Convex Optimization in Computer Graphics

L. Mattos Da Silva[1]

[1]Massachusetts Institute of Technology, USA

Abstract

A number of tasks in computer graphics can be conceived as critical point conditions for an optimization problem. These optimization problems, however, often involve nonlinear or nonconvex formulations that cannot be solved easily with standard tools. In this course, we will go over how convex relaxation techniques can make solving these optimization problems more efficient. In particular, we will explore how convex optimization is used to solve for shape matching, contour models, geodesic distances, PDEs, and optimal transport tasks in computer graphics. We will also cover modern convex optimization software tools. The goal of the course is to equip students with a beginner's toolkit to apply convex optimization strategies to problems that they might encounter in their own research. All course materials will be available at https://convex-optimization-graphics.github.io/.

CCS Concepts

• ***Computing methodologies*** **→ *Computer graphics;*** • ***Mathematics of computing*** **→ *Convex optimization; Solvers;***

1. Introduction

Computer graphics is an area full of challenging mathematical problems. From finding shape correspondence maps and solving optimal transport problems to computing geodesic distances on meshes, many algorithms and pipelines in graphics and geometry processing rely on having robust numerical methods. A noteworthy feature of problems in our area is that they often can be reformulated as optimization problems.

However, simply formulating a problem as an optimization task does not guarantee a solution. Many geometric problems are inherently non-convex, leading to local minima and sensitivity to initialization. This tutorial provides a comprehensive introduction to *convex relaxation*, a powerful strategy whereby difficult non-convex problems are approximated by convex ones. Convexity offers several significant advantages: for instance, it provides certificates of global optimality, and it allows the researcher to leverage a vast library of efficient algorithms and modern off-the-shelf software.

This tutorial bridges the gap between the theory of convex optimization and practical applications in graphics and geometry processing. We will show how convex relaxation techniques enable the researcher to prototype faster, debug more effectively, and solve large-scale problems with relative ease.

2. Tutorial Length and Outline

This is a half-day tutorial (2×90 minutes), structured as follows:

Part I: Basics of Convex Optimization

1 Fundamentals of optimization problems
2 Defining convexity for objectives and constraints
3 Motivating convexity
4 Standard convex problems: LP, QP, SOCP, SDP

Coffee Break

Part II: Convex Relaxation in Graphics

5 Viscosity solutions for parabolic PDE
6 Mapping problems: Procrustes analysis and soft maps
7 Optimal transport for effective interpolation
8 Sum-of-squares relaxation for mesh repair
9 Hidden convex substructures for physical simulation

3. Intended Audience

The intended audience is primarily graduate students. We expect students to have a working knowledge of linear algebra (e.g., matrices, eigensystems, etc.) and multivariate calculus (e.g., gradients, Hessians, etc.). Familiarity with basic optimization concepts and fundamentals of geometry processing is helpful but not required. The tutorial is designed to be accessible to young researchers looking to add convex optimization methods to their toolkit.

4. Syllabus of Instruction

Part I: Basics of Convex Optimization

In the first half of the course, we will review the basics of optimization, define the notion of convexity, and discuss its importance. We will then examine a range of standard convex optimization problems, organized according to a clear taxonomy. For each problem,

DOI: 10.2312/egt.20261005

we will present its standard formulation and provide representative examples of interest in computer graphics.

4.1. The Fundamentals

We begin by formalizing the general notion of an optimization problem. We introduce canonical historical examples from the graphics literature, such as snake models [KWT88]. We will discuss standard approaches like gradient descent and highlight its known drawbacks, particularly its sensititivity to initialization when applied to non-convex minimization problems.

4.2. Defining Convexity

We define convexity as a property of both the objective function and the constraint set. For an optimization problem to be convex, we require that it has: (1) a convex objective function, and (2) a convex set of constraints, or equivalently a convex feasible region. We will provide mathematical and pictorial examples to clarify what distinguishes convex problems from non-convex problems.

4.3. Why Convexity?

We will discuss the practical benefits of having convex formulations for optimization problems. First, convexity provides a certificate of global optimality, eliminating ambiguity from the result of an optimization. Second, a vast array of algorithms are designed specifically for convex optimization, allowing for efficient solutions to large and complex problems.

We will also discuss the utility of modern convex optimization software (e.g., CVXPY, Mosek, etc.). These tools enable researchers to experiment with research ideas quickly, without having to implement heavy-weight custom solvers from scratch. We will demonstrate how these tools can be used for debugging: comparing a custom solver's output against a reference solution from an off-the-shelf convex tool can help identify bugs in implementation. We will showcase a live example using a Jupyter notebook, solving an optimization problem with just a few lines of code.

4.4. Standard Convex Problems

Having introduced the notion of convexity and motivated its study, we will then present the following simplified taxonomy of standard convex problems:

- **Linear Programming (LP)**: Problems with linear objectives and constraints. We will discuss applications in optimal transport, such as Earth mover's distance [SRGB14].
- **Quadratic Programming (QP)**: Problems with positive semidefinite quadratic objectives and linear constraints. As an example, we will examine linear blending weights for the transformation of animated characters [JBPS11].
- **Quadratically Constrained Quadratic Programming (QCQP)**: Problems with positive semidefinite quadratic objectives and constraints. We will discuss how QCQPs arise in the resolution of Procrustes matching problems [MDK*16].
- **Second-Order Cone Programming (SOCP)**: Problems with linear objectives and linear and cone constraints. We explore SOCP in the context of optimal transport with quadratic cost [LCCS18].
- **Semidefinite Programming (SDP)**: Problems with a linear objectives and linear matrix inequality constraints. We review SDPs involving linear matrix inequalities, citing examples in volumetric mesh deformation [KABL14].

Part II: Convex Relaxation in Graphics

In the second half, following the coffee break, we will delve into specific convex relaxation strategies used in computer graphics research. These strategies are categorized into broad groups based on the applications for which they are most relevant.

4.5. Viscosity Solutions

We apply convex relaxation techniques to the problem of computing geodesic distances. A typical approach involves solving the Eikonal equation, a nonliner PDE that can exhibit non-smooth gradients and shocks. We discuss the work [BF20], where a convex optimization method is proposed to approximate the solution of the Eikonal equation on discrete meshes. We also discuss the notion of regularized geodesic distance functions [EGSBC23] and convex relaxation strategies for solving associated second-order parabolic PDEs, such as the Fokker–Planck equation [MDSSS24].

4.6. Mapping Problems

We revisit the Procrustes matching problem, explaining how this non-convex quadratically constrained quadratic problem can be relaxed into a semidefinite program (SDP) at the cost of exactness [MDK*16]. This constitutes an example in which the solution to the relaxed problem is not necessarily the solution of the original problem. Additionally, we discuss the notion of soft maps [SNB*12], where instead of trying to compute point-to-point correspondences between surfaces, the idea is to relax the problem by allowing soft or probabilistic maps.

4.7. Optimal Transport

We analyze the transition from the non-convex Monge problem for computing deterministic maps to the convex Kantorovich relaxation, which permits couplings characterized by mass splitting [SRGB14, SDGP*15]. The big idea here is that replacing the transport maps by couplings π yields an easier-to-solve convex linear program. We further discuss the Benamou–Brenier dynamical formulation and its discretization on triangle meshes [LCCS18], which results in a convex optimization problem involving physically meaningful quantities like densities and momenta.

4.8. SOS Relaxation

We discuss sum-of-squares (SOS) relaxation, which is a powerful technique for converting polynomial inequalities into semidefinite programs. The idea is that nonnegative functions can be approximated using sums of squares of polynomials. We illustrate this method with the problem of repairing hexahedral meshes, which often have tangled or degenerate elements that are hard to fix using traditional optimization [MPZS20]. SOS relaxation can be used to certify positivity of mesh element Jacobians and correct elements if necessary.

4.9. Convex Substructures

Finally, we discuss the notion of hidden convexity. Some non-convex problems possess convex substructures that are not apparent at first glance; for instance, an optimization problem could be convex in a particular choice of variables even if the full problem remains non-convex. We explain how techniques like polar decomposition can isolate non-convexities, such as rotations, from convex parts, such as symmetric stretch matrices. This strategy is applicable to commonly used distortion energies like ARAP and Symmetric Dirichlet [SLS22] and variational schemes for elastodynamic simulation [MDSSPTS25].

5. Course Notes

Comprehensive course materials, including slides, notes, interactive Jupyter notebook, and a bibliography of suggested reading, will be provided. After the conference, all materials will be made available at `https://convex-optimization-graphics.github.io/`.

6. Resume of Presenter

Leticia Mattos Da Silva is a Ph.D. student at MIT, working in the Geometric Data Processing Group under the supervision of Justin Solomon. Her research focuses on developing numerical algorithms to solve nonlinear PDEs that arise in geometry processing. Her paper [MDSSS24] demonstrates how a broad class of parabolic PDEs can be solved via convex relaxation strategies compatible with modern solvers. Even more recently, her work [MDSSPTS25] has uncovered "hidden" convexity in optimization problems resulting from variational schemes for nonlinear elastic deformation.

Ms. Mattos Da Silva is the recipient of two MathWorks Fellowships, as well as a Google Fellowship. She holds a B.S. in Mathematics from the University of California, Los Angeles, and in summers past, she has conducted research as part of internships at the Flatiron Institute, at Adobe, Inc. and at the Fields Institute.

Email: `leticiam@mit.edu`
Website: `https://www.lmattos.com/`

7. Related Tutorials

The most recent related tutorial at Eurographics was *An Introduction to Optimization Techniques in Computer Graphics* [IGG*14], which focused on broader optimization techniques. In contrast, this tutorial will be dedicated to convex optimization and techniques related to convexity.

Related courses outside Eurographics include *Optimization in Geometry Processing* by Solomon and Claici [SC19] and an earlier edition by Bommes and Solomon [BS16]. These courses also focused on broader optimization techniques. Since then, convex optimization techniques and their applications have evolved in the field of geometry processing as well as graphics.

Earlier edition. A similar version of this tutorial was presented at the Symposium on Geometry Processing (SGP) 2025 (Course page: `https://school.geometryprocessing.org/summerschool-2025/index.html#course2`). That iteration of the tutorial was limited to a brief examples due to time constraints (90 minutes total). The Eurographics 2026 offering described herein is expanded to include a detailed module on using modern solver software like CVXPY, an interactive Jupyter notebook that allows the audience to code along, and deeper coverage of recent papers in the field.

8. Conclusion

The goal of this tutorial is to introduce researchers in graphics and geometry processing to important convex relaxation strategies and hopefully inspire them to also make use of these important techniques in their own research. By the end of the tutorial, attendees should have an understanding of why we would care about convex relaxation, how to use it in practice, and what applications in computer graphics have benefited from the use of these strategies.

References

[BF20] Belyaev A., Fayolle P.-A.: An admm-based scheme for distance function approximation. *Numerical Algorithms 84*, 3 (2020), 983–996. 2

[BS16] Bommes D., Solomon J.: Optimization in geometry processing. URL: `https://school.geometryprocessing.org/summerschool-2016/summerschool/presentations.html#course3`. 3

[BV04] Boyd S., Vandenberghe L.: *Convex optimization*. Cambridge university press, 2004.

[CWW13] Crane K., Weischedel C., Wardetzky M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (TOG) 32*, 5 (2013), 1–11.

[EGSBC23] Edelstein M., Guillen N., Solomon J., Ben-Chen M.: A convex optimization framework for regularized geodesic distances. In *ACM SIGGRAPH 2023 Conference Proceedings* (2023), pp. 1–11. 2

[IGG*14] Ihrke I., Granier X., Guennebaud G., Jacques L., Goldlücke B.: An Introduction to Optimization Techniques in Computer Graphics. In *Eurographics 2014 - Tutorials* (2014), Holzschuch N., Myszkowski K., (Eds.), The Eurographics Association. `doi:/10.2312/egt.20141019`. 3

[JBPS11] Jacobson A., Baran I., Popović J., Sorkine O.: Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph. 30*, 4 (July 2011). URL: `https://doi.org/10.1145/2010324.1964973`, `doi:10.1145/2010324.1964973`. 2

[KABL14] Kovalsky S. Z., Aigerman N., Basri R., Lipman Y.: Controlling singular values with semidefinite programming. *ACM Trans. Graph. 33*, 4 (July 2014). URL: `https://doi.org/10.1145/2601097.2601142`, `doi:10.1145/2601097.2601142`. 2

[KLF11] Kim V. G., Lipman Y., Funkhouser T.: Blended intrinsic maps. *ACM transactions on graphics (TOG) 30*, 4 (2011), 1–12.

[KWT88] Kass M., Witkin A., Terzopoulos D.: Snakes: Active contour models. *International Journal of Computer Vision 1*, 4 (1988), 321–331. URL: https://doi.org/10.1007/BF00133570, doi:10.1007/BF00133570. 2

[LCCS18] LAVENANT H., CLAICI S., CHIEN E., SOLOMON J.: Dynamical optimal transport on discrete surfaces. *ACM Transactions on Graphics (TOG) 37*, 6 (2018), 1–16. 2

[MDK*16] MARON H., DYM N., KEZURER I., KOVALSKY S., LIPMAN Y.: Point registration via efficient convex relaxation. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 1–12. 2

[MDSSPTS25] MATTOS DA SILVA L., SELLÁN S., PACHECO-TALLAJ N., SOLOMON J.: Variational elastodynamic simulation. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (2025), pp. 1–11. 3

[MDSSS24] MATTOS DA SILVA L., STEIN O., SOLOMON J.: A framework for solving parabolic partial differential equations on discrete domains. *ACM Trans. Graph. 43*, 5 (June 2024). URL: https://doi.org/10.1145/3666087, doi:10.1145/3666087. 2, 3

[MMP87] MITCHELL J. S., MOUNT D. M., PAPADIMITRIOU C. H.: The discrete geodesic problem. *SIAM Journal on Computing 16*, 4 (1987), 647–668.

[MPZS20] MARSCHNER Z., PALMER D., ZHANG P., SOLOMON J.: Hexahedral mesh repair via sum-of-squares relaxation. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 133–147. 3

[MZPS21] MARSCHNER Z., ZHANG P., PALMER D., SOLOMON J.: Sum-of-squares geometry processing. *ACM Transactions on Graphics (TOG) 40*, 6 (2021), 1–13.

[SC19] SOLOMON J., CLAICI S.: Optimization in geometry processing. URL: https://school.geometryprocessing.org/summerschool-2019/index.html#course2. 3

[SDGP*15] SOLOMON J., DE GOES F., PEYRÉ G., CUTURI M., BUTSCHER A., NGUYEN A., DU T., GUIBAS L.: Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (ToG) 34*, 4 (2015), 1–11. 2

[SLS22] STEIN O., LI J., SOLOMON J.: A splitting scheme for flip-free distortion energies. *SIAM Journal on Imaging Sciences 15*, 2 (2022), 925–959. 3

[SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1617–1626. 2

[SRGB14] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Earth mover's distances on discrete surfaces. *ACM Transactions on Graphics (ToG) 33*, 4 (2014), 1–12. 2

[SSK*05] SURAZHSKY V., SURAZHSKY T., KIRSANOV D., GORTLER S. J., HOPPE H.: Fast exact and approximate geodesics on meshes. *ACM transactions on graphics (TOG) 24*, 3 (2005), 553–560.

[SSVS25] SILVA L. M. D., SELLÁN S., VARGAS F., SOLOMON J.: Mirror bridges between probability measures, 2025. URL: https://arxiv.org/abs/2410.07003, arXiv:2410.07003.

[ZMST23] ZHANG P., MARSCHNER Z., SOLOMON J., TAMSTORF R.: Sum-of-squares collision detection for curved shapes and paths. In *ACM SIGGRAPH 2023 conference proceedings* (2023), pp. 1–11.