

Machine learning for optimization (3/5)

Clément W. Royer

M2 MODO - 2025/2026

December 8, 2026



An example (not for a lab session)

- Continuous relaxations of SAT problems.
- Differentiable solvers (Wang et al '19).

A more toy example (easier for a lab session)

- Learning solver hyperparameters.
- Unfolded ISTA (Ablin et al '19).

- 1 SATNet
- 2 Learning with unfolded solvers

- 1 SATNet
- 2 Learning with unfolded solvers

Classical example: MaxCut (Goemans, Williamson '95).

- Problem: Given graph (V, E) with weighted edges, find a cut with maximum edge weight.

Classical example: MaxCut (Goemans, Williamson '95).

- Problem: Given graph (V, E) with weighted edges, find a cut with maximum edge weight.
- Using graph Laplacian matrix $\mathbf{L} \in \mathbb{R}^{n \times n}$, can be formulated as

$$\underset{\mathbf{x} \in \{-1, 1\}^n}{\text{maximize}} \mathbf{x}^T \mathbf{L} \mathbf{x}.$$

Classical example: MaxCut (Goemans, Williamson '95).

- Problem: Given graph (V, E) with weighted edges, find a cut with maximum edge weight.
- Using graph Laplacian matrix $L \in \mathbb{R}^{n \times n}$, can be formulated as

$$\underset{x \in \{-1, 1\}^n}{\text{maximize}} \quad x^T L x.$$

- Equivalent to the continuous program

$$\underset{X \in \mathcal{S}^{n \times n}}{\text{maximize}} \quad \text{trace}(L^T X) \quad \text{subject to} \quad X_{ii} = 1, X \succeq 0, \text{rank}(X) = 1.$$

Classical example: MaxCut (Goemans, Williamson '95).

- Problem: Given graph (V, E) with weighted edges, find a cut with maximum edge weight.
- Using graph Laplacian matrix $L \in \mathbb{R}^{n \times n}$, can be formulated as

$$\underset{x \in \{-1, 1\}^n}{\text{maximize}} \quad x^T L x.$$

- Equivalent to the continuous program

$$\underset{X \in \mathcal{S}^{n \times n}}{\text{maximize}} \quad \text{trace}(L^T X) \quad \text{subject to} \quad X_{ii} = 1, X \succeq 0, \text{rank}(X) = 1.$$

- Remove rank constraint: Get the SDP relaxation!

$$\underset{X \in \mathcal{S}^{n \times n}}{\text{maximize}} \quad \text{trace}(L^T X) \quad \text{subject to} \quad X_{ii} = 1, X \succeq 0.$$

From the relaxation to a solution

- 1 Solve Max-Cut SDP \Rightarrow Solution $\mathbf{X}^* \succeq \mathbf{0}$, $\mathbf{X}_{ii}^* = 1$.

From the relaxation to a solution

- 1 Solve Max-Cut SDP \Rightarrow Solution $\mathbf{X}^* \succeq \mathbf{0}$, $\mathbf{X}_{ii}^* = 1$.
- 2 Write $\mathbf{X}^* = [\mathbf{v}_i^T \mathbf{v}_j]$ with $\mathbf{v}_1, \dots, \mathbf{v}_n$ unit vectors in \mathbb{R}^n .

From the relaxation to a solution

- 1 Solve Max-Cut SDP \Rightarrow Solution $\mathbf{X}^* \succeq \mathbf{0}$, $\mathbf{X}_{ii}^* = 1$.
- 2 Write $\mathbf{X}^* = [\mathbf{v}_i^T \mathbf{v}_j]$ with $\mathbf{v}_1, \dots, \mathbf{v}_n$ unit vectors in \mathbb{R}^n .
- 3 Draw \mathbf{u} uniformly at random in the unit sphere, and set

$$\forall i = 1, \dots, n, \quad x_i^* = \begin{cases} -1 & \text{if } \mathbf{u}^T \mathbf{v}_i \leq 0 \\ 1 & \text{if } \mathbf{u}^T \mathbf{v}_i > 0. \end{cases}$$

From the relaxation to a solution

- 1 Solve Max-Cut SDP \Rightarrow Solution $\mathbf{X}^* \succeq \mathbf{0}$, $\mathbf{X}_{ii}^* = 1$.
- 2 Write $\mathbf{X}^* = [\mathbf{v}_i^T \mathbf{v}_j]$ with $\mathbf{v}_1, \dots, \mathbf{v}_n$ unit vectors in \mathbb{R}^n .
- 3 Draw \mathbf{u} uniformly at random in the unit sphere, and set

$$\forall i = 1, \dots, n, \quad x_i^* = \begin{cases} -1 & \text{if } \mathbf{u}^T \mathbf{v}_i \leq 0 \\ 1 & \text{if } \mathbf{u}^T \mathbf{v}_i > 0. \end{cases}$$

Guarantees

- Randomized rounding above finds an 0.87856-approximation!
- Similar guarantees can be obtained for other problems, such as MAXSAT.
- Challenge: SDPs are difficult to solve at scale.

MAXSAT problem

Given m vectors $\{\tilde{\mathbf{s}}_i\} \subset \{-1, 0, 1\}^m$, solve

$$\underset{\tilde{\mathbf{v}} \in \{-1, 1\}^n}{\text{maximize}} \sum_{j=1}^m \bigvee_{i=1}^n \mathbf{1} \{ \tilde{s}_{ij} \tilde{v}_i > 0 \}$$

MAXSAT problem


Given m vectors $\{\tilde{\mathbf{s}}_i\} \subset \{-1, 0, 1\}^m$, solve

$$\underset{\tilde{\mathbf{v}} \in \{-1, 1\}^n}{\text{maximize}} \sum_{j=1}^m \bigvee_{i=1}^n \mathbf{1} \{ \tilde{s}_{ij} \tilde{v}_i > 0 \}$$

Continuous SDP relaxation

$$\underset{\mathbf{V} \in \mathbb{R}^{k \times (n+1)}}{\text{minimize}} \langle \mathbf{S}^T \mathbf{S}, \mathbf{V}^T \mathbf{V} \rangle \quad \text{s.t.} \quad \|\mathbf{v}_i\| = 1 \forall i = 1, \dots, n+1.$$

- Relax \tilde{v}_i into $\mathbf{v}_i \in \mathbb{R}^k$, $\|\mathbf{v}_i\| = 1$.
- Add a variable \mathbf{v}_0 to apply randomized rounding.
- \mathbf{S} built from the $\tilde{\mathbf{s}}_i$ with scaling.
- If $k > \sqrt{2n}$, recovers the original solution.



Optimization solver

- Use vector representation of SDP matrix.
- Cheap update, one vector at a time.
- Amenable to batch parallelism.
- Can differentiate through the solver!

SDP layer

- Careful encoding of backpropagation.
- Continuous relaxation and randomized rounding encoded through probability distributions.

Cool SATNet example

SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver

Model	Train	Test	Model	Train	Test	Model	Train	Test
ConvNet	72.6%	0.04%	ConvNet	0%	0%	ConvNet	0.31%	0%
ConvNetMask	91.4%	15.1%	ConvNetMask	0.01%	0%	ConvNetMask	89%	0.1%
SATNet (ours)	99.8%	98.3%	SATNet (ours)	99.7%	98.3%	SATNet (ours)	93.6%	63.2%

(a) Original Sudoku. (b) Permuted Sudoku. (c) Visual Sudoku. (Note: the theoretical “best” test accuracy for our architecture is 74.7%.)

Table 1. Results for 9×9 Sudoku experiments with 9K train/1K test examples. We compare our SATNet model against a vanilla convolutional neural network (ConvNet) as well as one that receives a binary mask indicating which bits need to be learned (ConvNetMask).

- Setup: Learn rules and fill out Sudoku grids, represented as vectors.
- Convolutional networks treat grids as images, must learn the masked bits.
- Permuting the inputs does not change the rules to learn \Rightarrow Clear advantage of SATNet.

- 1 SATNet
- 2 Learning with unfolded solvers

Goal: Sparse representation/coding

$$\underset{z \in \mathbb{R}^m}{\text{minimize}} f(z) = \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1.$$

- $x \in \mathbb{R}^n$: Target vector.
- $D \in \mathbb{R}^{n \times m}$: Dictionary.
- $\|z\|_1 = \sum_{j=1}^d |z_j|$: Promotes sparsity of z .

An algorithm: ISTA

Problem minimize $_{\mathbf{z} \in \mathbb{R}^m}$ $f(\mathbf{z}) = \frac{1}{2} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|^2 + \lambda \|\mathbf{z}\|_1$.

ISTA iteration

$$\mathbf{z}_{k+1} = \text{ST}(\mathbf{z}_k - \alpha_k \mathbf{D}^T(\mathbf{D}\mathbf{z}_k - \mathbf{x}); \lambda\alpha_k)$$

- Gradient descent step+ Soft-thresholding (ST) (\equiv Proximal gradient)

$$\text{ST}(t; \mu) = \begin{cases} t - \mu & \text{if } t > \mu \\ t + \mu & \text{if } t < \mu \\ 0 & \text{otherwise.} \end{cases}$$

An algorithm: ISTA

Problem minimize $_{\mathbf{z} \in \mathbb{R}^m}$ $f(\mathbf{z}) = \frac{1}{2} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|^2 + \lambda \|\mathbf{z}\|_1$.

ISTA iteration

$$\mathbf{z}_{k+1} = \text{ST}(\mathbf{z}_k - \alpha_k \mathbf{D}^T(\mathbf{D}\mathbf{z}_k - \mathbf{x}); \lambda \alpha_k)$$

- Gradient descent step+ Soft-thresholding (ST) (\equiv Proximal gradient)

$$\text{ST}(t; \mu) = \begin{cases} t - \mu & \text{if } t > \mu \\ t + \mu & \text{if } t < -\mu \\ 0 & \text{otherwise.} \end{cases}$$

- Key parameter: Stepsize α_k .
- Popular choice: $\alpha_k = \frac{1}{L}$ for every k , where $L = \|\mathbf{D}\|^2$ (Lipschitz constant for the gradient of $\frac{1}{2} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|^2$).

Neural net representation of ISTA

- One “layer” = K iterations of ISTA starting from z_0 .
- Network parameters: $\alpha_0, \dots, \alpha_{K-1}$.
- Fixed parameters: D, λ .

Neural net representation of ISTA

- One “layer” = K iterations of ISTA starting from z_0 .
- Network parameters: $\alpha_0, \dots, \alpha_{K-1}$.
- Fixed parameters: D, λ .

Training

- Forward: Feed x , get $f(z_K(x))$ from **explicit ISTA iterations**.
- Backward: Automatic differentiation (nothing to do!).
- Training data: Unsupervised, try to minimize

$$\frac{1}{N} \sum_{i=1}^N f(z_K(x^i)).$$

- P. Ablin, T. Moreau, M. Massias, A. Gramfort, *Learning step sizes for unfolded sparse coding*, NeurIPS, 2019.
- M. X. Goemans and D. P. Williamson, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, J. ACM, 1995.
- P.-W. Wang, P. L. Donti, B. Wilder and J. Z. Kolter, *SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver*, ICML, 2019.