

Discriminative Sequence Models and Conditional Random Fields

Mark Gales

Lent 2014



Machine Learning for Language Processing: Lecture 6

MPhil in Advanced Computer Science

MPhil in Advanced Computer Science

Sequence Models


- So far examined the **hidden Markov model** (HMM) as a sequence model
 - generative model of the data sequence, $P(\mathbf{x}_1, \dots, \mathbf{x}_T | q_0, \dots, q_{T+1})$,
 - use Bayes' rule to yield “class sequence” posteriors $P(\mathbf{y} | \mathbf{x}_1, \dots, \mathbf{x}_T)$
 - here $\mathbf{y} = \{y_0, \dots, y_{T+1}\}$ (the states are associated with classes)
- HMM parameters usually trained using **maximum likelihood**
 - possible to also use **discriminative training** criteria to estimate parameters λ
 - **conditional maximum likelihood**, maximise label posterior, $P(\mathbf{y} | \mathbf{x}_1, \dots, \mathbf{x}_T)$

$$\hat{\lambda} = \operatorname{argmax}_{\lambda} \left\{ \sum_{r=1}^R \log \left(\frac{P(\mathbf{y}^{(r)}) P(\mathbf{x}_1^{(r)}, \dots, \mathbf{x}_{T_r}^{(r)} | \mathbf{y}^{(r)}, \lambda)}{\sum_{\mathbf{q} \in \mathcal{Q}_{T_r}} P(\mathbf{q}) P(\mathbf{x}_1^{(r)}, \dots, \mathbf{x}_{T_r}^{(r)} | \mathbf{q}, \lambda)} \right) \right\}$$

- R sequences, labels $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(R)}$
- sequence r is of length T_r , with observations $\mathbf{x}_1^{(r)}, \dots, \mathbf{x}_{T_r}^{(r)}$

What about discriminative sequence models?

Discriminative Sequence Models



- Simple generative model (left) and discriminative model (right)
 - right BN a **maximum entropy Markov model** (q_{T+1} dropped for simplicity)

$$P(q_0, \dots, q_T | \mathbf{x}_1, \dots, \mathbf{x}_T) = \prod_{t=1}^T P(q_t | q_{t-1}, \mathbf{x}_t)$$

state posterior probability given by (Z_t normalisation term at time t)

$$P(q_t | q_{t-1}, \mathbf{x}_t) = \frac{1}{Z_t} \exp \left(\sum_{i=1}^D \lambda_i f_i(q_t, q_{t-1}, \mathbf{x}_t) \right)$$

Sequence Maximum Entropy Models

- State posteriors modelled in the Maximum Entropy Markov model
 - could extend to the complete sequence


$$P(q_0, \dots, q_T | \mathbf{x}_1, \dots, \mathbf{x}_T) = \frac{1}{Z} \exp \left(\sum_{i=1}^D \lambda_i f_i(q_0, \dots, q_T, \mathbf{x}_1, \dots, \mathbf{x}_T) \right)$$

- Problem is that there are a vast number of possible features

What features to extract from the state/observation sequence?

- need to be able to handle variations in length of the sequence
- keep the number of model parameters λ reasonable

(Simple) Linear Chain Conditional Random Fields




- Extract features based on undirected graph
 - conditional independence assumptions similar to HMM (though undirected)

- Posterior model becomes

$$P(q_0, \dots, q_T | \mathbf{x}_1, \dots, \mathbf{x}_T) = \frac{1}{Z} \exp \left(\sum_{t=1}^T \left(\sum_{i=1}^{D_t} \lambda_i^t f_i(q_t, q_{t-1}) + \sum_{i=1}^{D_a} \lambda_i^a f_i(q_t, \mathbf{x}_t) \right) \right)$$

- D_t number of transition style features with parameters λ^t
- D_a number of acoustic style features with parameters λ^a
- This has some relationships to HMMs for particular forms of features (though training different)

Linear Chain Conditional Random Fields



- Extract features based on undirected graph
 - conditional independence assumptions extended to previous state

- Posterior model becomes


$$P(q_0, \dots, q_T | \mathbf{x}_1, \dots, \mathbf{x}_T) = \frac{1}{Z} \exp \left(\sum_{t=1}^T \left(\sum_{i=1}^D \lambda_i f_i(q_t, q_{t-1}, \mathbf{x}_t) \right) \right)$$

- More interesting than HMM-like features
 - features the same as MaxEnt Markov model
 - **BUT** normalised **globally** not **locally**

Normalisation term


- Need to be able to compute the normalisation term efficiently
 - initially consider the **simple linear chain case**

$$Z = \sum_{q \in Q_T} \exp \left(\sum_{t=1}^T \left(\sum_{i=1}^{D_t} \lambda_i^t f_i(q_t, q_{t-1}) + \sum_{i=1}^{D_a} \lambda_i^a f_i(q_t, \mathbf{x}_t) \right) \right)$$



- Consider same topology and observation sequence $\mathbf{x}_1, \dots, \mathbf{x}_7$ as the HMM

Total Path Cost to a State/Time



- Red possible partial paths
- Green state of interest

$$\text{LAdd}(a, b) = \log(\exp(a) + \exp(b))$$

$$\exp(\text{LAdd}(a, b)) = \exp(a) + \exp(b)$$

- Total path cost to state s_i at time t is $\alpha_i(t)$
 - total path cost to state s_4 at time 5 given by (compare to Viterbi)

$$\alpha_4(5) = \text{LAdd} \left(\alpha_3(4) + \sum_{i=1}^{D_t} \lambda_i^t f_i(\mathbf{s}_4, \mathbf{s}_3), \alpha_4(4) + \sum_{i=1}^{D_t} \lambda_i^t f_i(\mathbf{s}_4, \mathbf{s}_4) \right) + \sum_{i=1}^{D_a} \lambda_i^a f_i(\mathbf{s}_4, \mathbf{x}_5)$$



Forward-Backward Algorithm

- α is related to the **forward-"probability"** that is used to train HMMs
 - recursion for this form of model can be expressed as

$$\alpha_j(t) = \log \left(\sum_{k=1}^N \exp \left(\alpha_k(t-1) + \sum_{i=1}^{D_t} \lambda_i^t f_i(\mathbf{s}_j, \mathbf{s}_k) \right) \right) + \sum_{i=1}^{D_a} \lambda_i^a f_i(\mathbf{s}_j, \mathbf{x}_t)$$

- normalisation term can then be expressed as $Z = \exp(\alpha_N(T))$
- There's also a term related to the **backward-"probability"**
 - consider observation at time t **given** state \mathbf{s}_j , $\beta_j(t)$

$$\beta_j(t) = \log \left(\sum_{k=1}^N \exp \left(\beta_k(t+1) + \sum_{i=1}^{D_t} \lambda_i^t f_i(\mathbf{s}_k, \mathbf{s}_j) + \sum_{i=1}^{D_a} \lambda_i^a f_i(\mathbf{s}_k, \mathbf{x}_{t+1}) \right) \right)$$

- designed so that $Z = \sum_{i=1}^N \exp(\alpha_i(t) + \beta_i(t))$



(Aside) HMM-Training using EM

- The forward-backward algorithm used in EM training of HMMs
 - enables latent variable posteriors $P(\mathbf{Z}|\mathbf{X}, \boldsymbol{\lambda})$ to be computed
 - similar form to simple linear chain CRF

$$\sum_{i=1}^{D_t} \lambda_i^t f_i(q_t = \mathbf{s}_j, q_{t-1} = \mathbf{s}_i) : \log(P(q_t = \mathbf{s}_j, q_{t-1} = \mathbf{s}_i)) = \log(a_{ij})$$

$$\sum_{i=1}^{D_a} \lambda_i^a f_i(q_t = \mathbf{s}_j, \mathbf{x}_t) : \log(p(\mathbf{x}_t | q_t = \mathbf{s}_j)) = \log(b_j(\mathbf{x}_t))$$

- (Log) forward $\alpha_j(t)$ and (log) backward probabilities, $\beta_j(t)$:

$$\alpha_j(t) = \log(p(\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = \mathbf{s}_j)) = \log \left(\sum_{k=1}^N a_{kj} \exp(\alpha_k(t-1)) \right) + \log(b_j(\mathbf{x}_t))$$

$$\beta_j(t) = \log(p(\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = \mathbf{s}_j)) = \log \left(\sum_{k=1}^N a_{jk} b_k(\mathbf{x}_{t+1}) \exp(\beta_k(t+1)) \right)$$



(Aside) HMM-Update Formulae

- Forward and backward probabilities can be used to derive posteriors
 - at iteration l

$$\gamma_j^{[l]}(t) = P(q_t = \mathbf{s}_j | \mathbf{x}_1, \dots, \mathbf{x}_T, \boldsymbol{\lambda}^{[l]}) = \exp \left(\alpha_j^{[l]}(t) + \beta_j^{[l]}(t) - \alpha_N^{[l]}(T) \right)$$

- Update formulae with Gaussian state output distribution $b_j(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$

$$\boldsymbol{\mu}_j^{[l+1]} = \frac{\sum_{t=1}^T \gamma_j^{[l]}(t) \mathbf{x}_t}{\sum_{t=1}^T \gamma_j^{[l]}(t)}$$

$$\boldsymbol{\Sigma}_j^{[l+1]} = \frac{\sum_{t=1}^T \gamma_j^{[l]}(t) \mathbf{x}_t \mathbf{x}_t^\top}{\sum_{t=1}^T \gamma_j^{[l]}(t)} - \boldsymbol{\mu}_j^{[l+1]} \boldsymbol{\mu}_j^{[l+1]\top}$$



General Sequence CRFs


- The general form of CRF uses an undirected graphical model to define features
 - need to be able to handle sequence data - **dynamic CRF**
 - undirected graph **repeated** each time instance - set of cliques is \mathcal{C}
- The posterior probability for this form of model is

$$P(q_0, \dots, q_T | \mathbf{x}_1, \dots, \mathbf{x}_T) = \frac{1}{Z} \exp \left(\sum_{t=1}^T \sum_{\mathcal{C} \in \mathcal{C}} \lambda_{\mathcal{C}}^T \mathbf{f}(q_{\mathcal{C}t}, \mathbf{x}_1, \dots, \mathbf{x}_T, t) \right)$$

- $\lambda_{\mathcal{C}}^T$ **time-independent** parameters associated with clique \mathcal{C}
- $\mathbf{f}(q_{\mathcal{C}t}, \mathbf{x}_1, \dots, \mathbf{x}_T, t)$ **time-dependent** features extracted from clique \mathcal{C} with **time-dependent** label sequence $q_{\mathcal{C}t}$



Example of a Sequence CRF



- Cliques associated with linear CRF

$$\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2\}$$

1. **transitions:** $\mathcal{C}_1 = \{q_t, q_{t-1}\}$
2. **acoustics:** $\mathcal{C}_2 = \{q_t, \mathbf{x}_t\}$

- Posterior model for the simple linear chain CRF

$$\begin{aligned} P(q_0, \dots, q_T | \mathbf{x}_1, \dots, \mathbf{x}_T) &= \frac{1}{Z} \exp \left(\sum_{t=1}^T \sum_{\mathcal{C} \in \mathcal{C}} \boldsymbol{\lambda}_{\mathcal{C}}^{\top} \mathbf{f}(q_{\mathcal{C}t}, \mathbf{x}_1, \dots, \mathbf{x}_T, t) \right) \\ &= \frac{1}{Z} \exp \left(\sum_{t=1}^T (\boldsymbol{\lambda}^{\text{t}} \mathbf{f}(q_t, q_{t-1}) + \boldsymbol{\lambda}^{\text{a}} \mathbf{f}(q_t, \mathbf{x}_t)) \right) \end{aligned}$$

Training CRFs

- Training for CRFs is normally fully observed

training observation sequence $\mathbf{x}_1, \dots, \mathbf{x}_T$
 training label sequence y_1, \dots, y_T

– where $y_\tau \in \{\omega_1, \dots, \omega_K\}$

- No need to use EM (or related approaches)
 - extension to CRFs includes additional latent variables **hidden CRFs**
 - training data for HCRFs only **partially observed**
- Need to find the model parameters λ so that

$$\begin{aligned}
 \hat{\lambda} &= \operatorname{argmax}_{\lambda} \{P(y_1, \dots, y_T | \mathbf{x}_1, \dots, \mathbf{x}_T, \lambda)\} \\
 &= \operatorname{argmax}_{\lambda} \left\{ \frac{1}{Z} \exp \left(\sum_{i=1}^D \lambda_i f_i(\mathbf{x}_1, \dots, \mathbf{x}_T, y_1, \dots, y_T) \right) \right\}
 \end{aligned}$$



Generalised Iterative Scaling for CRFs

- CRF (also MaxEnt model) training is a **convex optimisation** problem
 - one solution to train parameters is **generalised iterative scaling**

$$\lambda_i^{[k+1]} = \lambda_i^{[k]} + \frac{1}{C} \log \left(\frac{f_i(\mathbf{x}_1, \dots, \mathbf{x}_T, y_1, \dots, y_T)}{\sum_{\mathbf{q} \in \mathcal{Q}_T} P(\mathbf{q} | \mathbf{x}_1, \dots, \mathbf{x}_T, \boldsymbol{\lambda}^{[k]}) f_i(\mathbf{x}_1, \dots, \mathbf{x}_T, \mathbf{q})} \right)$$

- iterative approach (parameters at iteration k are $\boldsymbol{\lambda}^{[k]}$)
- (strictly) requires that the features add up to a constant

$$\sum_{i=1}^D f_i(\mathbf{x}_1, \dots, \mathbf{x}_T, \mathbf{q}) = C, \quad \forall \mathbf{q} \in \mathcal{Q}_T$$

- extensions relaxes this requirements, e.g. **improved iterative scaling**



Inference with CRFs

- Recognition with CRFs involves finding the most probable label sequence \hat{q}

$$\begin{aligned}\hat{q} &= \operatorname{argmax}_{\mathbf{q} \in \mathcal{Q}_T} \{P(\mathbf{q} | \mathbf{x}_1, \dots, \mathbf{x}_T)\} \\ &= \operatorname{argmax}_{\mathbf{q} \in \mathcal{Q}_T} \left\{ \sum_{i=1}^D \lambda_i f_i(\mathbf{x}_1, \dots, \mathbf{x}_T, \mathbf{q}) \right\}\end{aligned}$$

– normalisation term Z not used as it is the same for **all** label sequences

- The **Viterbi algorithm** is often used to perform recognition
 - for the simple linear chain CRF relationship to HMM Viterbi clear:

$$\hat{q} = \operatorname{argmax}_{\mathbf{q} \in \mathcal{Q}_T} \left\{ \sum_{t=1}^T \left(\sum_{i=1}^{D_t} \lambda_i^t f_i(q_t, q_{t-1}) + \sum_{i=1}^{D_a} \lambda_i^a f_i(q_t, \mathbf{x}_t) \right) \right\}$$

