# Bayesian Hierarchical Classification 

Dmitry Kondrashkin (HSE) in collaboration with Dr. Christoph Lampert (IST Austria)

October 23, 2015

## Hierarchical classification

## Current approaches:

- Non-Bayesian: hierarchy of classifiers.
- Bayesian: hierarchy of priors.

Our approach:

- Use information provided by hierarchy to adjust model complexity.


## Notation

- $(X, Y)=\left\{\left(x_{n}, y_{n}\right)\right\}_{n=1}^{N}$,
- $y_{n} \in\{1, \ldots, K\}$,
- $M$ is the number of nodes in the hierarchy tree,
- $K$ is the number of classes (or leaf nodes).
- $Z=\left\{z_{l}\right\}_{l=1}^{M-K}, z_{l} \in\{0,1\}$ is binary latent variable,
- $F=\left\{f_{m}\right\}_{m=1}^{M}, f_{m}$ is Gaussian Process, i.e. $f_{m} \sim \mathcal{N}\left(0, K_{m}\right)$, $K_{m} \in \mathbb{R}^{N \times N}$,
- path $(i)$ is the set of all the nodes in the path from $i$-th node to the root,
- $\mathrm{cl}(i)$ is the set of all the leaf nodes (classes) which are in the subtree with root in the $i$-th node.


## Example: binary latent variables

- $K=5$,

- $M=9$,
- $Z=\left\{z_{1}, z_{2}, z_{3}, z_{4}\right\}$.


## Example: Gaussian Processes



- $K=5$,
- $M=9$,
- $Z=\left\{z_{1}, z_{2}, z_{3}, z_{4}\right\}$,
- $F=\left\{f_{1}, f_{2}, \ldots, f_{9}\right\}$.


## Probabilistic model

Complete likelihood:

$$
p(Y, Z, F \mid X, \Theta)=\prod_{n=1}^{N} p\left(y_{n}, z_{n} \mid F\right) \prod_{m=1}^{M} \mathcal{N}\left(f_{m} \mid 0, K\left(\theta_{m}\right)\right)
$$

- $\Theta=\{\theta\}_{m=1}^{M}$ are hyperparameters,
- $K(\cdot)$ is the kernel function.


## Probabilistic model

Likelihood for one object:

$$
\begin{aligned}
& p\left(y_{n}=k, z_{n} \mid F\right) \propto \exp \left\{f_{n k} \prod_{i \in \operatorname{path}(k)}\left(1-z_{i}\right)\right. \\
& \left.\quad+\sum_{j \in \operatorname{path}(k)}\left(f_{n j}+\rho_{j}\right) z_{j} \prod_{i \in \operatorname{path}(j)}\left(1-z_{i}\right)\right\}
\end{aligned}
$$

- $\rho_{k}=-\ln |\operatorname{cl}(k)|$ is a penalty term,
- $z_{i}=1$ means that all the classes from $\mathrm{cl}(i)$ are merged into one class.


## Example

Consider the following class hierarchy:


Probability for each class:

- $p(y=1) \propto \exp \left\{f_{1}\left(1-z_{1}\right)\left(1-z_{3}\right)+\left(f_{5}-\ln 2\right) z_{1}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$,
- $p(y=2) \propto \exp \left\{f_{2}\left(1-z_{1}\right)\left(1-z_{3}\right)+\left(f_{5}-\ln 2\right) z_{1}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$,
- $p(y=3) \propto \exp \left\{f_{3}\left(1-z_{2}\right)\left(1-z_{3}\right)+\left(f_{6}-\ln 2\right) z_{2}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$,
- $p(y=4) \propto \exp \left\{f_{4}\left(1-z_{2}\right)\left(1-z_{3}\right)+\left(f_{6}-\ln 2\right) z_{2}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$.


## Example

Consider the following class hierarchy:


If all $z_{i}=0$ :

- $p(y=1) \propto \exp \left\{f_{1}\right\}$,
- $p(y=2) \propto \exp \left\{f_{2}\right\}$,
- $p(y=3) \propto \exp \left\{f_{3}\right\}$,
- $p(y=4) \propto \exp \left\{f_{4}\right\}$.


## Example

Consider the following class hierarchy:


Probability for each class:

- $p(y=1) \propto \exp \left\{f_{1}\left(1-z_{1}\right)\left(1-z_{3}\right)+\left(f_{5}-\ln 2\right) z_{1}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$,
- $p(y=2) \propto \exp \left\{f_{2}\left(1-z_{1}\right)\left(1-z_{3}\right)+\left(f_{5}-\ln 2\right) z_{1}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$,
- $p(y=3) \propto \exp \left\{f_{3}\left(1-z_{2}\right)\left(1-z_{3}\right)+\left(f_{6}-\ln 2\right) z_{2}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$,
- $p(y=4) \propto \exp \left\{f_{4}\left(1-z_{2}\right)\left(1-z_{3}\right)+\left(f_{6}-\ln 2\right) z_{2}\left(1-z_{3}\right)+\left(f_{7}-\ln 4\right) z_{3}\right\}$.


## Example

Consider the following class hierarchy:


If all $z_{i}=1$ :

- $p(y=1) \propto \exp \left\{f_{7}-\ln 4\right\}=0.25$,
- $p(y=2) \propto \exp \left\{f_{7}-\ln 4\right\}=0.25$,
- $p(y=3) \propto \exp \left\{f_{7}-\ln 4\right\}=0.25$,
- $p(y=4) \propto \exp \left\{f_{7}-\ln 4\right\}=0.25$.


## Example

Consider a configuration: $z_{1}=1, z_{2}=0, z_{3}=0$ :


- $p(y=1) \propto \exp \left\{f_{5}+\ln 0.5\right\}=0.5 \exp \left\{f_{5}\right\}$,
- $p(y=2) \propto \exp \left\{f_{5}+\ln 0.5\right\}=0.5 \exp \left\{f_{5}\right\}$,
- $p(y=3) \propto \exp \left\{f_{3}\right\}$,
- $p(y=4) \propto \exp \left\{f_{4}\right\}$.


## Inference: Variational Bayes

We use variational Bayes to approximate the posterior over $Z$ and $F$ :

$$
\ln p(Y \mid X, \Theta) \geq \sum_{Z} \int q(Z) q(F) \ln \frac{p(Y, Z \mid F) p(F \mid \Theta)}{q(Z) q(F)} d F
$$

- Here we approximate the posterior $p(F, Z) \approx q(F) q(Z)$,
- $p(F \mid \Theta)$ is the prior,
$p(Y, Z \mid F)$ is a softmax, thus the integral is intractable, there are two options to overcome this issue:
- Use local variational bounds like Jaakkola-Jordan bound to obtain a quadratic lower bound for the softmax to make this integral tracktable.
- Use stochastic optimization and reparametrization trick to compute gradient of the lower bound.


## Inference: Variational Bayes

We would like to find $q(F)=\prod_{m=1}^{M} \mathcal{N}\left(f_{m} \mid \mu_{m}, \Sigma_{m}\right)$ and $q(Z)$ in the family of discrete distributions.
Variational Bayes update rules:

- $q(Z) \propto \exp \left\{\mathbb{E}_{q(F)} \ln p(Y, Z \mid F) p(F \mid \Theta)\right\}$
- $q(F) \propto \exp \left\{\mathbb{E}_{q(Z)} \ln p(Y, Z \mid F) p(F \mid \Theta)\right\}$

The issues:

- The first integration is intractable because of softmax in likelihood.
- We could compute $\mathbb{E}_{q(Z)}$ by explicit summation only if $|Z|$ is small.


## Solving the first issue: local variational bound

1. bound log-sum-exp:

$$
\ln \sum_{k=1}^{K} e^{x_{k}} \leq \alpha+\sum_{k=1}^{K} \ln \left(1+e^{x-\alpha}\right), \forall \alpha \in \mathbb{R}
$$

2. bound log-sigmoid using Jaakkola-Jordan bound:

$$
\ln \left(1+e^{x}\right) \leq \lambda(\xi)\left(x^{2}-\xi^{2}\right)+(x-\xi) / 2+\ln \left(1+e^{\xi}\right), \forall \xi \in \mathbb{R},
$$

where $\lambda(\xi)=\frac{1}{2 \xi}\left(1 /\left(1+e^{-\xi}\right)-0.5\right)$.

## Solving the second issue

Update rule for $q(F)$ after applying the bound:

$$
\begin{gathered}
\ln q(F)=\text { const }+ \text { LogPrior }+ \\
\sum_{n=1}^{N} \sum_{k=1}^{K}\left(\left[y_{n}=k\right]\left(f_{n k} \mathbb{E}_{Z} \tilde{z}_{k}^{0}+\sum_{j \in \operatorname{path}(k)} \tilde{f}_{n j} \mathbb{E}_{Z} \tilde{z}_{k}\right)\right. \\
-0.5\left(f_{n k} \mathbb{E}_{Z} \tilde{z}_{k}^{0}+\sum_{j \in \operatorname{path}(k)} \tilde{f}_{n j} \mathbb{E}_{Z} \tilde{z}_{k}\right) \\
\left.-\lambda\left(\xi_{n k}\right)\left(f_{n k}^{2} \mathbb{E}_{Z} \tilde{z}_{k}^{0}+\sum_{j \in \operatorname{path}(k)} \tilde{f}_{n j}^{2} \mathbb{E}_{Z} \tilde{z}_{k}+\ldots\right)\right)
\end{gathered}
$$

where

- $\tilde{f}_{n j}=f_{n j}+\rho_{j}$
- $\tilde{z}_{k}^{0}=\prod_{j \in \operatorname{path}(k)}\left(1-z_{j}\right)$
- $\tilde{z}_{k}=z_{k} \prod_{j \in \operatorname{path}(k)}\left(1-z_{j}\right)$


## Solving the second issue: message passing

Note that

$$
\begin{gathered}
\mathbb{E}_{Z} \tilde{z}_{k}^{0}=\mathbb{E}_{Z} \prod_{j \in \operatorname{path}(k)}\left(1-z_{j}\right)=q\left(z_{j_{1}}=0, \ldots, z_{j_{s}}=0\right) \\
\mathbb{E}_{Z} \tilde{z}_{k}=\mathbb{E}_{Z} z_{k} \prod_{j \in \operatorname{path}(k)}\left(1-z_{j}\right)=q\left(z_{k}=1, z_{j_{1}}=0, \ldots, z_{j_{s}}=0\right)
\end{gathered}
$$

where $\left\{j_{1}, \ldots, j_{s}\right\}=\operatorname{path}(k)$. These are the marginals of $q(Z)$, they could be efficiently computed using message passing.

## Inference: local bounds

## Pros:

- We could derive closed form updates for $\xi_{n k}, \alpha_{n}, q(F)$ and $q(Z)$.
- We could use effective dynamic programming approach to deal with discrete distribution $q(Z)$.
Cons:
- Comlicated implementation.
- We introduce $N \times K+N$ additional variational parameters $\left\{\xi_{n k}\right\}_{n=1, k=1}^{N, K}$ and $\left\{\alpha_{n}\right\}_{n=1}^{N}$.
- The bound might be untight.


## Inference: stochastic optimization

Pros:

- Very simple implementation using Theano.
- Using enough samples we could obtain very good approximation.

Cons:

- Reparametrization trick doesn't work for discrete distributions: only explicit summation available.
- Convergence issues.
- Slow.


## Experiments: 10 ImageNet classes





## Experiments: 10 ImageNet classes, 5 objects per class

- 0 kit fox

- 1 English setter
- 2 Siberian husky
- 3 Australian terrier
- 4 English springer
- 5 grey whale
(2) 6 red panda
- 7 Egyptian cat
- 8 ibex
- 9 Persian cat


## Experiments: 20 ImageNet classes



## Conclusion

In this project we have used the following ideas:

- Marginal likelihood optimization to adjust model complexity.
- Local variational bounds to perform inference.
- Message passing to handle discrete distribution $q(Z)$.

