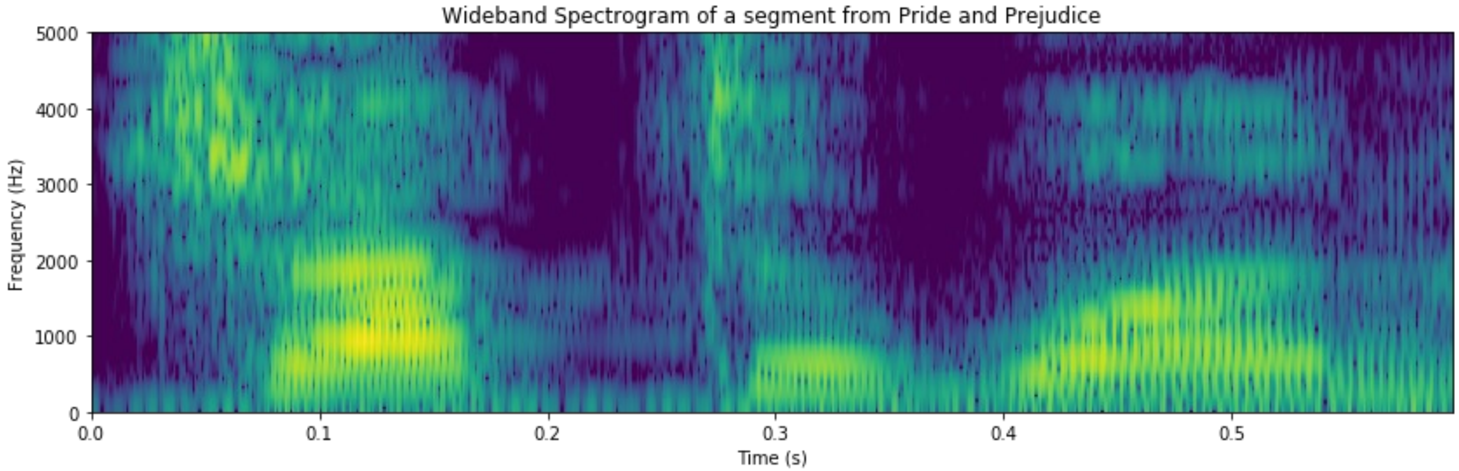

- Here's a spectrogram of the utterance "chapter one."
- Each column is the Fourier transform of 0.02s of audio, spaced 0.01s apart. Let's call the spectral vector X_t, where t is time in centiseconds
- The speech sounds follow a sequence: silence for a while, then /sh/ for a while, then /ae/ for a while, then…. Let's denote the speech sound at time t as Y_t

Hidden Markov Model

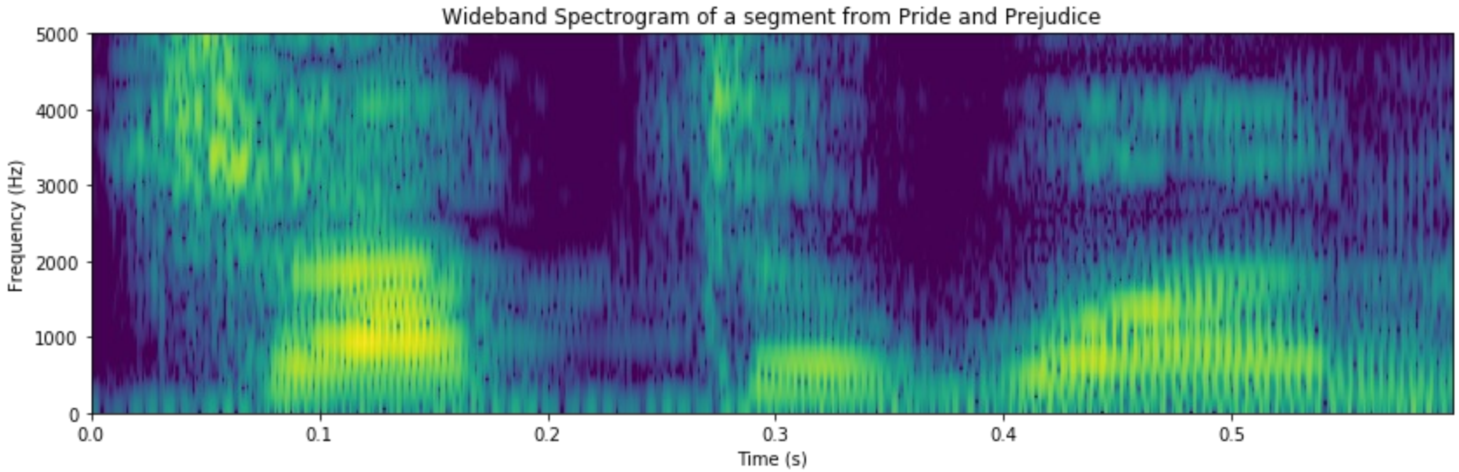

A Hidden Markov Model is a Bayes Network with these assumptions:

- Y_t depends only on Y_{t-1}
- X_t depends only on Y_t

Hidden Markov Model

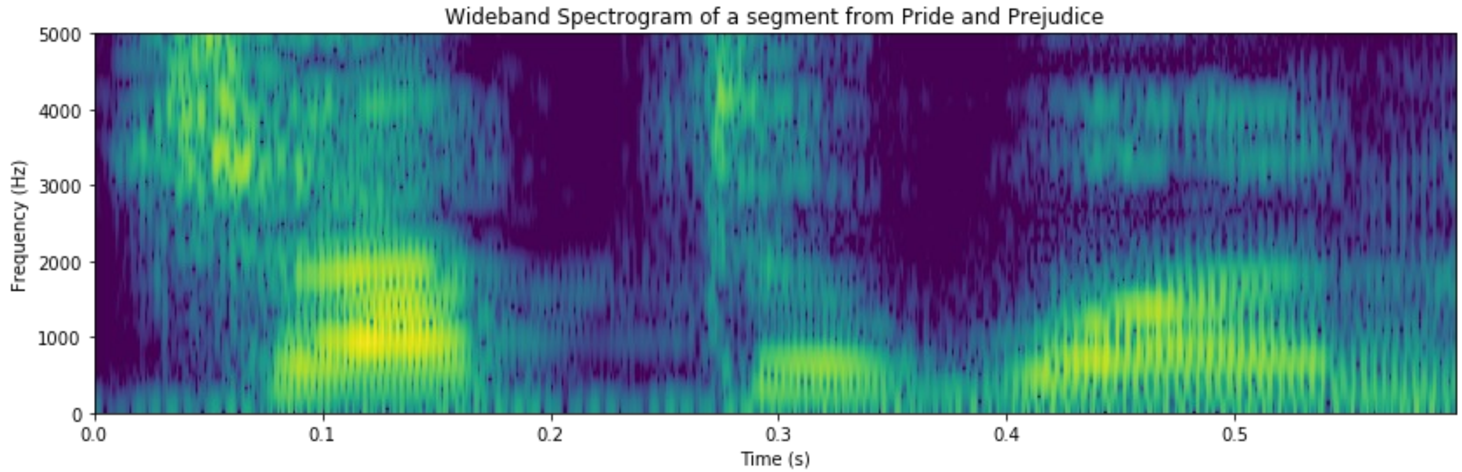

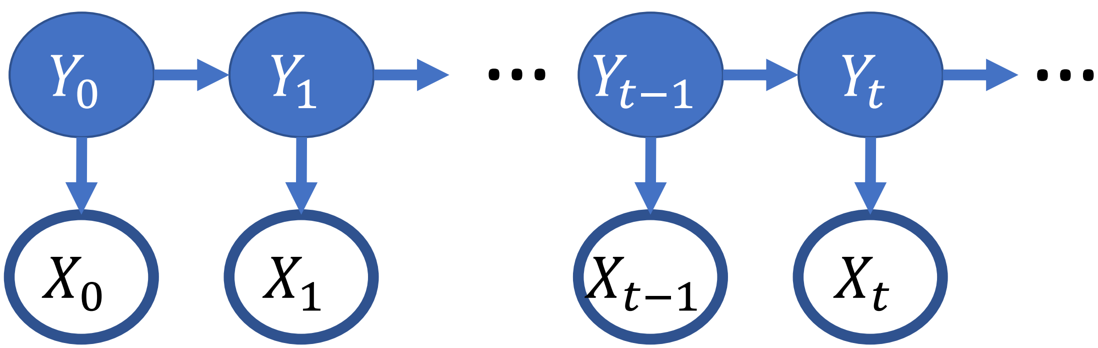

Hidden Markov Model

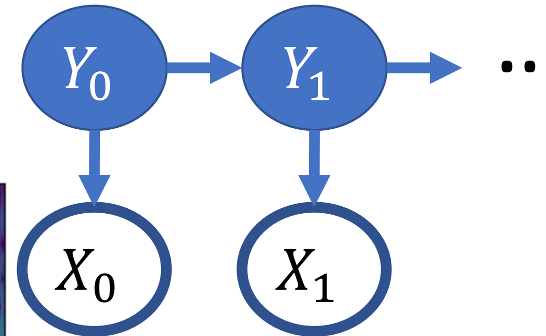

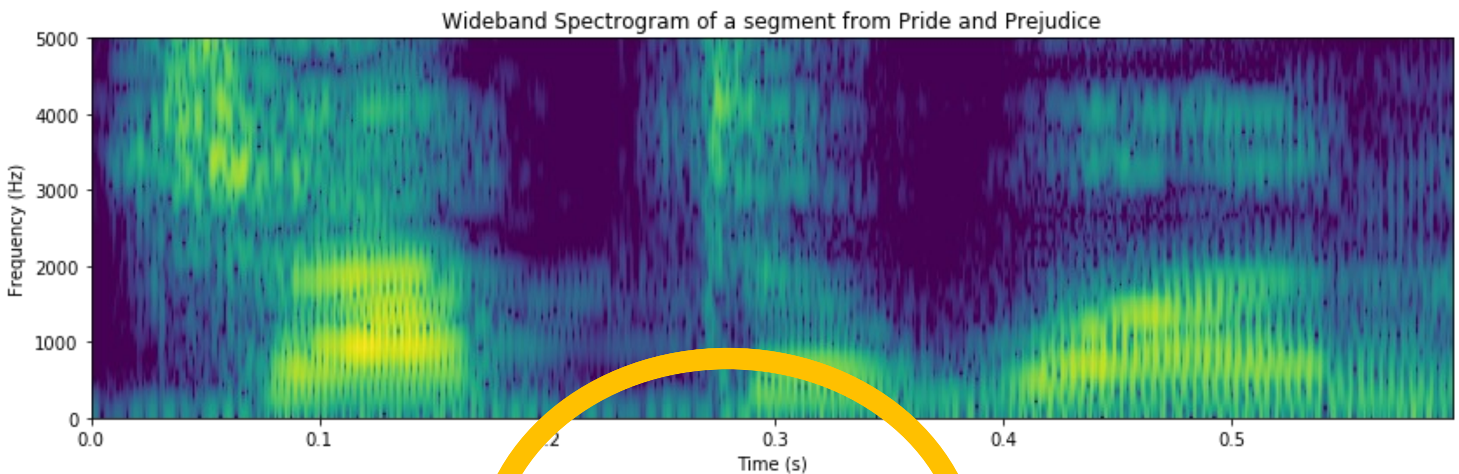

$$P(x_0, \ldots, x_{d-1}) = \sum_{y_0} \sum_{y_1} \cdots \sum_{y_{d-1}} P(y_0) P(x_0|y_0) P(y_1|y_0) P(x_1|y_1) P(y_2|y_1) \cdots$$

$\mathcal{O}\{|\mathcal{Y}|^d\}$ terms in this summation. Does this mean time-complexity is $\mathcal{O}\{|\mathcal{Y}|^d\}$?

Hidden Markov Model

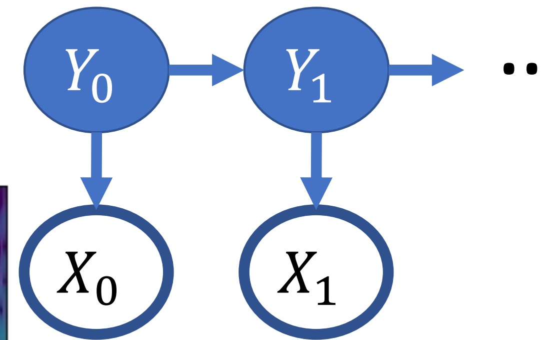

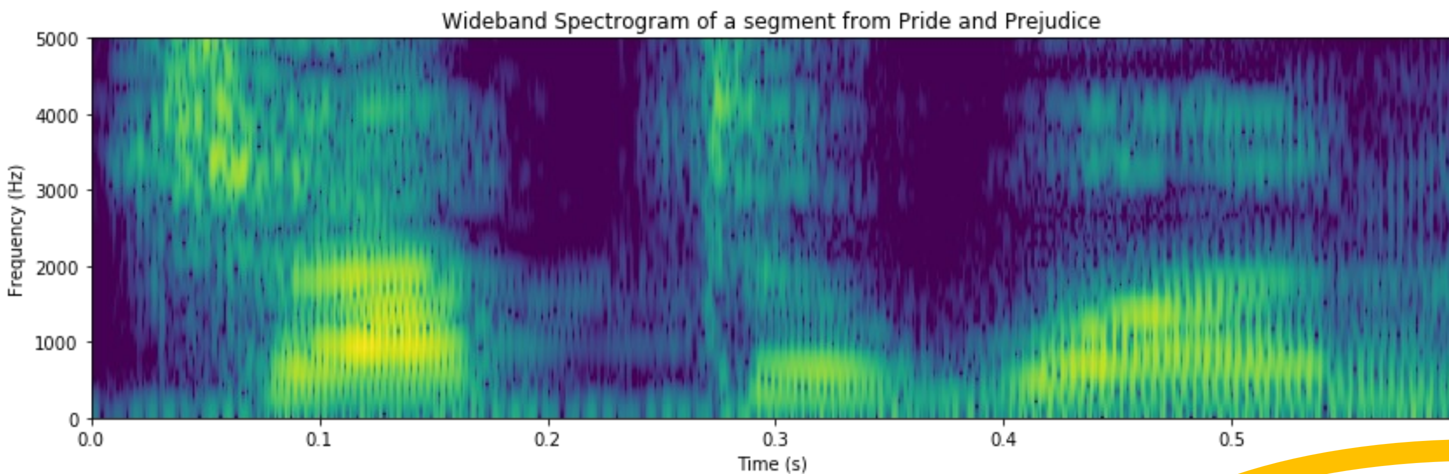

$$P(x_0, \ldots, x_{d-1}) = \sum_{y_{d-1}} \cdots \sum_{y_1} P(y_2|y_1)P(x_1|y_1) \sum_{y_0} P(y_1|y_0)P(x_0|y_0)P(y_0)$$

- There are only $|\mathcal{Y}|$ terms in this summation $(0 \leq y_0 \leq |\mathcal{Y}| - 1)$.
- This summation needs to be computed $|\mathcal{Y}|$ times (once for each value y_1)
- Total complexity: $\mathcal{O}\{|\mathcal{Y}|^2\}$

Two views of an HMM

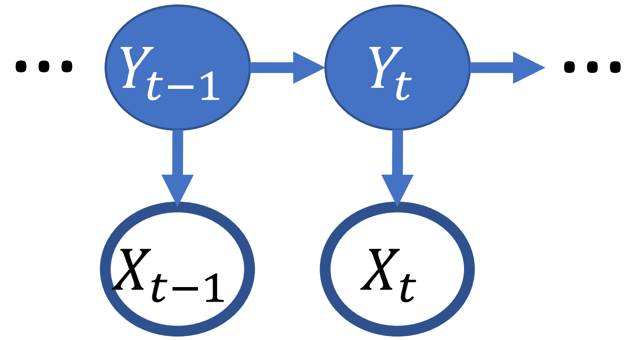

- Bayesian Network diagram shows:
  - Time
  - Y_t depends only on Y_{t-1}
  - X_t depends only on Y_t
- Finite State Machine diagram shows:
  - All of the possible values that Y_t can take
  - The time-complexity of inference is $\mathcal{O}\{|\mathcal{Y}|^2\}$ (because that's the number of edges in the FSM diagram)

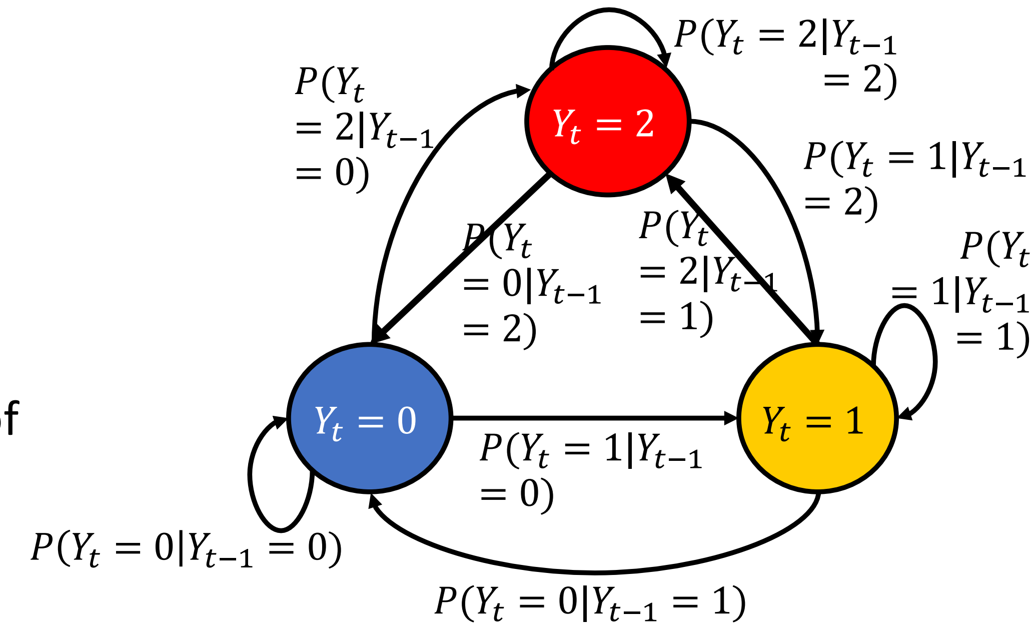

Key advantage of a hidden Markov model

Key advantage: Polynomial-time complexity

- Suppose there are $|\mathcal{Y}|$ different speech sounds in English, and the length of the utterance is d centiseconds ($|\mathcal{Y}| \approx 50, d \approx 100$)
- Without the HMM assumptions, $\text{argmax} P(y_0, \ldots, y_{d-1} | x_0, \ldots, x_{d-1})$ is $\mathcal{O}\{|\mathcal{Y}|^d\}$
- With an HMM, each variable has only one parent, so inference is $\mathcal{O}\{|\mathcal{Y}|^2\}$

…and it works much better than naïve Bayes:

Claude Shannon (1948) gave these examples:

- Text generated by a naïve Bayes model (unigram model):

Representing and speedily is an good apt or come can different natural here he the a in came the to of to expert gray come to furnishes the line message had be these…

- Text generated by a HMM (bigram model):

The head and in frontal attack on an English writer that the character of this point is therefore another for the letters that the time of who ever told the problem for an unexpected…

Applications of HMMs

- Speech recognition HMMs:
  - Observations are acoustic signals (continuous valued)
  - States are specific positions in specific words (so, tens of thousands)

- Machine translation HMMs:
  - Observations are words (tens of thousands)
  - States are cross-lingual alignments

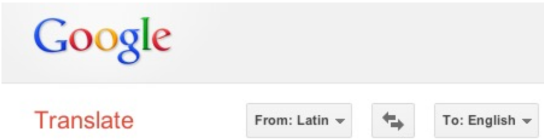

- Robot tracking:
  - Observations are range readings (continuous)
  - States are positions on a map

Source: Tamara Berg

Outline

- Review: Bayesian classifer, Bayesian networks
- HMM: Probabilistic reasoning over time
- Viterbi algorithm

Viterbi Algorithm Key Concept

Given a particular sequence of observations, what is the most likely underlying sequence of states?

In other words, given a particular sequence of observations, what is the most probable path through the trellis?

Example: Speech Recognition

- Observations: X_t = spectrum of 25ms frame of the speech signal.
- State: Y_t = phoneme or letter being currently produced

The goal of speech recognition: find the most probable sequence of text characters $\{y_1, \ldots, y_T\}$ given observations $\{X_1 = x_1, \ldots, X_T = x_T\}$.

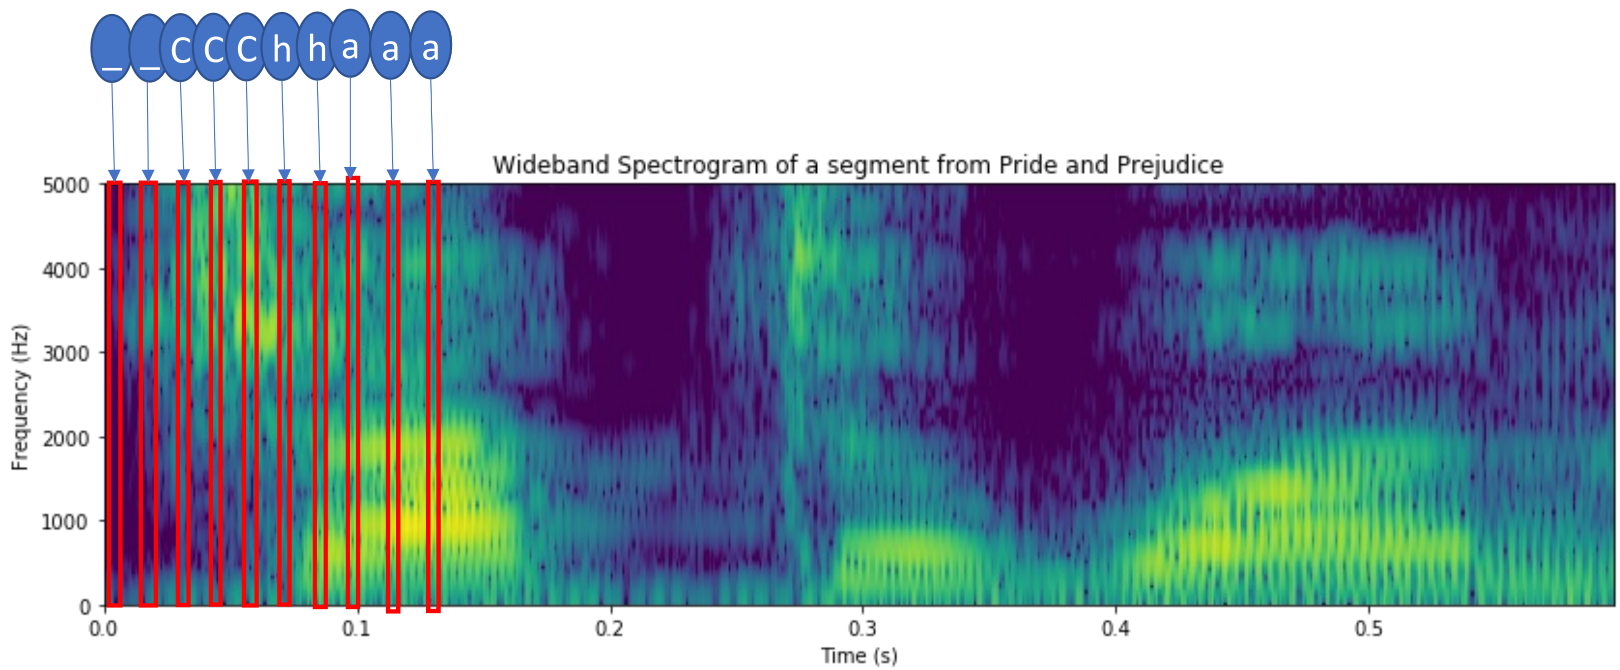

Viterbi Algorithm

Basic concept:

$$\max P(y_0, \dots, y_{d-1}, x_0, \dots, x_{d-1})$$

$$= \max_{y_{d-1}} \cdots \max_{y_1} P(y_2|y_1)P(x_1|y_1) \max_{y_0} P(y_1|y_0)P(x_0|y_0)P(y_0)$$

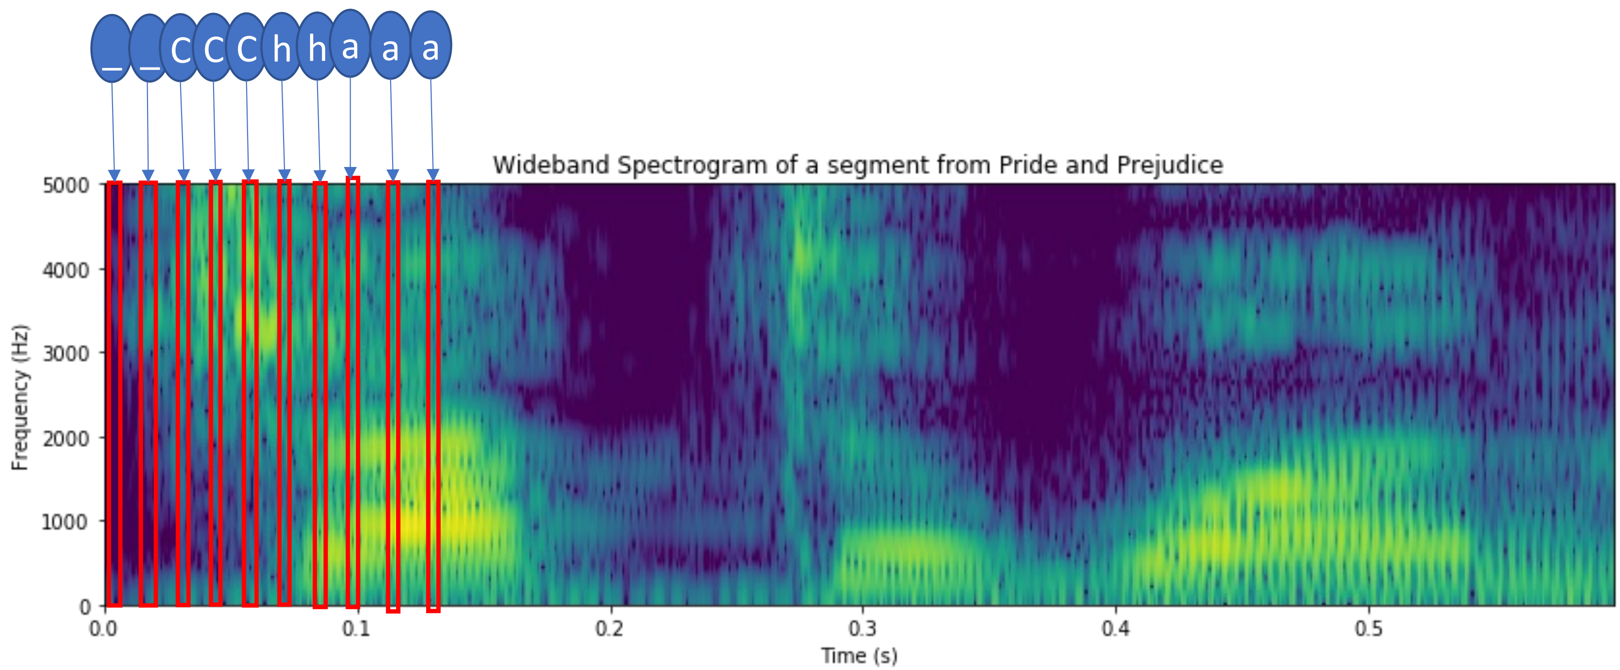

Notation

- Initial State Probability:
$$\pi_i = P(Y_0 = i)$$
- Transition Probabilities:
$$a_{i,j} = P(Y_t = j | Y_{t-1} = i)$$
- Observation Probabilities:
$$b_j(x_t) = P(X_t = x_t | Y_t = j)$$

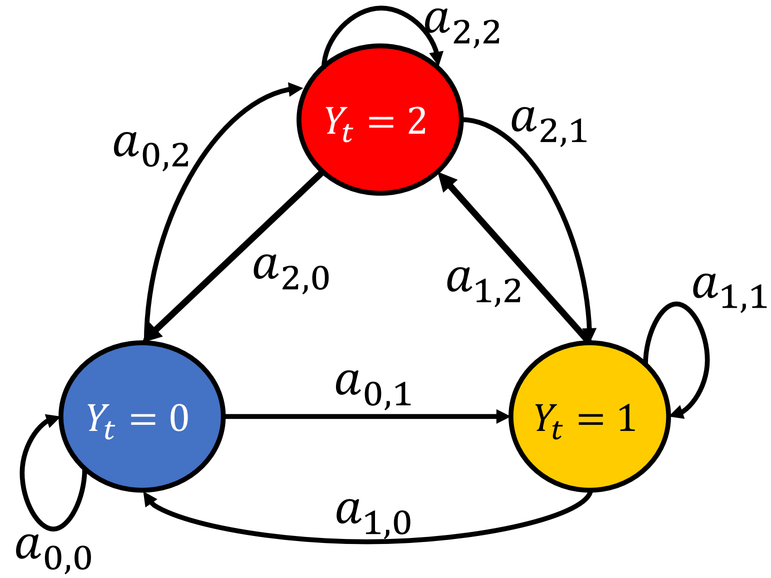

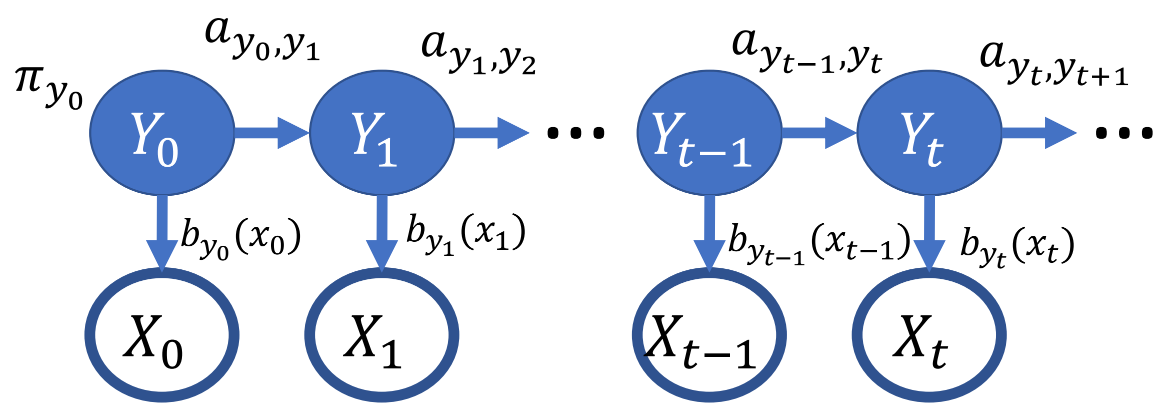

The Trellis

- Time is on the horizontal axis
- State number on the vertical axis:
  $$\pi_i = P(Y_0 = i)$$
- Edges show state transitions: $a_{i,j}$
- States show observation probabilities: $b_j(x_t)$

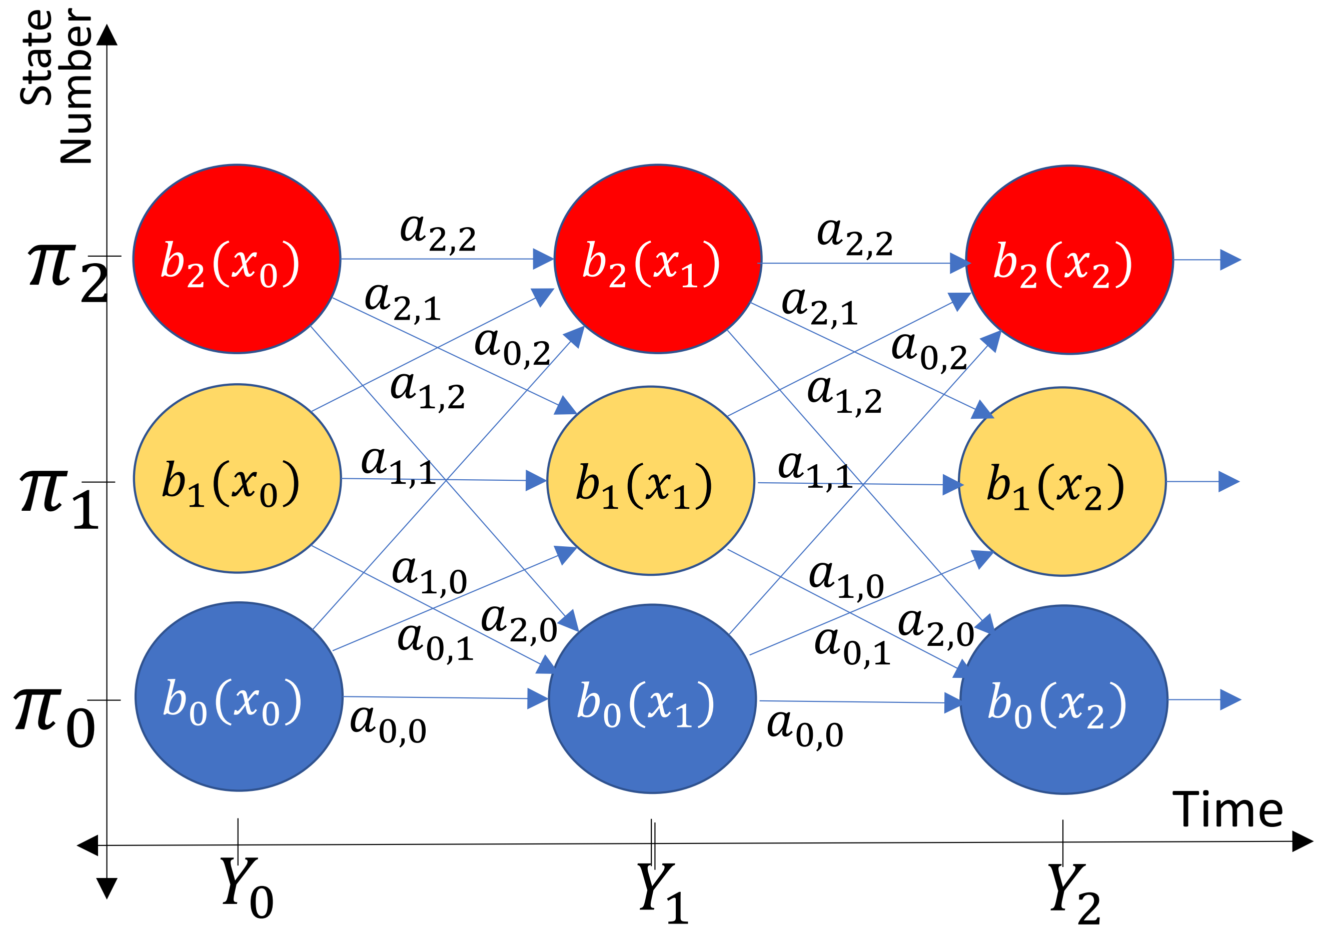

The Trellis

- A sequence of state variables is a path through the trellis.

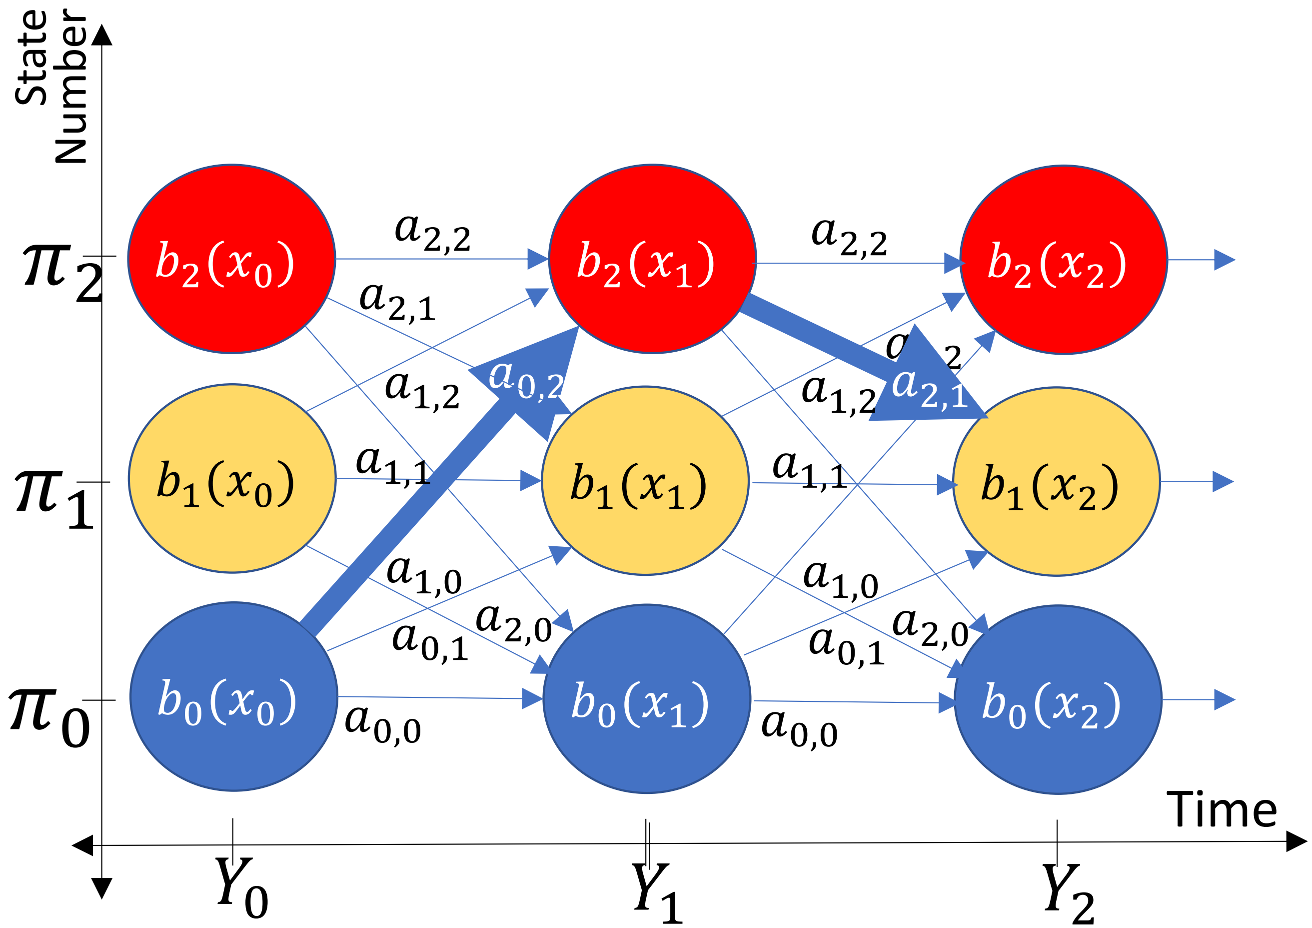

- For example:

$$P(Y_0 = 0, Y_1 = 2, Y_2 = 1, X_0 = x_0, X_1 = x_1, X_2 = x_2|) = \pi_0 b_0(x_0) a_{0,2} b_2(x_1) a_{2,1} b_1(x_2)$$

Node probabilities and backpointers

- **Node Probability**: Probability of the best path until node j at time t

$$v_t(j) = \max_{y_0,\dots,y_{t-1}} P(Y_0 = y_0 \dots, Y_{t-1} = y_{t-1}, Y_t = j, X_0 = x_0, \dots, X_t = x_t)$$

- **Backpointer**: which node, i, precedes node j on the best path?

$$\psi_t(j) = \underset{y_0,\dots,i}{\operatorname{argmax}} P(Y_0 = y_0 \dots, Y_{t-1} = i, Y_t = j, X_0 = x_0, \dots, X_t = x_t)$$

- **B**

Viterbi Algorithm

- Initialization: for $i \in \mathcal{Y}$:
$$v_0(i) = \pi_i b_i(x_0)$$
- Iteration: for $1 \leq t < d$, for $j \in \mathcal{Y}$:
$$v_t(j) = \max_{i \in \mathcal{Y}} v_{t-1}(i)\, a_{i,j} b_j(x_t)$$
$$\psi_t(j) = \operatorname*{argmax}_{i \in \mathcal{Y}} v_{t-1}(i)\, a_{i,j} b_j(x_t)$$
- Termination:
$$y_{d-1} = \operatorname*{argmax}_{i \in \mathcal{Y}} v_{d-1}(i)$$
- Back-Trace:
$$y_t = \psi_{t+1}(y_{t+1})$$

Initialization

$$v_0(i) = \pi_i b_i(x_0)$$

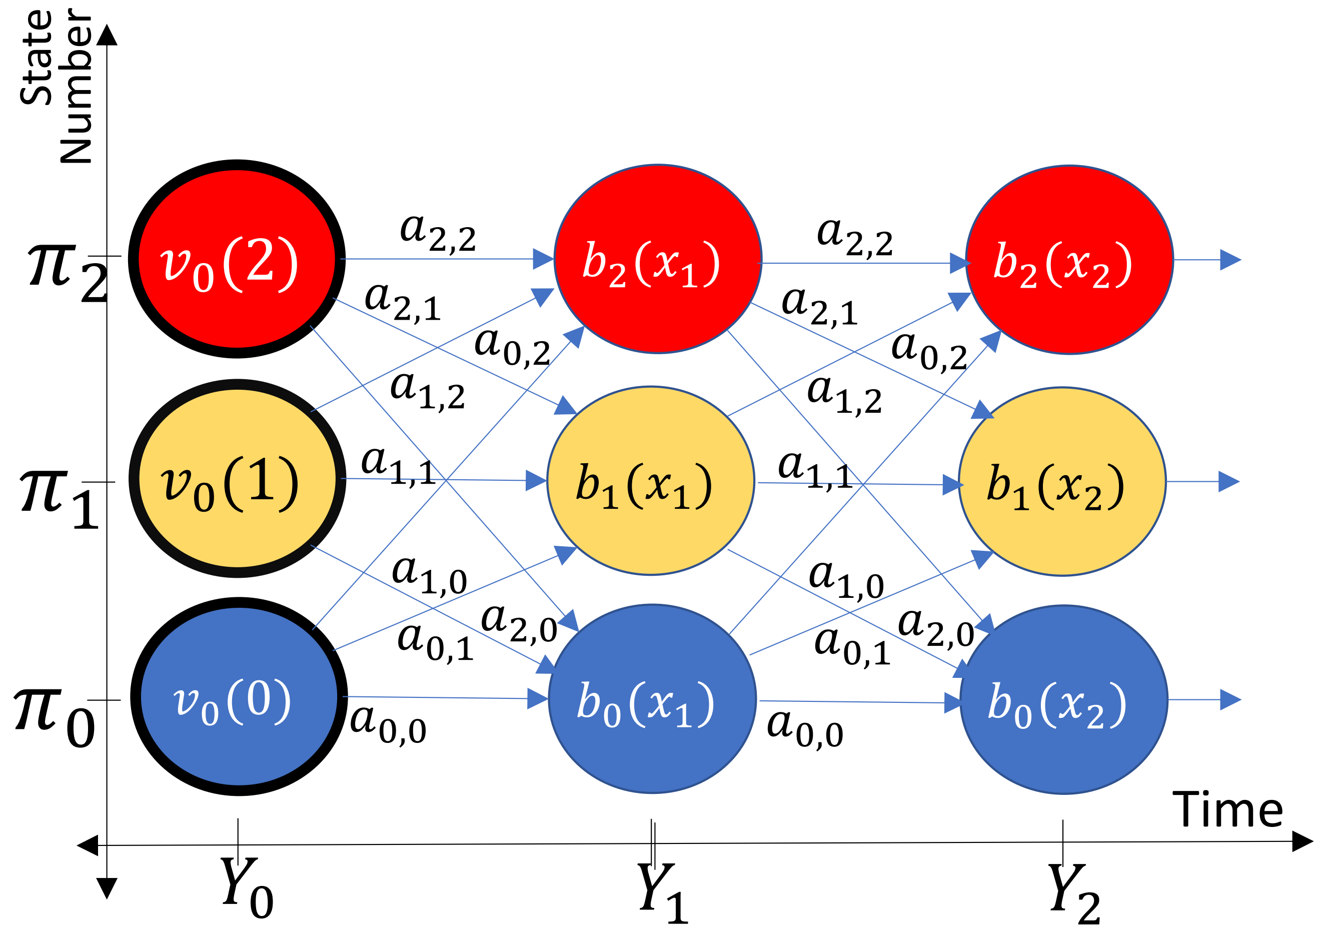

Iteration

Each node now has a value:

$$v_t(j) = \max_{i \in \mathcal{Y}} v_{t-1}(i)\, a_{i,j} b_j(x_t)$$

… and there is exactly one backpointer, from every node, to exactly one node in the previous time step:

$$\psi_t(j) = \underset{i \in \mathcal{Y}}{\operatorname{argmax}}\, v_{t-1}(i)\, a_{i,j} b_j(x_t)$$

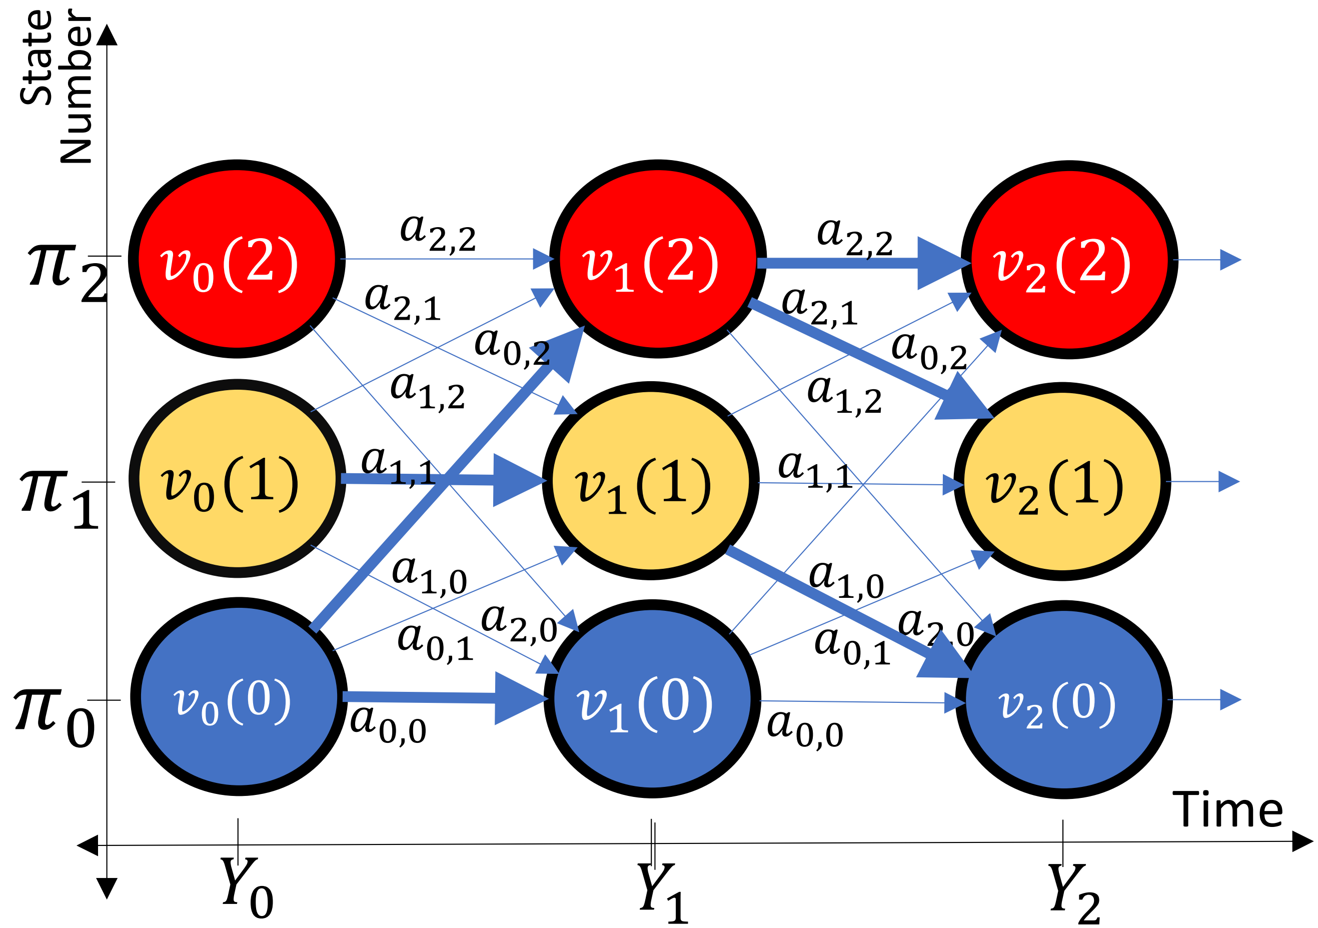

Termination

The best path is the one that ends with the highest-value node:

$$y_{d-1} = \operatorname*{argmax}_{i \in \mathcal{Y}} v_{d-1}(i)$$

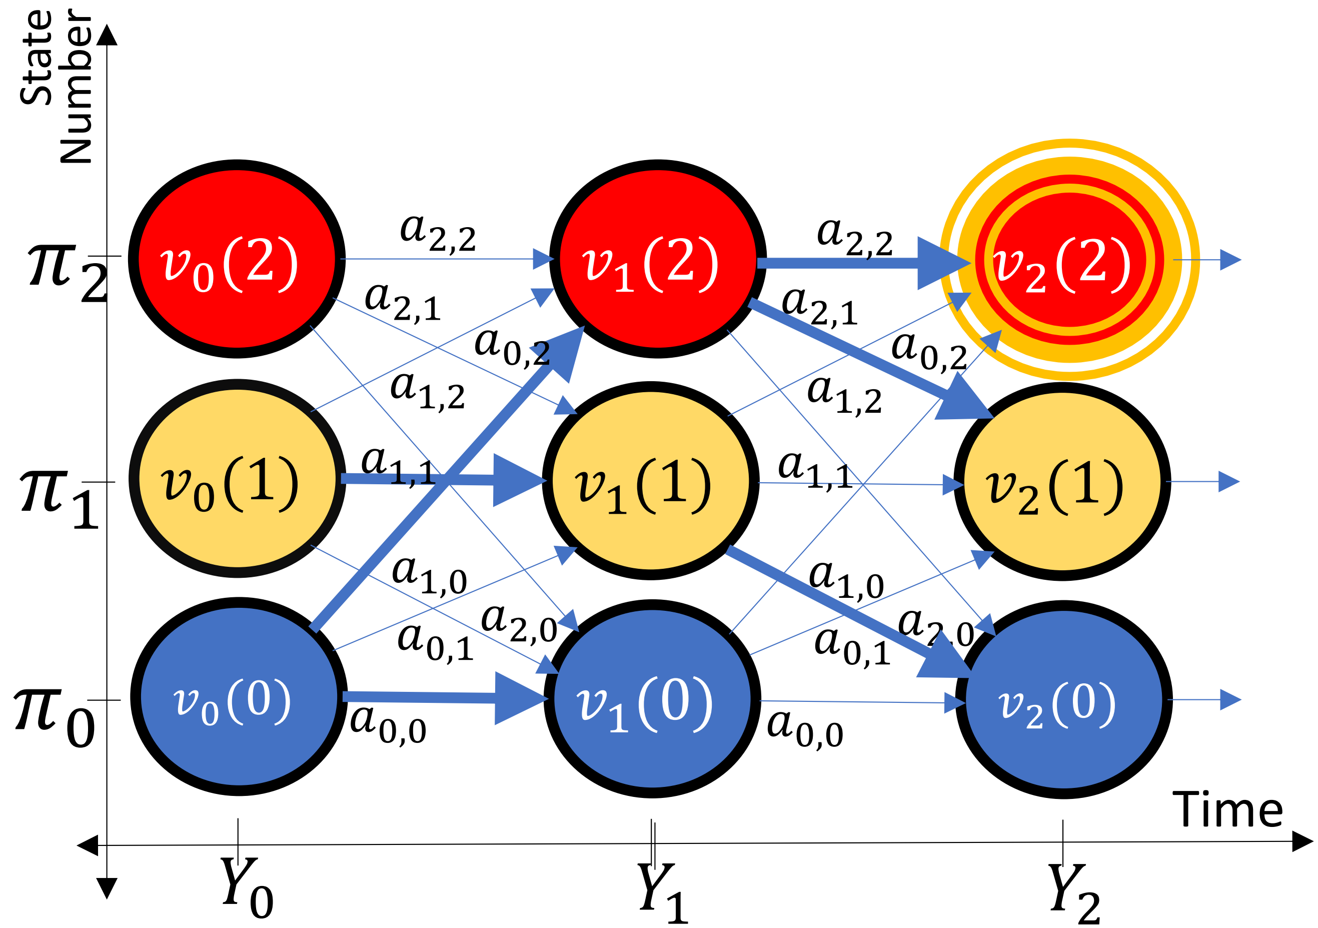

Back-Tracing

The most likely state sequence is the one that ends with the highest-value node:

$$y_t = \psi_{t+1}(y_{t+1})$$

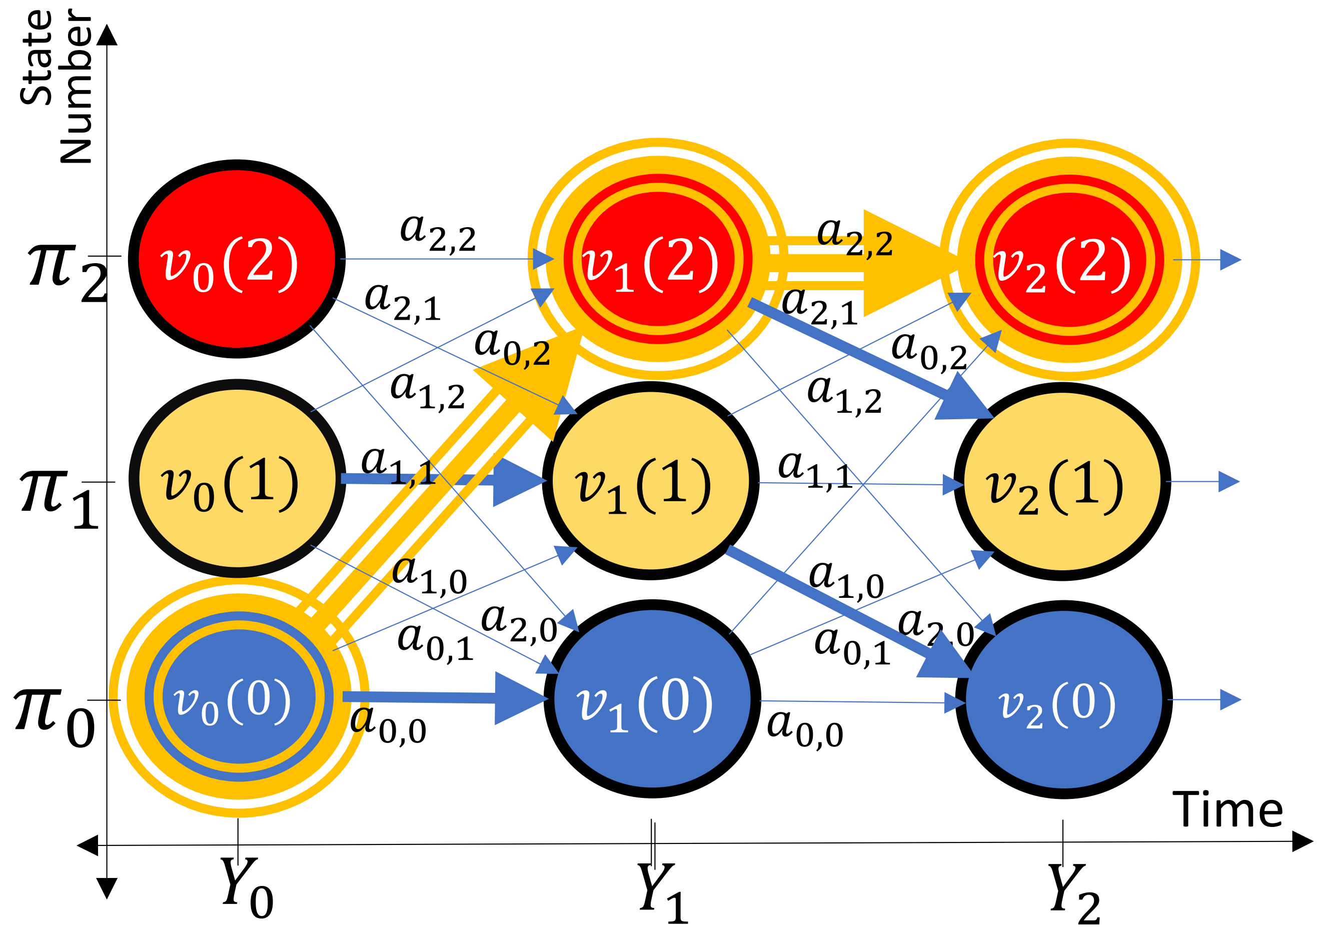

Viterbi Algorithm Computational Complexity

- Initialization: for $i \in \mathcal{Y}$:

$$v_0(i) = \pi_i b_i(x_0)$$

$\mathcal{O}\{|\mathcal{Y}|\}$

- Iteration: for $1 \leq t < d$, for $j \in \mathcal{Y}$:

$$v_t(j) = \max_{i \in \mathcal{Y}} v_{t-1}(i)\, a_{i,j} b_j(x_t)$$

$$\psi_t(j) = \operatorname{argmax}_{i \in \mathcal{Y}} v_{t-1}(i)\, a_{i,j} b_j(x_t)$$

$\mathcal{O}\{d|\mathcal{Y}|^2\}$

- Termination:

$$y_{d-1} = \operatorname{argmax}_{i \in \mathcal{Y}} v_{d-1}(i)$$

$\mathcal{O}\{|\mathcal{Y}|\}$

- Back-Trace:

$$y_t = \psi_{t+1}(y_{t+1})$$

$\mathcal{O}\{d\}$

Total: $\mathcal{O}\{d|\mathcal{Y}|^2\}$

Try the quiz!

- Quiz: https://us.prairielearn.com/pl/course_instance/129874/assessment/2337143

Outline

- Review: Bayesian classifer, Bayesian networks

$$\hat{y} = \underset{y}{\operatorname{argmax}}\, P(Y = y | X = x)$$

- HMM: Probabilistic reasoning over time

$$\pi_i = P(Y_0 = i)$$
$$a_{i,j} = P(Y_t = j | Y_{t-1} = i)$$
$$b_j(x_t) = P(X_t = x_t | Y_t = j)$$

- Viterbi algorithm

$$v_t(j) = \max_{i \in \mathcal{Y}} v_{t-1}(i)\, a_{i,j} b_j(x_t)$$
$$\psi_t(j) = \underset{i \in \mathcal{Y}}{\operatorname{argmax}}\, v_{t-1}(i)\, a_{i,j} b_j(x_t)$$