

Neural networks

Training CRFs - discriminative vs. generative learning

GENERATIVE VS. DISCRIMINATIVE


Topics: discriminative learning, generative learning

- In discriminative learning, we optimize the conditional likelihood $-\log p(\mathbf{y}|\mathbf{X})$
 - ▶ CRFs are discriminative
- In generative learning, we optimize the joint log-likelihood:
$$-\log p(\mathbf{y}, \mathbf{X}) = -\log (p(\mathbf{y}|\mathbf{X})p(\mathbf{X})) = -\log p(\mathbf{y}|\mathbf{X}) - \log p(\mathbf{X})$$
 - ▶ HMMs are usually trained generatively
 - ▶ $-\log p(\mathbf{X})$ is similar to a regularizer

GENERATIVE VS. DISCRIMINATIVE

Topics: generative learning, discriminative learning

- It can be shown that:
 - ▶ if model is well-specified (i.e. is the true model) generative learning is better

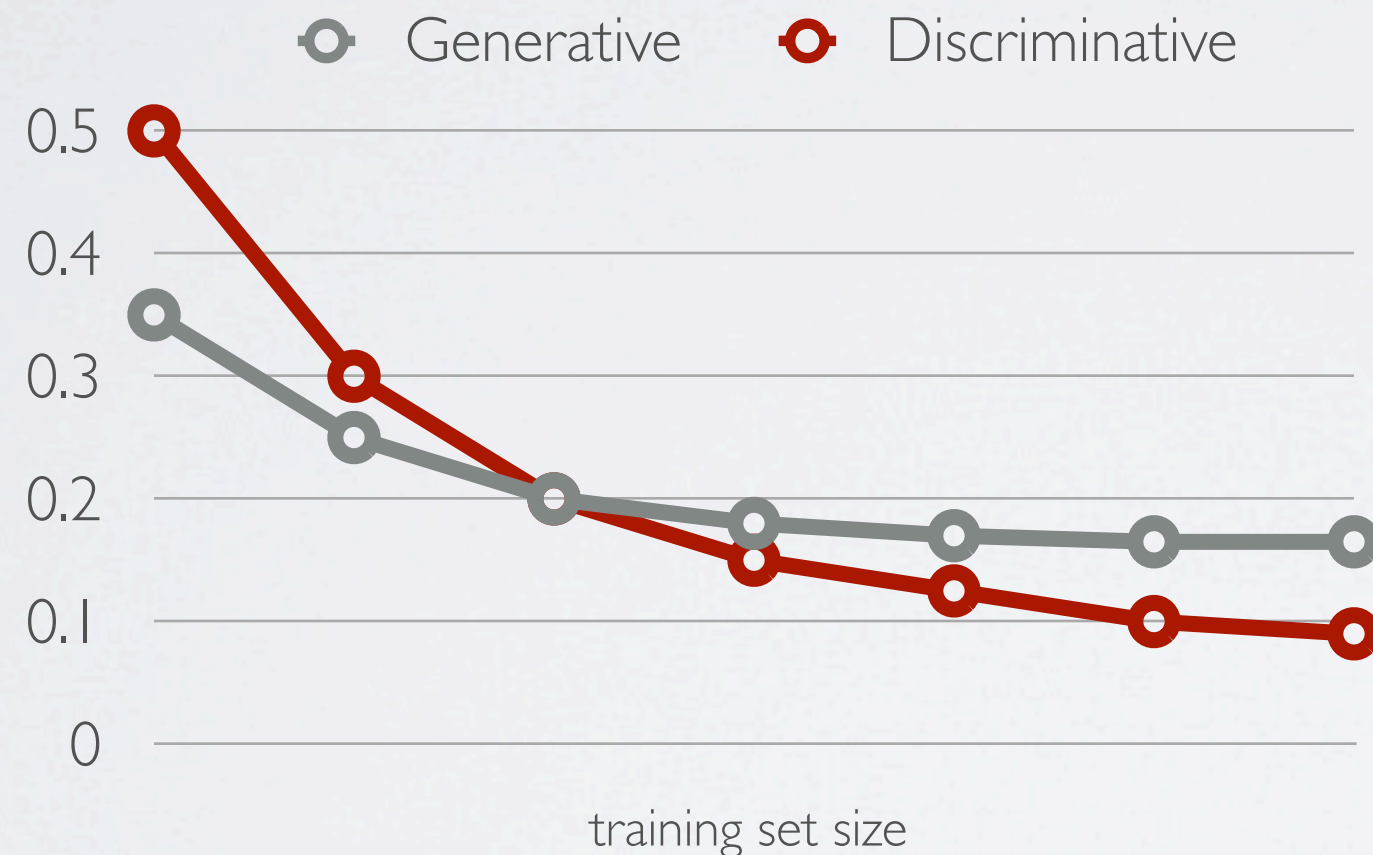


GENERATIVE VS. DISCRIMINATIVE

Topics: generative learning, discriminative learning

- It can be shown that:

- ▶ if model is not well-specified (i.e. most of the time), it depends:



- ▶ See these papers for more details:

- On Discriminative vs. Generative classifiers: A comparison of logistic regression and naive Bayes. Andrew Ng and Michael Jordan, 2001