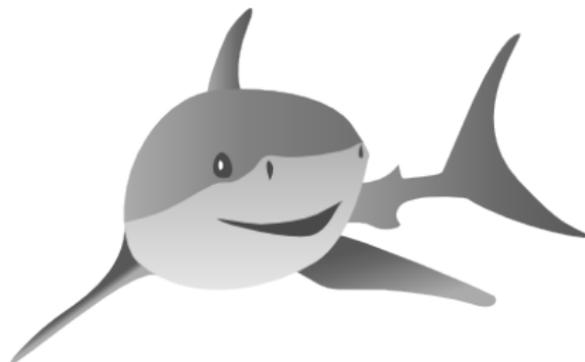


Restricted Boltzmann Machines in Shark

Practical Session

Oswin Krause

Department of Computer Science



Shark 3.0

- General machine learning library
- Written in C++
 - 80k source LOC
 - 10 Years of development
 - LGPL
- Targeted towards research
- Trade-offs
 - ⊕ performance
 - ⊕ genericity
 - ⊖ complexity
 - ⊖ knowledge about C++
- C. Igel, V. Heidrich-Meisner, and T. Glasmachers. Shark. JMLR 9, pp. 993-996, 2008.



What does it offer?

- Does not offer ready made solutions
- Shark is a toolbox. . .
 - Classification
 - Regression
 - Optimization
- . . . and a toolbox to make tools
 - Algorithms are modular
 - Library includes many bits and pices. . .
 - . . . and code to connect them
 - Easy to add new pieces



Examples of Implemented Algorithms

- SVMs: linear, nonlinear, multi-class
- Neural Networks: classification, regression, autoencoder
- Classification and Regression Trees, Random Forests
- Nearest-Neighbor algorithms: linear and kernel
- Restricted Boltzman Machines
- Clustering: (kernel) k-means
- Multi-Objective Optimization (MO-CMA-ES)
- CMA-ES
- General optimisation Methods: CG, BFGS, ...
- ...



Restricted Boltzmann Machines

Restricted Boltzmann Machines (RBMS)

- $p(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} e^{-\frac{1}{T} E(\mathbf{v}, \mathbf{h})}$ with $Z = \sum_{\mathbf{v}, \mathbf{h}} e^{-\frac{1}{T} E(\mathbf{v}, \mathbf{h})}$
- Generalized formula for $E(\mathbf{v}, \mathbf{h})$

$$E(\mathbf{v}, \mathbf{h}) = f_h(\mathbf{h}) + f_v(\mathbf{v}) + \phi_h(\mathbf{h})^T W \phi_v(\mathbf{v})$$

- Allows for a wide set of RBMs used in practice (binary, exponential-binary, Gaussian-binary, semi-RBMs, . . .)
- f_h and ϕ_h depend on the choice of hidden distributions (same for visible)
- $f_h(\mathbf{h}) = \mathbf{b}^T \mathbf{h}$, $f_v(\mathbf{v}) = \mathbf{c}^T \mathbf{v}$ and $\phi_h(\mathbf{h}) = \mathbf{h}$, $\phi_v(\mathbf{v}) = \mathbf{v}$ leads to

$$E(\mathbf{v}, \mathbf{h}) = \mathbf{b}^T \mathbf{h} + \mathbf{c}^T \mathbf{v} + \mathbf{h}^T W \mathbf{v}$$

The binary-binary RBM for $\mathbf{v} \in \{0, 1\}^n$ and $\mathbf{h} \in \{0, 1\}^m$.



Restricted Boltzmann Machines in Shark

- Shark implements the generalized energy formula.
- Supports binary, Gaussian, bipolar and truncated exponential distributions.
- Variety of gradient approximators:
 - Analytic (exponential time!)
 - Contrastive Divergence
 - Persistent Contrastive Divergence
 - Parallel Tempering
- Estimators for Z
 - Analytic (exponential time!)
 - AIS
 - Bridge-Sampling



Take a look at it

Code-Time!



Some Assignments

Play around with the code:

- a) BinaryRBM.cpp is about training an RBM
 - Try to get a high likelihood/ prevent divergence
 - Change hyper parameters(k , learning rate, momentum, regularization)
 - What happens with low/high values?
 - Try Parallel Tempering instead of CD

```
size_t chains = 10;  
BinaryParallelTempering estimator(&rbm);  
estimator.chain().setUniformTemperatureSpacing(chains)
```

- Momentum is easier to use when multiplying the learning rate by $(1 - v)$, where $v \in [0, 1]$ is the strength of the momentum



Some Assignments

- b) `BinaryRBMFeatureClassification.cpp` is a semi-supervised task
 - First an RBM is trained and then an LDA is trained on the learned features
 - Goal: Reduce classification error
 - Does RBM model quality affect performance?
 - Do not forget to regularize LDA
- c) `BinaryRBMClassification.cpp` is also semi-supervised
 - First an RBM is trained and then transformed into an FFNet which is then fine tuned
 - Goal: Reduce classification error
 - Does RBM model quality affect performance?
 - (Try the other tasks first!)

