

Towards a neurosymbolic approach to the MAX-SAT problem using LBM

Kristian Aguilar, Gustavo Oliveira, Jônatas Lima, Fernando Ventilari, Edjard Mota, Son Tran, Artur Garcez

{kristian.aguilar,gustavo.oliveira,jonatas.azevedo,fernando.neder,edjard}@icomp.ufam.edu.br, son.tran@deakin.edu.au,a.garcez@city.ac.uk



Introduction

In this poster we utilize the Logical Boltzmann Machine (LBM) introduced by Son [5] to tackle the MAX-SAT problem. A LBM is a neurosymbolic system that can represent any logical proposition in strict disjunctive normal form (SDNF) as a Restricted Boltzmann Machine (RBM). As an energy-based model, we can minimize its free energy function to arrive at the truth value assignment that maximises the number of satisfied clauses. We benchmark our model with the incomplete solver Loandra, one of the highest-rated solutions in the 2023 Max-SAT Evaluation. Although results show average cost of the LBM is lower than Loandra, the latter still performs better at current MAX-SAT benchmarks with more than 20.000 clauses.

Related Works

Son introduced a neurosymbolic system (Logical Boltzmann Machine) for reasoning about symbolic knowledge at scale, showing equivalence between energy minimization and logical satisfiability [5]. Furthermore, Son showed equivalence between propositional logic and restricted Boltzmann machines and future work focuses on scaling up applications to SAT and end-to-end learning and reasoning [6]. Hernandez et al. [2] introduces High Order Boltzmann Machine (HOBM) applications and maps the MAX-SAT problem into an HOBM combinatorial optimization framework, offering an approximate solution to SAT.

Conclusion

Although the average cost of the LBM is lower than Loandra, we cannot conclude that it is better at solving MAX-SAT due to the high standard deviation on both values. It is necessary to develop our current implementation further to be more competitive with current symbolic solutions. A possible alternative would

References

- [1] Josep Argelich, Chu Min Li, Felip Manyà, and Jordi Planes. Max-sat 2016 - eleventh max-sat evaluation, 2016.
- [2] C Hernández, F.X Albizuri, A D’Anjou, M Graña, and F.J. Torrealdea. Efficient solution of max-sat and sat via higher order boltzmann machines. *Investigación Operacional*, 22(1), Jun. 2023.
- [3] Geoffrey E Hinton. A practical guide to training restricted boltzmann machines. In *Neural Networks: Tricks of the Trade: Second Edition*, pages 599–619. Springer, 2012.
- [4] Matti Järvisalo, Jeremias Berg, Ruben Martins, and Andreas Niskanen. Maxsat evaluation 2023, 2023.
- [5] Son N. Tran and Artur d’Avila Garcez. Logical boltzmann machines, 2021.
- [6] Son N. Tran and Artur d’Avila Garcez. Neurosymbolic reasoning and learning with restricted boltzmann machines. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(5):6558–6565, Jun. 2023.
- [7] David Warde-Farley, Vinod Nair, Yujia Li, Ivan Lobov, Felix Gimeno, and Simon Osindero. Solving maxsat with matrix multiplication, 2023.

Acknowledgements

We thank UFAM for partially supporting this work.

MAX-SAT problem definition

MAX-SAT is a generalization of the SAT problem (Boolean satisfiability problem), where the problem is to determine the truth value assignment, also called model, that maximises the number of satisfied clauses of a given boolean formula, typically in conjunctive normal form (CNF).

Example: $(x_1 \vee x_2) \wedge (\neg x_1 \vee x_2) \wedge (x_1 \vee \neg x_2) \wedge (\neg x_1 \vee \neg x_2)$

This formula is unsatisfiable, but the model $x_1 = F, x_2 = F$ results in three clauses being satisfied, which is the maximum number of satisfiable clauses.

Algorithm Explanation

A Restricted Boltzmann Machine (RBM), shown in Figure 1, is a bidirectional graph model that captures the probability distribution of a given dataset.




Figure 1: A graphical representation of an RBM. Source: Nguyễn Văn Lĩnh

As an energy-based model, its purpose is to encode dependencies between variables using the following energy function [5]:

$$E(x, \mathbf{h}) = - \sum_{i,j} w_{ij} x_i h_j - \sum_i a_i x_i - \sum_j b_j h_j \quad (1)$$

Where a_i and b_j are the biases of the input node x_i and hidden unit h_j respectively, and w_{ij} is the connection weight between x_i and h_j . Son [5] proved that any logical proposition can be represented by an RBM’s energy function by first converting the proposition into the Strict Disjunctive Normal Form, which is a DNF with at most one conjunctive clause that maps to True.

One way of finding the MAX-SAT model is utilizing the free energy function of the RBM:

$$\mathcal{F}_B = - \sum_j (-\log(1 + \exp(c \sum_{i \in \mathcal{A} \cup \mathcal{B}} w_{ij} x_i + b_j))) \quad (2)$$

Where the term $-\log(1 + \exp(c \sum_{i \in \mathcal{A} \cup \mathcal{B}} w_{ij} x_i + b_j))$ is a negative softplus function scaled by a non-negative value c called confidence value. The more likely a truth assignment is to satisfy the proposition, the lower the output. Using a local search method named dual annealing, the algorithm can find value assignments that minimize the free energy of the RBM. The general solution is summarized below:

LBM(cnf,c,t):

cnf: A logical proposition in conjunctive normal form

c: Confidence value

t: timeout for the optimizing algorithm, given in seconds

sdnf \leftarrow toSDNF(*cnf*)

rbm \leftarrow RBM(*sdnf*, *c*)

cost, *model* \leftarrow FEMin(*rbm*, *t*)

return *cost*, *model*

▷ Converts CNF to SDNF

▷ Builds the RBM model

▷ Minimizes the RBM’s free energy function

Results

	Cost (avg.)
Loandra	171.4±8.45
LBM	169.4±4.22

Table 1: Performance of LBM against Loandra.

To benchmark the effectiveness of LBM against modern MAX-SAT solutions we compared our algorithm to Loandra, one of the top solvers in MAX-SAT Evaluations 2023 [4]. The chosen dataset comes from the MAX-SAT Evaluation 2016 incomplete category, named *maxcut*. We chose an older set of data due to the processing limitations in our algorithm, that couldn’t handle propositions with more than 20.000 clauses. Each algorithm passed through the same 5 different instances of the problem, with the final cost averaged out. The cost is the number of clauses that were not satisfied by the truth value assignment given by the solvers.