CISC371: Nonlinear Optimization

Fall, 2019

Calendar Description:

Methods for computational optimization, particularly examining nonlinear functions of vectors. Topics may include: unconstrained optimization; first-order methods; second-order methods; convex problems; equality constraints; inequality constraints; applications in machine learning.

Detailed Description:

When we use machine learning or data analytics methods on structured numerical data, we often want to find an optimal answer. For example, if we want to model the data as being linearly related, we might want to minimize the error between the model and the data.

In this course, we will explore data analysis where we have a single number that indicates how well the model represents the data. This number is the *objective* that we want to optimize. Some of the models that we will study include: neural networks; linear relationships within the data where we have a nonlinear objective, such as ridge regression and the "lasso"; and dividing data into two classes based on how the data is labeled, such as the support vector machine or SVM.

A guiding principle for this course will be that linear algebra is our main way of describing and implementing nonlinear analysis of structured data. Most other courses use a "calculus-first" principle; instead, we will use a limited amount of calculus to derive the linear algebra that we need for writing code. This "algebra-first" principle will allow us to explore complicated ideas in a concise and scalable manner by using linear algebra as the "language" of structured data.

The beginning part of the course will deal with mathematical preliminaries. We will explore two simple ways to optimize a function that depends on just one variable; these are approximation methods and search methods. Next, we will explore a few of the basic ways to optimize a function that depends on a vector of variables; these methods combine the one-variable methods into ways of performing *unconstrained* optimization. The first part ends with neural networks, where we will describe and implement the back-propagation algorithm by using matrix computations instead of the usual combination of calculus terms and summation loops.

The middle part of the course will deal with *constraints*, which are strict requirements that we place on the acceptable solutions to our problems. We will explore how to write simple constraints as "expanded" matrix problems, and then we will convert the constraints into "contracted" matrix problems. These ideas are based on *Lagrange multipliers* and *duality*, respectively. The middle part ends by extending duality to nonlinear constraints using the *KKT conditions*.

The final part of the course will deal with the application of constraints to problems that arise widely in data analysis. We will place constraints on simple regression problems that will help us to understand how ridge regression and the "lasso" solve these problems. We will explore how a single constraint – part of a family of *regularization* methods – can help us when we solve seemingly complicated problems in data modeling. The idea of duality will help us to use linear algebra to describe constrained solutions to problems of data classification. We will evaluate our results using current statistical methods for validation of regression and assessment of classification. The course ends by using kernels in an SVM to find nonlinear classifications of structured data.

Learning Outcomes:

By the end of the course, a successful student will be able to:

- **LO1:** Assess their implementation of shallow neural networks that are trained on complete data sets
- **LO2:** Interpret basic constrained optimization using Lagrange multipliers and the KKT conditions
- **LO3:** Validate constrained regression using randomized tests
- **LO4:** Interpret the sensitivity and specificity of classification methods
- **LO5:** Verbally present and defend the choice and method for solving a problem

These learning outcomes will be assessed by computer implementations, written English assignments, test questions, and video presentation by each student.

Textbook:

There is no required textbook for this course. A student may wish to purchase "Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB Algebra", by Amir Beck (SIAM Press 2014; ISBN: 978-1-61197-364-8). We will use other sources that are available through permissions granted to the Queen's University Library. Additional instructional material may be provided, as needed.

Grading Scheme:

The grading scheme will be finalized by the second week of the course. A draft is:

Assignments	A1@5%; A2, A3, A4 @11% each	38%
Video Presentation	Brief description and defense	2%
Tests, in class	lowest@6%; middle, highest @12% each	30%
Final examination		30%

	Tentative Test Dates
Test #1	Oct 4
Test #2	Oct 30
Test #3	Nov 22

Calculator Policy:

In accordance with Academic Regulation 9.2, only Casio 991 series calculators are pre-approved in tests. Any other calculator must be approved by the instructor at least one day before the test. A calculator needs only the basic five functions and must not have storage capability or communication capability. The use of a calculator is recommended but the tests are written with the intent that they can be answered without the use of a calculator.

Additional Syllabus Information

Please see the course's "Syllabus Addendum" for additional information on academic integrity, accommodation requests, extenuating circumstances, and relevant copyright protection.

Assignments:

There will be several compulsory written assignments, based on MATLAB coding and a clear English answer to the assignment questions, each assignment receiving a numerical grade. The assignments will have varying numbers of marks that reflect the expected difficulty of the assignments, with easier assignments having fewer marks that contribute towards the final grade. Each assignment will be graded by a teaching assistant (TA) so a TA should be contacted first for any question regarding grading.

The policy for late assignments will start immediately after the submission deadline, with the deadline being specified in the statement of the assignment. Marks will be deducted at a rate of 20% off the maximum assignment value for each 24 hours that the assignment is late. Academic consideration for assignments will be managed on an individual basis.

Requests for re-grading of an assignment must be made within 72 hours of completion of grading, which will be determined by the time in onQ that general feedback for the assignment is provided. This will help us to ensure that we address each student's concerns in a timely manner.

No Syllabus Changes:

Academic Regulation 7.2.1 requires that the instructor makes no changes to grade calculation "if the changes will disadvantage any student in the class". Accordingly, *no request for re-weighting or additional work will be granted* other than for extenuating circumstances as described in the Faculty policy on academic consideration.

Turnitin Statement:

This course makes use of Turnitin, a third-party application that helps maintain standards of excellence in academic integrity. Normally, students will be required to submit their course assignments through onQ to Turnitin. In doing so, students' work will be included as source documents in the Turnitin reference database, where they will be used solely for the purpose of detecting plagiarism.

Turnitin is a suite of tools that provide instructors with information about the authenticity of submitted work and facilitates the process of grading. Turnitin compares submitted files against its extensive database of content, and produces a similarity report and a similarity score for each assignment. A similarity score is the percentage of a document that is similar to content held within the database. Turnitin does not determine if an instance of plagiarism has occurred. Instead, it gives instructors the information they need to determine the authenticity of work as a part of a larger process.

Please read Turnitin's Privacy Pledge, Privacy Policy, and Terms of Service, which governs users' relationship with Turnitin. Also, please note that Turnitin uses cookies and other tracking technologies; however, in its service contract with Queen's Turnitin has agreed that neither Turnitin nor its third-party partners will use data collected through cookies or other tracking technologies for marketing or advertising purposes. For further information about how you can exercise control over cookies, see Turnitin's Privacy Policy.

Turnitin may provide other services that are not connected to the purpose for which Queen's University has engaged Turnitin. Your independent use of Turnitin's other services is subject solely to Turnitin's Terms of Service and Privacy Policy, and Queen's University has no liability for any independent interaction you choose to have with Turnitin.