
Distributed Optimization of Binary Autoencoders Using Auxiliary Coordinates

Miguel Á. Carreira-Perpiñán Mehdi Alizadeh
EECS, University of California, Merced
<http://eecs.ucmerced.edu>

Introduction Many powerful machine learning systems are based on the composition of multiple processing layers, such as deep nets, which gives rise to nonconvex objective functions. A general, recent approach to optimize such “nested” functions is the *method of auxiliary coordinates (MAC)* [3]. This decomposes the optimization into steps that alternate between training single layers and updating the coordinates. It has the advantage that it reuses existing single-layer algorithms, introduces parallelism, and does not need to use chain-rule gradients, so it works with nondifferentiable layers. With large-scale problems, or when distributing the computation is necessary for faster training, the dataset may not fit in a single machine. It is then essential to limit the amount of communication between machines so it does not obliterate the benefit of parallelism. We describe a general way to achieve this, *ParMAC*, and illustrate it in learning binary autoencoders for fast information retrieval.

MAC for binary autoencoders Although MAC (and ParMAC) apply more generally to multiple layers, we focus on the case of the composition of two functions, in particular the *binary autoencoder (BA)* [1]. This is a usual autoencoder, with an encoder \mathbf{h} and a decoder \mathbf{f} , but with a binary code layer, trained by least squares: $E_{\text{BA}}(\mathbf{h}, \mathbf{f}) = \sum_{n=1}^N \|\mathbf{x}_n - \mathbf{f}(\mathbf{h}(\mathbf{x}_n))\|^2$, given a training set of D -dimensional patterns $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$, where $\mathbf{h}(\mathbf{x}) = \sigma(\mathbf{W}\mathbf{x})$ applies a step function $\sigma(\cdot)$ to a linear (or nonlinear) function. Note this problem is nonconvex and chain-rule gradients are inapplicable because of the binary outputs of \mathbf{h} . MAC introduces an *auxiliary coordinate* $\mathbf{z}_n = \mathbf{h}(\mathbf{x}_n) \in \{0, 1\}^L$ corresponding to the L -bit code vector (in the bottleneck layer), and uses a penalty method to optimize the constrained problem. This results in optimizing

$$E_Q(\mathbf{h}, \mathbf{f}, \mathbf{Z}; \mu) = \sum_{n=1}^N \left(\|\mathbf{x}_n - \mathbf{f}(\mathbf{z}_n)\|^2 + \mu \|\mathbf{z}_n - \mathbf{h}(\mathbf{x}_n)\|^2 \right) \text{ s.t. } \mathbf{z}_n \in \{0, 1\}^L, n = 1, \dots, N$$

and driving the penalty parameter $\mu \rightarrow \infty$. We alternate two steps over \mathbf{Z} and (\mathbf{h}, \mathbf{f}) :

M step (“model” or “minimization”) Over (\mathbf{h}, \mathbf{f}) for fixed \mathbf{Z} , we obtain $L + D$ independent problems for each of the L single-bit hash functions and each of the D decoders. With linear \mathbf{h} and \mathbf{f} this simply involves fitting L SVMs to (\mathbf{X}, \mathbf{Z}) and D linear regressors to (\mathbf{Z}, \mathbf{X}) .

C step (“coordination”) Over \mathbf{Z} for fixed (\mathbf{h}, \mathbf{f}) , the problem separates for each of the N codes $\mathbf{z}_1, \dots, \mathbf{z}_N$. This can be solved by enumeration if L is small or by alternating optimization.

Training BAs with MAC beats approximate approaches such as relaxing the codes or the step function in the encoder. It yields state-of-the-art binary hash functions \mathbf{h} [1], which can be used for fast approximate nearest-neighbor search using Hamming distances (e.g. $N = 10^9$ points with $D = 500$ take 2 TB, but only 8 GB using $L = 64$ bits, which easily fits in RAM).

Computationally, MAC achieves embarrassing parallelism within each step: $L + D$ independent models in the M step and N independent codes in the C step. We now show how to turn this into a distributed, low-communication ParMAC algorithm.

ParMAC for binary autoencoders With large datasets in distributed systems, it is imperative to minimize data movement over the network because the communication time generally far exceeds the computation time in modern architectures. In MAC we have 3 types of data: the original training

data \mathbf{X} , the auxiliary coordinates \mathbf{Z} , and the model parameters. Usually, the latter type is far smaller. In ParMAC, we never communicate training or coordinate data; each machine keeps a disjoint portion of (\mathbf{X}, \mathbf{Z}) corresponding to a subset of the points. Only model parameters are communicated (for BAs, the hash functions \mathbf{h} and the decoder \mathbf{f}). Using P machines, ParMAC iterates as follows:

M step The hash functions w_l and decoders f_d visit each machine. This implies we train them with stochastic gradient descent, where one “epoch” for model w_l corresponds to w_l having visited all P machines (within each machine, data are split into minibatches as well).

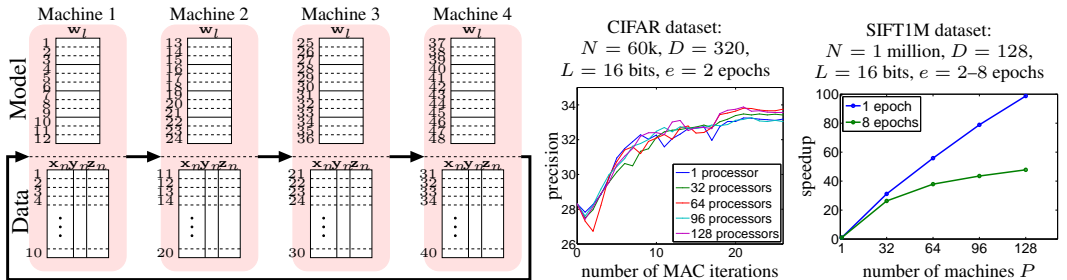
C step Identical to MAC, each training point’s coordinates z_n are optimized independently, in parallel over machines (since each machine contains x_n, z_n , and all the model parameters).

No communication at all occurs within the C step. In the M step, all w_l are communicated asynchronously in parallel. With e epochs, the entire model parameters are communicated $e + 1$ times. The last round of communication is needed to ensure each machine has the most updated version of the model for the C step. This communication can be implemented with a circular topology. This introduces no bottlenecks, unlike the use of parameter servers that gather parameters from the workers, update them, and broadcast them back to the workers.

The only approximation to the original MAC algorithm is incurred by using SGD in the M step. Since we can guarantee convergence of SGD under certain conditions, we can recover the original convergence guarantees for MAC [2]. ParMAC has additional advantages: 1) *Data shuffling* in each epoch is easy by randomizing the sequence in which machines are visited for each model (the circular topology). 2) *Load balancing* can be controlled in advance, by loading more data in faster machines, or on the fly, by having faster machines run more iterations. 3) *Fault tolerance* is easy by replacing a faulty machine with a new one where the corresponding \mathbf{X} are loaded. 4) *Streaming* can be done by adding, removing or replacing machines and their data on the fly.

In general with multilayer models, e.g. deep nets, ParMAC locks a portion of the input, output and coordinate data $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ inside each machine, and each hidden unit weight vector is an independent model that is communicated in the M step. Each weight vector may only depend on a subset of the dimensions of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, and each dimension of z_n may only depend on a subset of the weights.

Experiments We have trained BAs with ParMAC (implemented in C with MPI) in both distributed machines (UCSD Triton Shared Computer Cluster, with 128 nodes having 4 GB RAM per core) and shared-memory machines (with 32 cores and a total of 256 GB of RAM), achieving near-linear speedups even with small datasets. We are able to train BAs with linear and kernel hash functions on the SIFT1B dataset (<http://corpus-texmex.irisa.fr>), containing $N = 1$ billion images represented as $D = 128$ SIFT features, using $L = 64$ bits and $P = 128$ machines, in around a day. This would take months in a single machine with enough RAM to hold the data and parameters (which would require several TB).



Acknowledgments Supported by NSF award IIS-1423515 and a Google Faculty Research Award.

References

[1] M. Á. Carreira-Perpiñán and R. Raziperchikolaei. Hashing with binary autoencoders. In *CVPR*, 2015.
 [2] M. Á. Carreira-Perpiñán and W. Wang. Distributed optimization of deeply nested systems. arXiv:1212.5921 [cs.LG], Dec. 24 2012.
 [3] M. Á. Carreira-Perpiñán and W. Wang. Distributed optimization of deeply nested systems. *AISTATS*, 2014.