Statistical Machine Learning https://cvml.ist.ac.at/courses/SML_W18

Christoph Lampert

# Institute of Science and Technology 

Winter Semester 2018/2019
Lecture 9
(lots of material courtesy of S. Nowozin, http://www.nowozin.net)

## Standard Regression/Classification:

$$
f: \mathcal{X} \rightarrow \mathbb{R}
$$

- inputs $\mathcal{X}$ can be any kind of objects
- output $y \in \mathcal{Y}$ is a number (real or integer)


## Structured Prediction:

$$
f: \mathcal{X} \rightarrow \mathcal{Y}
$$

- inputs $\mathcal{X}$ can be any kind of objects
- outputs $y \in \mathcal{Y}$ are complex (structured) objects


## What is structured data?

Ad hoc definition: data that consists of several parts, and not only the parts themselves contain information, but also the way in which the parts belong together.

Text


Documents/HyperText


Molecules / Chemical Structures


Images

## What is structured output prediction?

Ad hoc definition: predicting structured outputs from input data (in contrast to predicting just a single number, like in classification or regression)

- Natural Language Processing:
- Automatic Translation (output: sentences)
- Bioinformatics:
- Secondary Structure Prediction (output: bipartite graphs)
- Speech Processing:
- Text-to-Speech (output: audio signal)
- Robotics:
- Planning (output: sequence of actions)
- Information Retrieval:
- Ranking (output: ordered list of documents)

This lecture: mainly examples from Computer Vision

## Example: Human Pose Estimation



$$
x \in \mathcal{X}
$$


$y \in \mathcal{Y}$

- Given an image, where is a person and how is it articulated?

$$
f: \mathcal{X} \rightarrow \mathcal{Y}
$$

- Image $x$, but what is $y \in \mathcal{Y}$ precisely?


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$


## Example: Human Pose Estimation



Example $y_{\text {head }}$

- Body Part: $y_{\text {head }}=(u, v, \theta)$ where $(u, v)$ center, $\theta$ rotation - $(u, v) \in\{1, \ldots, M\} \times\{1, \ldots, N\}, \theta \in\left\{0,45^{\circ}, 90^{\circ}, \ldots\right\}$
- same for torso, left arm, right arm, ...
- Entire Body: $y=\left(y_{\text {head }}, y_{\text {torso }}, y_{\text {left-lower-arm }}, \ldots\right) \in \mathcal{Y}$


## Example: Human Pose Estimation



- Idea: Have a head detector (CNN, SVM, RF, ...)

$$
f_{\text {head }}: \mathcal{X} \rightarrow \mathbb{R}
$$

## Example: Human Pose Estimation



- Idea: Have a head detector (CNN, SVM, RF, ...)

$$
f_{\text {head }}: \mathcal{X} \rightarrow \mathbb{R}
$$

- Evaluate for every possible location and record score
- Same construction for all other body parts


## Example: Human Pose Estimation



## Image $x \in \mathcal{X}$

- Put together body from individual parts

$$
y^{\text {best }}=\left(y_{\text {head }}^{\text {best }}, y_{\text {torso }}^{\text {best }}, \cdots\right)
$$

## Example: Human Pose Estimation



Image $x \in \mathcal{X}$


Prediction $y^{\text {best }} \in \mathcal{Y}$

- Put together body from individual parts

$$
y^{\text {best }}=\left(y_{\text {head }}^{\text {best }}, y_{\text {torso }}^{\text {best }}, \cdots\right)
$$

- Each part looks reasonable, but overall makes no sense


## Example: Human Pose Estimation



Image: Ben Sapp

## Enforce relations between parts

- For example, head must be connected to torso
- Problem:

$$
y^{\text {best }} \neq\left(y_{\text {head }}^{\text {best }}, y_{\text {torso }}^{\text {best }}, \cdots\right)
$$

independent decisions for each body part are not optimal anymore

## Example: Human Pose Estimation



## Enforce relations between parts

- For example, head must be connected to torso
- Problem:

$$
y^{\text {best }} \neq\left(y_{\text {head }}^{\text {best }}, y_{\text {torso }}^{\text {best }}, \cdots\right)
$$

independent decisions for each body part are not optimal anymore

- Needs structured output prediction function $f: \mathcal{X} \rightarrow \mathcal{Y}$


## The general recipe

Normal prediction function, $\mathcal{X}=$ anything, $\mathcal{Y}=\mathbb{R}$
Extract feature vector from $x$ and compute a number from it

$$
\text { e.g. } \quad f(x)=\langle w, \phi(x)\rangle+b
$$

## The general recipe

## Normal prediction function, $\mathcal{X}=$ anything, $\mathcal{Y}=\mathbb{R}$

Extract feature vector from $x$ and compute a number from it

$$
\text { e.g. } \quad f(x)=\langle w, \phi(x)\rangle+b
$$

Structured output prediction function, $\mathcal{X}=$ anything, $\mathcal{Y}=$ anything

1) Define auxiliary function, $g: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$,

$$
\text { e.g. } \quad g(x, y)=\prod_{i} \psi_{i}\left(y_{i}, x\right) \prod_{i \sim j} \psi_{i j}\left(y_{i}, y_{j}, x\right)
$$

2) Construct $f: \mathcal{X} \rightarrow \mathcal{Y}$ from $g$, e.g., $f(x)=\operatorname{argmax}_{y \in \mathcal{Y}} g(x, y)$

Challenges:

- how to learn $g(x, y)$ from training data?
- how to compute $f(x)$ from $g(x, y)$ ?


## Supervised Learning Problem

- Given training examples $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right) \in \mathcal{X} \times \mathcal{Y}$ $x \in \mathcal{X}$ : input, e.g. image
$y \in \mathcal{Y}$ : structured output, e.g. human pose, sentence


Images: HumanEva dataset

- How to make predictions for new inputs, i.e. learn a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ ?


## Supervised Learning Problem

- Given training examples $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right) \in \mathcal{X} \times \mathcal{Y}$ $x \in \mathcal{X}$ : input, e.g. image $y \in \mathcal{Y}$ : structured output, e.g. human pose, sentence
- How to make predictions for new inputs, i.e. learn $f: \mathcal{X} \rightarrow \mathcal{Y}$ ?


## Approach 1) Discriminative Probabilistic Learning

1) Use training data to obtain an estimate $p(y \mid x)$.
2) Use $f(x)=\operatorname{argmin}_{\bar{y} \in \mathcal{Y}} \sum_{y} p(y \mid x) \Delta(y, \bar{y})$ to make predictions. $\Delta: \mathcal{Y} \rightarrow \mathcal{Y} \rightarrow \mathbb{R}_{+}$is a structured loss function (later...)

## Approach 2) Loss-minimizing Parameter Estimation

1) Use training data to learn a compatibility function $g(x, y)$
2) Use $f(x):=\operatorname{argmax}_{y \in \mathcal{Y}} g(x, y)$ to make predictions.

## Probabilistic Graphical Models

## Refresher: Conditional Probability Distributions

## Binary Classification

$\mathcal{X}=\{$ anything $\}, \mathcal{Y}=\{ \pm 1\}$

- $p(y \mid x): 2$ values for each $x$, 1 degree of freedom
- learn one function: $\mathcal{X} \rightarrow \mathbb{R}$


## Multi-class prediction

$y \in \mathcal{Y}=\{1, \ldots, K\}$

- $p(y \mid x): K$ values for each $x$,
- learn $K-1$ functions, or $K$ functions with normalization


Structured objects: predicting $M$ variables jointly
$\mathcal{Y}=\{1, K\} \times\{1, K\} \cdots \times\{1, K\}$
For each $x$ :

- $K^{M}$ values, $K^{M}-1$ d.o.f.
$\rightarrow K^{M}$ functions



## Example: pose estimation

$$
\begin{aligned}
& \mathcal{Y}_{\text {part }}=\{1, \ldots, W\} \times\{1, \ldots, H\} \\
& \times\{1, \ldots, 360\} \\
& \mathcal{Y}=\mathcal{Y}_{\text {head }} \times \mathcal{Y}_{\text {left-arm }} \times \cdots \times \mathcal{Y}_{\text {right-foot }}
\end{aligned}
$$

For each $x$ :

- $(360 W H)^{\text {\#body parts }}$ values $\rightarrow$ many billions function



## Example: image denoising

$$
\mathcal{Y}=\{640 \times 480 \text { RGB images }\}
$$

For each $x$ :

# too much! 

$$
\begin{aligned}
& \left(255^{3}\right)^{640 \cdot 480} \text { values in } p(y \mid x), \\
& \rightarrow \quad \text { over } 10^{2,000,000} \text { functions }
\end{aligned}
$$

## Example: image denoising

$$
\mathcal{Y}=\{640 \times 480 \text { RGB images }\}
$$

For each $x$ :

## too much!

- $\left(255^{3}\right)^{640 \cdot 480}$ values in $p(y \mid x)$, $\rightarrow$ over $10^{2,000,000}$ functions

We cannot consider all possible distributions, we must impose structure.

## Probabilistic Graphical Models

A (probabilistic) graphical model defines a family of probability distributions over a set of random variables, by means of a graph.

## Probabilistic Graphical Models

A (probabilistic) graphical model defines a family of probability distributions over a set of random variables, by means of a graph.

Popular classes of graphical models,

- Undirected graphical models (Markov random fields),
- Directed graphical models (Bayesian networks),
- Factor graphs,
- Others: chain graphs, influence diagrams, etc.



## Probabilistic Graphical Models

A (probabilistic) graphical model defines

- a family of probability distributions over a set of random variables, by means of a graph.

Popular classes of graphical models,

- Undirected graphical models (Markov random fields),
- Directed graphical models (Bayesian networks),
- Factor graphs,
- Others: chain graphs, influence diagrams, etc.

The graph encodes conditional independence assumptions
 between the variables:

- Let $N(i)$ be the neighbors of node $i$ in the graph $(V, \mathcal{E})$. Then

$$
\begin{aligned}
& \qquad p\left(y_{i} \mid y_{V \backslash\{i\}}\right)=p\left(y_{i} \mid y_{N(i)}\right) \\
& \text { with } y_{V \backslash\{i\}}=\left(y_{1}, \ldots, y_{i-1}, y_{i+1}, y_{n}\right)
\end{aligned}
$$

## Example: Pictorial Structures for Articulated Pose Estimation



- All parts depend on each other.
- Knowing where the head is puts constraints on where the feet can be.
- But conditional independences as specified by the graph:
- If we fix where the left leg is, the left foot's position does not depend on the torso or the head position anymore, etc.

$$
p\left(y_{\text {left-foot }} \mid y_{\text {top }}, \ldots, y_{\text {torso }}, \ldots, y_{\text {right-foot }}, x\right)=p\left(y_{\text {left-foot }} \mid y_{\text {left-leg }}, x\right)
$$

## Factor Graphs

- Decomposable output $y=\left(y_{1}, \ldots, y_{|V|}\right)$
- Graph: $G=(V, \mathcal{F})$,
- variable nodes $V$,
- factor nodes $\mathcal{F}$,
- each factor $F \in \mathcal{F}$ connects a subset of nodes,
- write $F=\left\{v_{1}, \ldots, v_{|F|}\right\}$ and $y_{F}=\left(y_{v_{1}}, \ldots, y_{v_{|F|}}\right)$


Factor graph

## Factor Graphs

- Decomposable output $y=\left(y_{1}, \ldots, y_{|V|}\right)$
- Graph: $G=(V, \mathcal{F})$,
- variable nodes $V$,
- factor nodes $\mathcal{F}$,
- each factor $F \in \mathcal{F}$ connects a subset of nodes,
- write $F=\left\{v_{1}, \ldots, v_{|F|}\right\}$ and

$$
y_{F}=\left(y_{v_{1}}, \ldots, y_{v_{|F|}}\right)
$$



Factor graph

- Distribution factorizes into potentials $\psi$ at factors:

$$
p(y)=\frac{1}{Z} \prod_{F \in \mathcal{F}} \psi_{F}\left(y_{F}\right)
$$

- $Z$ is a normalization constant, called partition function:

$$
Z=\sum_{y \in \mathcal{Y}} \prod_{F \in \mathcal{F}} \psi_{F}\left(y_{F}\right)
$$

## Conditional Distributions

How to model $p(y \mid x)$ ?

- Potentials become also functions of (part of) $x: \psi_{F}\left(y_{F} ; x_{F}\right)$ instead of just $\psi_{F}\left(y_{F}\right)$

$$
p(y \mid x)=\frac{1}{Z(x)} \prod_{F \in \mathcal{F}} \psi_{F}\left(y_{F} ; x_{F}\right)
$$

- Partition function depends on $x_{F}$


Factor graph

$$
Z(x)=\sum_{y \in \mathcal{Y}} \prod_{F \in \mathcal{F}} \psi_{F}\left(y_{F} ; x_{F}\right)
$$

- Note: $x$ is treated just as an argument, not as a random variable.


## Conventions: Potentials and Energy Functions

Assume $\psi_{F}\left(y_{F}\right)>0$. Then

- instead of potentials, we can use energies:

$$
\begin{array}{rlr}
E_{F}\left(y_{F} ; x_{F}\right) & =-\log \left(\psi_{F}\left(y_{F} ; x_{F}\right)\right) & \text { for each factor } F . \\
E(y ; x) & =\sum_{F \in \mathcal{F}} E_{F}\left(y_{F} ; x_{F}\right) & \text { total energy }
\end{array}
$$

## Conventions: Potentials and Energy Functions

Assume $\psi_{F}\left(y_{F}\right)>0$. Then

- instead of potentials, we can use energies:

$$
\begin{array}{rlr}
E_{F}\left(y_{F} ; x_{F}\right) & =-\log \left(\psi_{F}\left(y_{F} ; x_{F}\right)\right) & \text { for each factor } F . \\
E(y ; x) & =\sum_{F \in \mathcal{F}} E_{F}\left(y_{F} ; x_{F}\right) & \text { total energy }
\end{array}
$$

- $p(y \mid x)$ can be written as Gibbs distribution

$$
\begin{aligned}
p(y \mid x) & =\frac{1}{Z(x)} \prod_{F \in \mathcal{F}} \psi_{F}\left(y_{F} ; x_{F}\right) \\
& =\frac{1}{Z(x)} \exp \left(-\sum_{F \in \mathcal{F}} E_{F}\left(y_{F} ; x_{F}\right)\right)=\frac{1}{Z(x)} \exp (-E(y ; x))
\end{aligned}
$$

In practice, one directly models the energy function
$\rightarrow$ the probability distribution is uniquely determined by it.

## Example: An Energy Function for Human Pose Estimation



$$
E(y ; x)=\sum_{i \in\{\text { head,torso }, \ldots\}} E_{i}\left(y_{i} ; x\right)+\sum_{(i, j)} E_{i j}\left(y_{i}, y_{j}\right)
$$

- unary factors (depend on one label): appearance
- e.g. $E_{\text {head }}(y ; x)$ "Does location $y$ in image $x$ look like a head?"
- pairwise factors (depend on two labels): geometry
- e.g. $E_{\text {head-torso }}\left(y_{\text {head }}, y_{\text {torso }}\right)$ "Is location $y_{\text {head }}$ above location $y_{\text {torso }}$ ?"


## Example: An Energy Function for Image Segmentation

Object segmentation: e.g. horse
$\mathcal{X}$ :


Energy function components ("Ising" model):

- $E_{i}\left(y_{i}=1, x_{i}\right)= \begin{cases}\text { low } & \text { if } x_{i} \text { is the right color, e.g. brown } \\ \text { high } & \text { otherwise }\end{cases}$
- $E_{i}\left(y_{i}=0, x_{i}\right)=-E_{i}\left(y_{i}=1, x_{i}\right)$
- $E_{i}\left(y_{i}, y_{j}\right)= \begin{cases}\text { low } & \text { if } y_{i}=y_{j} \\ \text { high } & \text { otherwise }\end{cases}$
prefer that neighbors have the same label $\rightarrow$ smooth labelings


## What to do with Structured Prediction Models?

Case 1) $p(y \mid x)$ is known

## MAP Prediction

Predict $f: \mathcal{X} \rightarrow \mathcal{Y}$ by optimization

$$
y^{*}=\underset{y \in \mathcal{Y}}{\operatorname{argmax}} p(y \mid x)=\underset{y \in \mathcal{Y}}{\operatorname{argmin}} E(y, x)
$$

## Probabilistic Inference

Compute marginal probabilities

$$
p\left(y_{F} \mid x\right)
$$

for any factor $F$, in particular, $p\left(y_{i} \mid x\right)$ for all $i \in V$.

## What to do with Structured Prediction Models?


input image

$\operatorname{argmax}_{y} p(y \mid x)$

$p\left(y_{i} \mid x\right)$ for $i=1, \ldots, 6$

- MAP makes a single (structured) prediction
- best overall pose
- Marginal probabilities $p\left(y_{i} \mid x\right)$ give us
- potential positions
- uncertainty
of the individual body parts.


## What to do with Structured Prediction Models?

Case 2) $p(y \mid x)$ is unknown, but we have training data

## Structure Learning

Learn graph structure from training data.

## Variable Learning

Learn, whether to use additional (latent) variables, and which ones. (input and output variables are fixed by the task we try to solve).

## Parameter Learning

Assume a fixed factor graph, learn parameters of the energy.

# Conditional Random Fields 

$$
\boldsymbol{\operatorname { m a x }}_{w} p(y \mid x ; w)
$$

## Conditional Random Field Learning

Goal: learn a conditional distribution

$$
p(y \mid x)=\frac{1}{Z(x)} e^{\left.-\sum_{F \in \mathcal{F}} E_{F}\left(y_{F} ; x\right)\right\rangle}
$$

with $\mathcal{F}=\{$ all factors $\}$ : all unary, pairwise, potentially higher order, ...

- parameterize each $E_{F}\left(y_{F} ; x\right)=\left\langle w_{F}, \phi_{F}\left(x, y_{F}\right)\right\rangle$.
- fixed feature functions $\left(\phi_{1}\left(x, y_{1}\right), \ldots, \phi_{|\mathcal{F}|}\left(x, y_{|\mathcal{F}|}\right)\right) \equiv: \phi(x, y)$
- weight vectors $\left(w_{1}, \ldots, w_{|\mathcal{F}|}\right) \equiv: w$

Result: log-linear model with parameter vector $w$

$$
\begin{aligned}
p(y \mid x ; w) & =\frac{1}{Z(x ; w)} e^{-\langle w, \phi(y, x)\rangle} \\
\text { with } \quad Z(x ; w) & =\sum_{\bar{y} \in \mathcal{Y}} e^{-\langle w, \phi(\bar{y}, x)\rangle} \quad(\text { "partition function") }
\end{aligned}
$$

New goal: find best parameter vector $w \in \mathbb{R}^{D}$.

## Probabilistic Learning

Maximize conditional likelihood, $p\left(\mathcal{D}_{y} \mid \mathcal{D}_{x} ; w\right)$, or maximum posterior, $p(w \mid \mathcal{D})$. Equivalently, minimize

$$
\begin{aligned}
\mathcal{L}(w) & =\frac{\lambda}{2}\|w\|^{2}-\sum_{n=1}^{N} \log p\left(y^{n} \mid x^{n} ; w\right) \\
& =\frac{\lambda}{2}\|w\|^{2}+\sum_{n=1}^{N}\left[\left\langle w, \phi\left(x^{n}, y^{n}\right)\right\rangle+\log \sum_{y \in \mathcal{Y}} e^{-\left\langle w, \phi\left(x^{n}, y\right)\right\rangle}\right]
\end{aligned}
$$

( $\lambda=0$ makes it unregularized)

Same optimization problem as for multi-class logistic regression.

- unconstrained
- smooth
- convex


## Solving the Training Optimization Problem in Practice

Task: Compute $v=\nabla_{w} \mathcal{L}\left(w_{\text {cur }}\right)$ and evaluate $\mathcal{L}\left(w_{\text {cur }}+\eta v\right)$ :

$$
\begin{aligned}
\mathcal{L}(w) & =\frac{\lambda}{2}\|w\|^{2}+\sum_{n=1}^{N}\left[\left\langle w, \phi\left(x^{n}, y^{n}\right)\right\rangle+\log \sum_{y \in \mathcal{Y}} e^{-\left\langle w, \phi\left(x^{n}, y\right)\right\rangle}\right] \\
\nabla_{w} \mathcal{L}(w) & =\lambda w+\sum_{n=1}^{N}\left[\phi\left(x^{n}, y^{n}\right)-\sum_{y \in \mathcal{Y}} p\left(y \mid x^{n} ; w\right) \phi\left(x^{n}, y\right)\right]
\end{aligned}
$$

## Solving the Training Optimization Problem in Practice

Task: Compute $v=\nabla_{w} \mathcal{L}\left(w_{\text {cur }}\right)$ and evaluate $\mathcal{L}\left(w_{\text {cur }}+\eta v\right)$ :

$$
\begin{aligned}
\mathcal{L}(w) & =\frac{\lambda}{2}\|w\|^{2}+\sum_{n=1}^{N}\left[\left\langle w, \phi\left(x^{n}, y^{n}\right)\right\rangle+\log \sum_{y \in \mathcal{Y}} e^{-\left\langle w, \phi\left(x^{n}, y\right)\right\rangle}\right] \\
\nabla_{w} \mathcal{L}(w) & =\lambda w+\sum_{n=1}^{N}\left[\phi\left(x^{n}, y^{n}\right)-\sum_{y \in \mathcal{Y}} p\left(y \mid x^{n} ; w\right) \phi\left(x^{n}, y\right)\right]
\end{aligned}
$$

Problem: $\mathcal{Y}$ typically is very (exponentially) large:

- binary image segmentation: $|\mathcal{Y}|=2^{640 \times 480} \approx 10^{92475}$
- ranking $N$ images: $|\mathcal{Y}|=N$ !, e.g. $N=1000:|\mathcal{Y}| \approx 10^{2568}$.

We must use the structure in $\mathcal{Y}$, otherwise we're lost.

## Solving the Training Optimization Problem in Practice

$$
\nabla_{w} \mathcal{L}(w)=\lambda w+\sum_{n=1}^{N}\left[\phi\left(x^{n}, y^{n}\right)-\underset{y \sim p\left(y \mid x^{n} ; w\right)}{\mathbb{E}} \phi\left(x^{n}, y\right)\right]
$$

Computing the Gradient (naive): $O\left(K^{M} N D\right)$

$$
\mathcal{L}(w)=\frac{\lambda}{2}\|w\|^{2}+\sum_{n=1}^{N}\left[\left\langle w, \phi\left(x^{n}, y^{n}\right)\right\rangle+\log Z\left(x^{n} ; w\right)\right]
$$

Line Search (naive): $O\left(K^{M} N D\right)$ per evaluation of $\mathcal{L}$

- $N$ : number of samples
- D: dimension of feature space
- $M$ : number of output variables $\approx 10$ s to $1,000,000$ s
- $K$ : number of possible labels of each output variables $\approx 2$ to 1000 s


## Solving the Training Optimization Problem in Practice

In a graphical model with factors $\mathcal{F}$, the features decompose:

$$
\begin{aligned}
\phi(x, y) & =\left(\phi_{F}\left(x, y_{F}\right)\right)_{F \in \mathcal{F}} \\
\underset{y \sim p(y \mid x ; w)}{\mathbb{E}} \phi(x, y) & =\left(\underset{y \sim p(y \mid x ; w)}{\mathbb{E}} \phi_{F}\left(x, y_{F}\right)\right)_{F \in \mathcal{F}} \\
& =\left(\underset{y_{F} \sim p\left(y_{F} \mid x ; w\right)}{\mathbb{E}} \phi_{F}\left(x, y_{F}\right)\right)_{F \in \mathcal{F}} \\
\underset{y_{F} \sim p\left(y_{F} \mid x ; w\right)}{\mathbb{E}} \phi_{F}\left(x, y_{F}\right) & =\underbrace{\sum_{y_{F} \in \mathcal{Y}_{F}}}_{K^{|F|} \text { terms }} \underbrace{p\left(y_{F} \mid x ; w\right)}_{\text {factor marginals }} \phi_{F}\left(x, y_{F}\right)
\end{aligned}
$$

Factor marginals $\mu_{F}=p\left(y_{F} \mid x ; w\right)$

- are much smaller than complete joint distribution $p(y \mid x ; w)$,
- compute/approximate them by probabilistic inference

