HMMS IN CONTEXT

JOHN ASHLEY BURGOYNE

MUMT 611

1 FEBRUARY 2007

HMMS: IS THERE MORE?

- HMMs are great, but there are other tools.
- HMMs are a generative model, i.e., they could be used to ‘generate’ new data.
- There are other generative models.
- Sometimes discriminative, i.e., data-defined, models are more appropriate.

OUTLINE

- Generative models
  - Generative model family
  - Grouped HMMs
  - Closest relatives of the HMM
- Discriminative models

BAYESIANS VS. FREQUENTISTS

- Frequentists estimate fixed parameters by maximum likelihood of generating data.
- Bayesians infer the maximum *a posteriori* value of hidden nodes conditioned on the observed data.
- Frequentists require more data; Bayesians require more assumptions.

Generative models

HMM VARIANTS

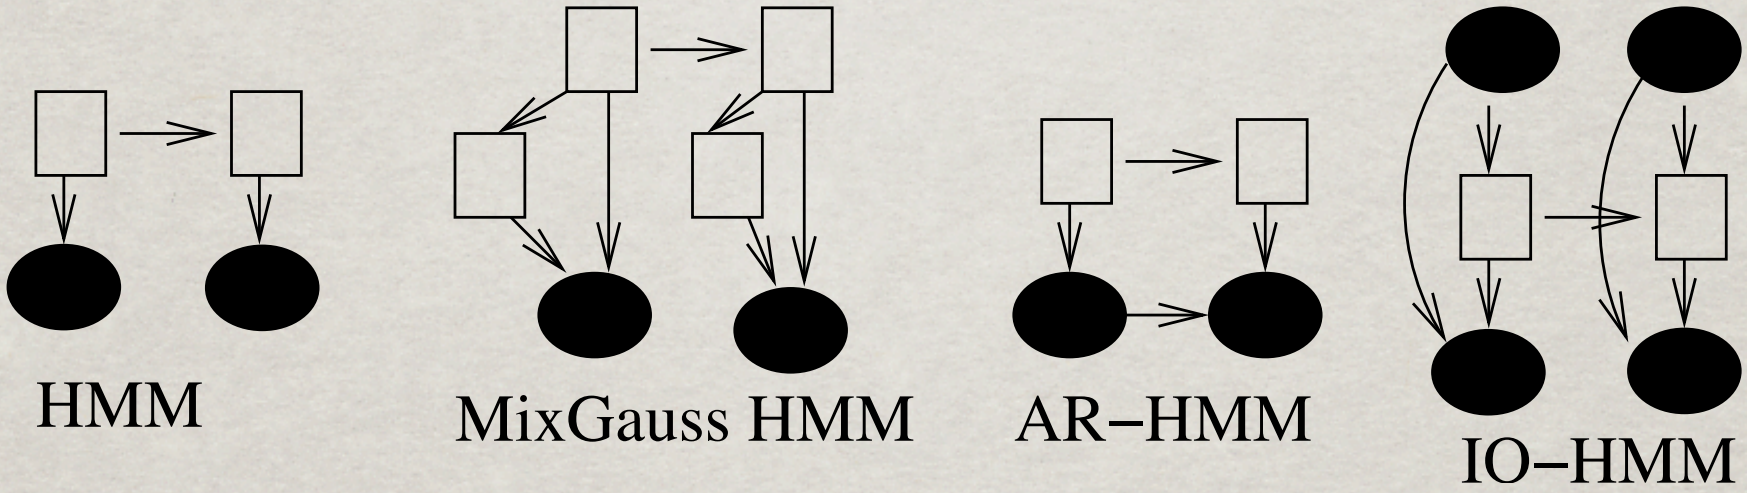

GENERATIVE MODELS

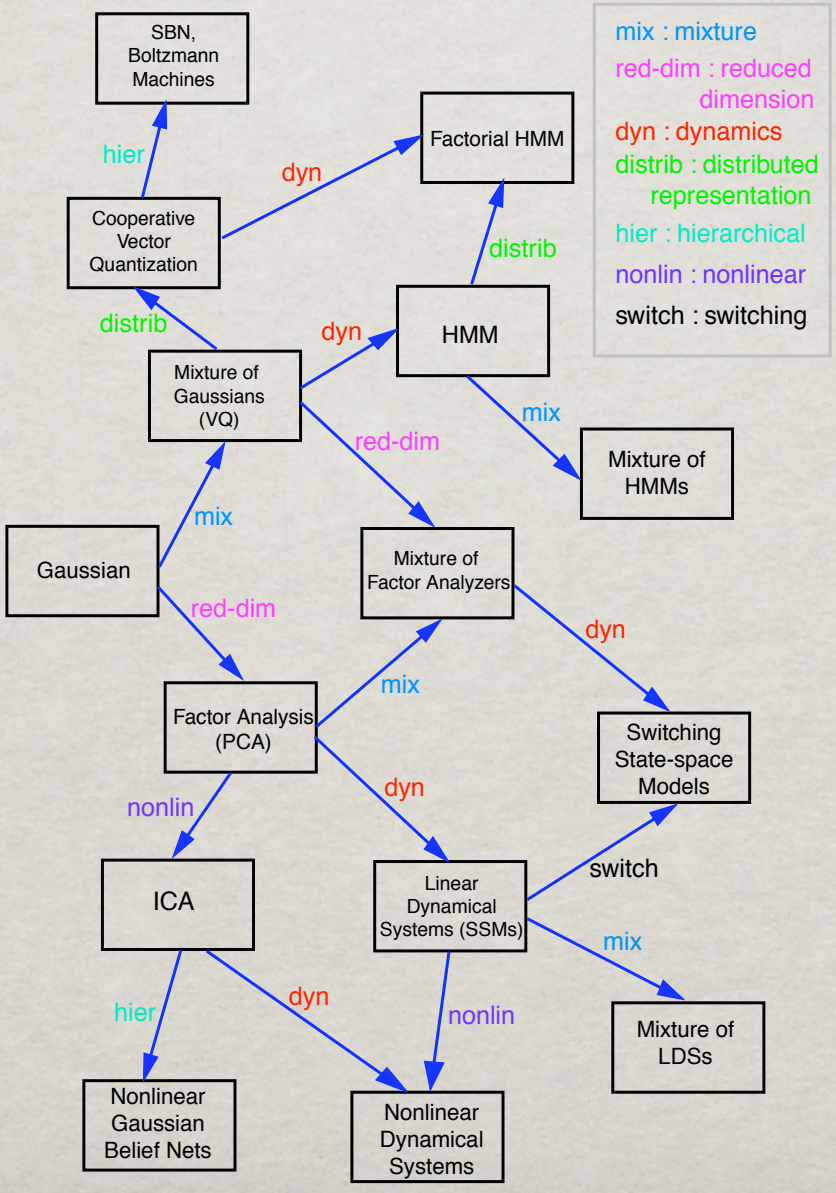

FACTORIAL HMM

- parallel HMMs sharing observations
- combinatorial state explosion
- approximate inference is necessary
- usually perform badly

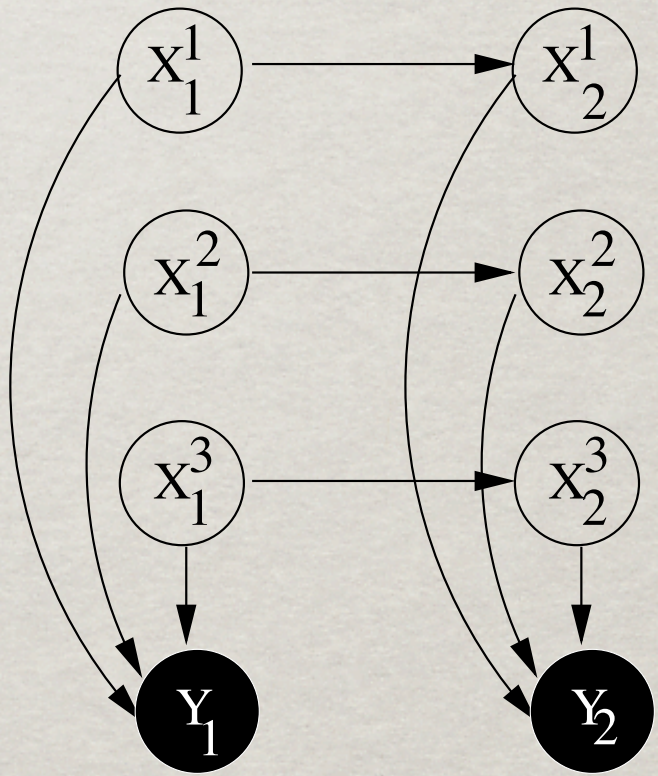

COUPLED HMM

- all present states connect with each other and all future states
- everything depends on everything
- enticing, but very difficult to compute

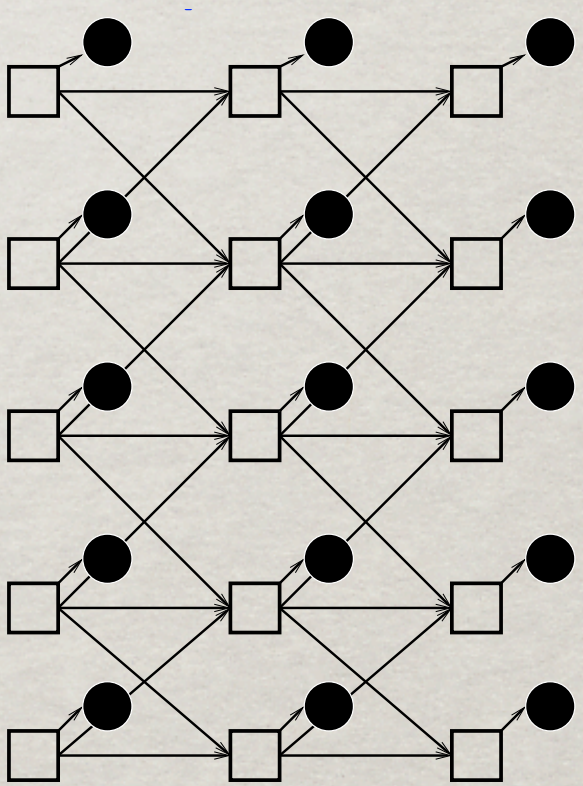

KALMAN FILTER

- HMM with continuous states
- transition matrix becomes a linear transformation
- state becomes a Gaussian (or mixture of Gaussians for more complex variants)
- commonly used in robotics for tracking postion or angle in space

DISCRIMINATIVE MODELS

LABEL-BIAS PROBLEM

- Generative models can only include short time dependencies – which is a problem for music!
- But because of their novelty and difficulty of implementation, discriminative models are not much used in music (yet).

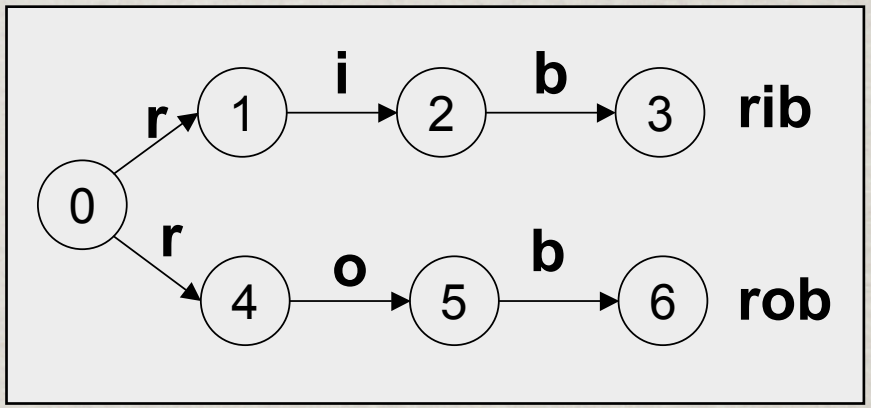

MAXIMUM ENTROPY MARKOV MODELS

- discriminative cousin of the HMM
- turns the observation dependency around
- excellent choice for segmentation
- training is difficult

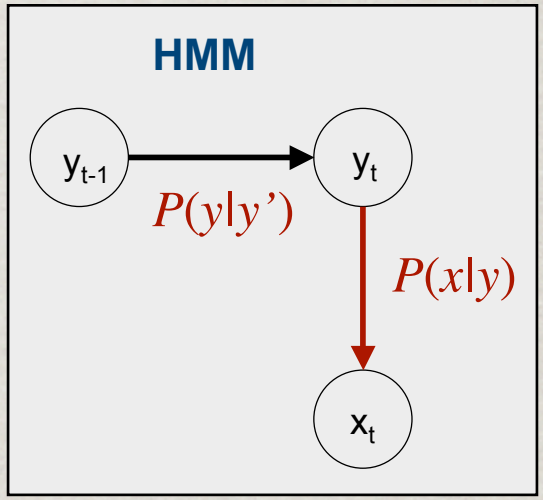

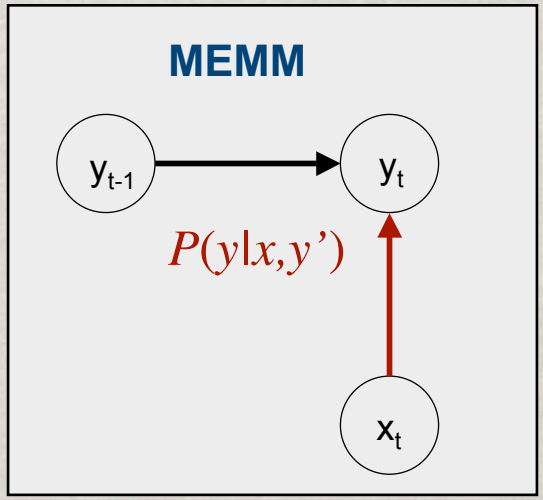

CONDITIONAL RANDOM FIELDS

- MEMM on a Markov random field instead of a Markov chain
- can accept large and disparate sets of observed features
- wildly successful for classification tasks
- difficult to implement and train
- forefront of research in sequence models

CONCLUSION

- Because music is sequential, HMMs often meet our needs as music technologists.
- Sometimes we need other approaches:
  - simpler or richer generative models
  - groups of HMMs
  - discriminative models for classification