

Supervised learning of convex piecewise linear approximations of optimization problems

Laurine Duchesne, Quentin Louveaux and Louis Wehenkel

University of Liege - Dept of EE&CS
Liege - Belgium

Abstract. We propose to use input convex neural networks (ICNN) to build convex approximations of non-convex feasible sets of optimization problems, in the form of a set of linear equalities and inequalities in a lifted space. Our approach may be tailored to yield both inner- and outer- approximations, or to maximize its accuracy in regions closer to the minimum of a given objective function. We illustrate the method on two-dimensional toy problems and motivate it by various instances of reliability management problems of large-scale electric power systems.

1 Introduction

Mathematical optimization is a very rich framework allowing us to model lots of practical problems. The tractability of an optimization problem depends on the properties of the objective function and the feasible set. Non-convex problems are often intractable whereas convex problems are tractable. In particular, linear optimization, which is a subclass of convex problems, is a mature field where problems with thousands of variables and constraints can be routinely solved with efficient and reliable solvers [1].

In this paper, we propose to develop machine learning approaches that would allow us to automatically build convex approximations of non-linear and/or non-convex feasible domains in the form of a set of linear (and thus convex) constraints, in order to exploit the extremely efficient methods and solvers already available for linear programs. If the obtained linear approximation of the feasible set is an inner approximation (i.e. all the solutions belonging to this approximation are feasible), it would allow us to generate feasible solutions and, in the case of a convex objective function, upper bounds of the minimal value of the original problem. If it is an outer approximation (i.e. the approximation contains all the feasible solutions of the original problem) it would provide lower bounds, also in the case of a convex objective function. Note that this constraint on the objective function is not restrictive because generically, any optimization problem $\min_x f(x)$ s.t. $x \in \mathcal{D}$ may be rewritten as $\min_{(x,z)} z$ s.t. $z \geq f(x)$, $x \in \mathcal{D}$ so that any non-linear and/or non-convex optimization problem may be reduced to the minimization of a linear (thus convex) function subject to a possibly non-convex feasible set.

The ICNN [2] is a neural network with constraints on its parameters and activation functions implying that the learnt input-output function $h(x, \theta)$ is a convex function of the inputs x . While this method was originally proposed in a regression context (e.g. to build convex approximations of objective functions of

optimization problems), we propose to use it in a classification setting in order to build convex approximations of feasible sets of optimization problems. Note that this approach can also be of interest in the context of convex feasible sets. Using an ICNN would indeed allow us to enforce the convexity property when using supervised learning to approximate a convex feasible set.

The rest of this paper is organized as follows. Section 2 presents the main idea, i.e. the feed-forward ICNN model and our proposed adaptation to represent and learn convex approximations of a feasible domain, and how this learnt convex approximation can be used effectively in an optimization problem if the ICNN architecture is using piecewise linear activation functions such as ReLU or leaky-ReLU. Section 3 presents some illustrative experiments, section 4 discusses related works, and section 5 possible real-world applications and directions for future work. We provide a “supplementary materials” document [3], which gives further details about simulation experiments and mathematical proofs.

2 ICNNs for convex classification and optimization

We consider feed-forward networks as shown in Fig. 1, where $x \in \mathbb{R}^n$ denotes the vector of inputs, $\theta = \{W_i^z, W_i^x, b_i\}_{i=0, \dots, k-1}$ the set of parameters, and g_i the activation function used in the i th layer. The relationship between the inputs and the outputs of layer i of such a model is thus recursively given by:

$$z_{i+1} = g_i(W_i^z \times z_i + W_i^x \times x + b_i) \text{ for } i = 0, \dots, k-1,$$

with $z_0 = 0$ and (the outputs) $h(x, \theta) = z_k$.




Fig. 1: Representation of the layers of an ICNN with two outputs h_0 and h_1 . The weights W^z , colored in green, are constrained to be non-negative.

Notice that compared to a classical feed-forward neural network, *pass-through layers* connecting directly the input vector x to each layer have been added to increase the representation power of these networks. In the ICNN model [2], the weights W_i^z for $i = 1, \dots, k-1$ are constrained to be non-negative and the activation functions g_i are constrained to be convex and non-decreasing. These two conditions are sufficient to guarantee that the activations z_i of each layer and hence the outputs $h(x, \theta)$ are convex functions of the input vector x . In the rest of this paper we will use (convex and non-decreasing) *piecewise linear* activation functions g_i (such as ReLU or leaky-ReLU).

2.1 Convex set representation by an ICNN with two outputs

Among several possibilities, we decided to use an ICNN with two (scalar) outputs h_0 and h_1 to create a binary classifier, where an input is associated to the target class 0 if $g(x, \theta) = h_1(x, \theta) - h_0(x, \theta) \leq 0$.

With this choice, it is clear that the set $\tilde{\mathcal{D}}$ of elements classified in class 0 by the ICNN is a convex subset of \mathbb{R}^n , as soon as $g(x, \theta)$ is a convex function of x . In order to ensure this, we use identity activation functions for the output layer ($g_{k-1}(x) = x$) and impose an additional constraint on the parameter vectors $W_{k-1}^{z,0}$ and $W_{k-1}^{z,1}$ feeding the output layer of the ICNN: they should satisfy component-wise the inequality

$$W_{k-1}^{z,1} \geq W_{k-1}^{z,0}. \quad (1)$$

Notice that if we feed such a network with an extended vector $x_e = (x, -x)$ (or, more generally $x_e = Ax$), the input-output relationship remains convex in x .

2.2 Building a family of nested convex sets

To build a convex approximation of a set $\mathcal{D} \subset \mathbb{R}^n$, we assume that we have (or that we can build) a dataset of input-output pairs $\hat{\mathcal{D}} = \{x^i, y^i\}_{i=1}^n$, where each input x^i describes the coordinates of a point in \mathbb{R}^n and where the corresponding output $y^i = 0$ if the point x^i belongs to \mathcal{D} and $y^i = 1$ otherwise.

We propose to learn from the dataset $\hat{\mathcal{D}}$ the parameters θ of an ICNN classifier which has as inputs $x_e = (x, -x)$, and by using the cross-entropy loss

$$\text{loss}(\theta, x, y) = -\log \left(\frac{\exp(h_y(\theta, x))}{\exp(h_0(\theta, x)) + \exp(h_1(\theta, x))} \right).$$

After training, we consider the whole family of (convex) sets

$$\tilde{\mathcal{D}}_\lambda = \{x \in \mathbb{R}^n | g(x, \theta) = h_1(x, \theta) - h_0(x, \theta) \leq \lambda\},$$

with $\lambda \in \mathbb{R}$, as candidate convex approximations of \mathcal{D} .

2.3 Exploitation in the context of optimization

A convex ICNN classifier can be used to approximate the feasible set \mathcal{D} of an optimization problem. If the objective $f(x)$ is convex, the approximated problem

$$\min_{x \in \tilde{\mathcal{D}}_\lambda} f(x), \quad (2)$$

is then also convex. Furthermore, if the objective function $f(x)$ is (piecewise) linear (and convex), and if all the activation functions g_i used in the ICNN are piecewise linear, convex, and non-decreasing functions (such as $\text{ReLU}(x) = \max(0, x)$, or $\text{leaky-ReLU}(x) = \max(0.01x, x)$), we can show that the resulting optimization problem reduces to a linear program (see section 2 of the supplementary material [3]). In general, we can use convex ICNN classifiers to approximate in a linear fashion any non-convex part of the constraints and/or objective function of any optimization problem. Solving (2) for an increasing sequence of λ values, would yield a decreasing sequence of optimal values.

3 Illustrations

We consider some (convex and non-convex) toy problems where $\mathcal{D} \subset \mathbb{R}^2$. For each one, we used a dataset of 20,000 labelled points. These points were sampled uniformly in a square of length 10 centered at $(0,0)$; 16,000 were used to train and 4,000 to test. We show results with ICNNs of 6 hidden layers and 50 neurons per layer, and ReLU activations. The ADAM optimizer [4] with a learning rate of 10^{-3} was used to update the network parameters at each epoch. To enforce the non negativity condition for the weights W^z , each negative element of the update computed with the optimizer is set to 0 before the next iteration. Similarly, to enforce the convexity constraint on the last layer, we set to 0 each pair of weights for which the convexity condition (1) is not met. Before training, the inputs are standardized based on their minimum and maximum values in the training set to be in the range $[0, 1]$.

Fig. 2 shows four regions \mathcal{D} and their approximation $\tilde{\mathcal{D}}_0$ with an ICNN. One can see that the convex region is well approximated by the ICNN. For the non-convex domains, the ICNN provides a convex approximation $\tilde{\mathcal{D}}_0$ of \mathcal{D} .




Fig. 2: The set \mathcal{D} is represented in red and the approximated set $\tilde{\mathcal{D}}_0$ in blue.




Fig. 3: Effect of λ on the $\tilde{\mathcal{D}}_\lambda$ set.

It is possible to play with the size of the approximated region by modifying the threshold λ in the definition $\tilde{\mathcal{D}}_\lambda = \{x \in \mathbb{R}^n | h_1(x) - h_0(x) \leq \lambda\}$. Fig. 3 shows the approximated region $\tilde{\mathcal{D}}_\lambda$ for various λ . We see that increasing the threshold allows us to find outer approximations of \mathcal{D} and decreasing the threshold allows us to find inner approximations. With this method, the model needs only to be learnt once and then the threshold can be manually adjusted to obtain inner or outer approximations.

3.1 Considering an objective function when learning the ICNN

In an optimization context, it is of interest to guide the learning of the ICNN in order to improve the approximation close to the minimum values of the objective function. For that, one possibility is to exploit the objective function $f(x)$ in the definition of the loss function used when training the classifier, by giving less weight to elements of the training set farther from the optimum, and thus induce the learnt approximation $\tilde{\mathcal{D}}_0$ to be tighter near to the sought optimum. More details about this procedure can be found in section 3 of the supplementary materials companion document [3].

Fig. 4 shows decision boundaries thus obtained for different objective functions. We see that the approximation is indeed tighter close to the unconstrained optimum and actually approaches the constrained optimum very well.




Fig. 4: Left parts of (a) and (b): training weights $w_{y,f}$ (blue/dark blue: $w_{y,f} = 1$; white: $w_{y,f} = 0.2$). Right parts: resulting $\tilde{\mathcal{D}}_0$, where red and black crosses indicate respectively the unconstrained and the constrained minimum of $f(x)$.

4 Related works

This work is not the first one to exploit the particular properties of ICNNs in the context of optimization problems. ICNNs have been used for complex physical systems control [5, 6, 7], to learn the objective function of an optimization problem and/or its constraints. In these papers, the considered case-studies are convex. Using ICNNs is therefore a way to exploit the prior knowledge of convexity while offering tractable control methods.

Similar to our method, ICNNs are used [8] in the context of non-convex optimization. The authors developed an algorithm called the Convex Difference Neural Network that expresses the learnt function as a difference of convex functions, so that they can use Difference of Convex programming techniques in problems where their algorithm has been used to learn objective functions and constraints of optimization problems.

Compared to this research, our method learns convex (actually linear) approximations of non-convex optimization problems, at the cost of possibly larger approximation errors but with the advantage that much more efficient and scalable optimization solvers can be used.

5 Conclusions and future works

We propose to use ICNNs to learn convex approximations of feasible sets of general optimization problems. This approximation reduces to a series of linear inequalities when ReLU activation functions are used, and we showed how the model may yield outer or inner approximations and/or tight approximations in regions near the optimum of a given objective function.

The next step is to test this method on practical non-convex optimization problems. Depending on the context in which this method could be applied, the found solution could be used directly or as a warm-start point for solving the non-convex problem. It can also be used to compute lower or upper bounds of the optimal solution when the learnt approximation is built so as to obtain an outer or an inner approximation of the feasible region.

Another direction of research is to consider the case where the feasible set \mathcal{D} to be approximated depends on some external parameters ξ . In electric power systems, for instance, the secure region of operation depends on load and renewable generation levels and so ξ could represent the realizations of these exogenous uncertainties. One could then learn a convex approximation of the parameterized domain $\mathcal{D}(\xi)$ with an ICNN that would have both x and ξ as inputs while being constrained to be convex only in x (in [2] such models are called *Partially Input Convex NNs*). If the ICNN is able to capture the relationship between the shape of the feasible region and the parameters ξ , then it could largely speed up the solving of this type of problems with varying ξ values, once it is learnt.

A further direction of research would consider distributed optimization problems, where the proposed approach could be used to enable various agents to learn and exchange convex approximations of their subsets of constraints and their sub-objectives. This would be of extremely high relevance to the field of multi-area electric power systems planning and operation.

References

- [1] Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge University Press, 2004.
- [2] Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. In *International Conference on Machine Learning*, pages 146–155. PMLR, 2017.
- [3] Laurine Duchesne, Quentin Louveaux, and Louis Wehenkel. Supplementary material of ‘Supervised learning of convex piecewise linear approximations of optimization problems’, 2021. https://people.montefiore.uliege.be/l Duchesne/research/icnn_appendix.pdf.
- [4] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [5] Yize Chen, Yuanyuan Shi, and Baosen Zhang. Optimal control via neural networks: A convex approach. *arXiv preprint arXiv:1805.11835*, 2018.
- [6] Yize Chen, Yuanyuan Shi, and Baosen Zhang. Data-driven optimal voltage regulation using input convex neural networks. *Electric Power Systems Research*, 189:106741, 2020.
- [7] Shu Yang and B Wayne Bequette. Optimization-based control using input convex neural networks. *Computers & Chemical Engineering*, 144:107143, 2021.
- [8] Parameswaran Sankaranarayanan and Raghunathan Rengaswamy. CDiNN-convex difference neural networks. *arXiv preprint arXiv:2103.17231*, 2021.