WHEN DEEP LEARNING MEETS CONSTRAINTS

Buyun Liang[1], *Ryan de Vera*[1], *Tim Mitchell*[2], *Ju Sun*[1]

[1]Computer Science & Engineering, University of Minnesota, Minneapolis, USA
[2] Queens College, City University of New York, New York City, USA

ABSTRACT

This tutorial focuses on practical ways to handle constraints in deep learning and its applications. We will start with constraints that can be absorbed into deep neural networks, then move to simple constraints that allow projected-gradient style algorithms. For nontrivial constraints, we will discuss standard numerical methods such as penalty methods and augmented Lagrangian methods. Our tutorial will culminate with the introduction of `NCVX` (`https://ncvx.org/`), a general-purpose optimization package we have built to solve generic constrained deep learning problems. We will draw concrete examples from various scientific and engineering domains such as robust recognition, structure design, physics-aware machine learning, imbalanced learning, fairness, and discrete neural-network training to help the audience to understand and apply these practical numerical methods.

1. RATIONALE

Imposing explicit constraints is relatively new but increasingly pressing in deep learning, stimulated by, e.g., trustworthy AI that performs robust optimization over complicated perturbation sets [1] and scientific and engineering applications that need to respect physical laws and constraints [2]. However, it can be hard to reliably solve constrained deep learning problems without optimization expertise. Existing deep learning frameworks, such as TensorFlow and PyTorch, do not admit constraints. General-purpose optimization packages can handle constraints but do not perform auto-differentiation and have trouble dealing with nonsmoothness.

In this tutorial, we will introduce various practical ways (e.g., projected gradient methods, penalty methods, augmented Lagrangian methods) to solve deep learning problems with constraints. In particular, we will highlight a user-friendly optimization package `NCVX` [3, 4] that we have built specifically for solving constrained deep learning painlessly, and discuss practical tricks to speed up its convergence in applications [5]. We will embed our description of the computational techniques into various applications of constrained deep learning in science and engineering.

As far as we are aware, this is the first time constrained deep learning—which prevails in modern AI foundations & applications—be discussed systematically together with practical solution recipes. The proposed tutorial is perfectly aligned with the theme of ICCASP'23: "signal processing in the artificial intelligence era". We believe our extensive research experience in solving constrained deep learning problems in computer vision, structure design, and imbalanced learning [1, 2], our strong integration of expertise in machine/deep learning, signal processing, numerical optimization, computer vision, inverse problems, and computational imaging, and our first-of-its-kind `NCVX` computing framework for constrained deep learning [5] guarantee a unique interaction and learning experience for the broad signal processing community, especially those interested in applying modern deep learning tools in any area of signal processing.

By the end of the tutorial, the audience is expected to be able to quickly choose appropriate ways to handle constraints they may encounter when applying deep learning, and in particular, become comfortable using the user-friendly `NCVX` framework to tackle complicated constraints.

2. TUTORIAL STRUCTURE

We will introduce both scientific and engineering applications leading to constrained deep learning problems with nontrivial constraints, and practical numerical methods to solve them.

1. Background and Motivation (30 mins)
   (a) Motivating examples: Robustness in vision recognition, AI for science (15 mins)
   (b) Challenges: Reliably solving them requires optimization expertise (15 mins)
2. Concrete Examples of Constrained Deep Learning & Tailored Numerical Methods for Solving Them (50 mins)
   (a) General ideas for handling constraints (20 mins): simple constraints into neural network models, projected gradient, penalty methods, min-max saddle point methods, augmented Lagrangian methods
   (b) Example I: Robustness in vision recognition (6 mins)

(c) Example II: Neural structural optimization (6 mins)
(d) Example III: Knowledge-aware machine learning (6 mins)
(e) Example IV: Orthogonal recurrent neural networks (6 mins)
(f) Example V: Learning with imbalanced data & Fairness (6 mins)

3. Break (20 min)
4. `NCVX`: A General-Purpose Software Package for Constrained Deep Learning (60 min)
   (a) Brief review of algorithms of `NCVX` (20 min)
   (b) Practical tricks to speed up convergence (20 min)
   (c) Constrained DL examples in `NCVX` (20 min)
5. Open Problems and Frontiers (20 min)
   (a) Challenges & open problems (10 min)
   (b) Questions and answers (10 min)

3. ABOUT THE PRESENTERS

Buyun Liang is a MS student of computer science at UMN, where he worked as a graduate researcher at the GLOVEX group, led by Prof. Ju Sun. Previously he obtained his bachelor's degree in physics at Nanjing University, and also a master's degree in materials science at UMN, where his research focus is about Monte-Carlo and molecular dynamics simulation. He is the lead author of `NCVX`, the general-purpose software package targeted at constrained deep learning. He also focuses on customizing `NCVX` for different practical problems, such as robustness for vision recognition and AI for science. See `https://buyunliang.org` for more info.

Ryan Devera is a first-year PhD student in Computer Science & Engineering, UMN, working with Prof. Ju Sun on constrained deep learning and AI for science and engineering at large. Before this, he worked for eight years as a senior data scientist, project manager, and technical mentor in various start-up companies. He holds a master degree in applied mathematics and bachelor degrees in mathematics and physics.

Prof. Tim Mitchell is an assistant professor of computer science at Queens College/CUNY. His research interests span the areas of optimization, numerical linear algebra, and scientific computing, with one focus being computing and optimizing robustness properties of linear dynamical systems. He is also interested in nonsmooth constrained optimization, machine learning, and model-order reduction. He was a postdoc at the Max Planck Institute in Magdeburg, Germany and the Courant Institute at NYU, which is where he did his PhD, and he previously worked at IBM Thomas J. Watson Research Center in Hawthorne, New York. For more info, see `http://www.timmitchell.com`.

Prof. Ju Sun is an assistant professor at the Department of Computer Science & Engineering, the University of Minnesota at Twin Cities. His research interests span computer vision, machine learning, numerical optimization, data science, computational imaging, and healthcare. His recent efforts are focused on the foundation and computation for deep learning and applying deep learning to tackle challenging science, engineering, and medical problems. Before this, he worked as a postdoc scholar at Stanford University (2016-2019), obtained his Ph.D. degree from Electrical Engineering of Columbia University in 2016 (2011-2016), and B.Eng. in Computer Engineering (with a minor in Mathematics) from the National University of Singapore in 2008 (2004-2008). He won the best student paper award from SPARS'15, honorable mention of doctoral thesis for the New World Mathematics Awards (NWMA) 2017, and AAAI New Faculty Highlight Programs 2021. For more info, see `https://sunju.org/`.

4. TUTORIAL DELIVERY EXPERIENCE

None of the presenters has delivered any tutorial of similar nature. Prof. Sun has given talks about the topic in several places. He has also taught graduate-level deep learning and machine learning courses at UMN over the past 3 years. Profs. Sun and Mitchell will cover all the theoretical parts, and Liang & de Vera will cover the hands-on coding examples.

5. HISTORY OF THE TUTORIAL

This tutorial will also be given at SIAM International Conference on Data Mining (SDM'23) in June 2023, with content customized for the data mining community. We also plan to propose and hold similar tutorials in other top computer vision, machine learning, and relevant science and engineering conferences, with application contents tailored to the respective domains.

6. REFERENCES

[1] Hengyue Liang, Buyun Liang, Ying Cui, Tim Mitchell, and Ju Sun, "Optimization for robustness evaluation beyond ℓ_p metrics," *arXiv preprint arXiv:2210.00621*, 2022.

[2] Buyun Liang, Ryan de Vera, Tim Mitchell, Ying Cui, Qizhi He, and Ju Sun, "Deep structural optimization with principled constrained optimization," *Forthcoming*, 2022.

[3] Buyun Liang, Tim Mitchell, and Ju Sun, "Ncvx: A user-friendly and scalable package for nonconvex optimization in machine learning," *arXiv preprint arXiv:2111.13984*, 2021.

[4] Buyun Liang, Tim Mitchell, and Ju Sun, "Ncvx: A general-purpose optimization solver for constrained machine and deep learning," *arXiv preprint arXiv:2210.00973*, 2022.

[5] Buyun Liang, Hengyue Liang, Ying Cui, Tim Mitchell, and Ju Sun, "Ncvx: A general-purpose optimization solver for machine learning, and practical tricks," *Forthcoming*, 2022.