Advanced Computational Methods in Economics and Finance

Syllabus for the course CS 5354/CS 4365, Fall 2018

Class time: MW 3-4:20 pm, Room BUSN 301

Instructor: Vladik Kreinovich, email vladik@utep.edu, office CCSB 3.0404, office phone (915) 747-6951

- The instructor's office hours are: Mondays 8:30-9 am and 2-3 pm, Wednesdays 10:30-12 pm and 4:30-5:30 pm, or by appointment.
- If you want to come during the scheduled office hours, there is no need to schedule an appointment.
- If you cannot come during the instructor's scheduled office hours, please schedule an appointment in the following way:
  - use the instructor's appointments page http://www.cs.utep.edu/vladik/appointments.html to find the time when the instructor is not busy (i.e., when he has no other appointments), and
  - send him an email, to vladik@utep.edu, indicating the day and time that you would like to meet.

  He will then send a reply email, usually confirming that he is available at this time, and he will place the meeting with you on his schedule.

Prerequisite:

- for graduate students: no special pre-requisites; graduate level standing is sufficient
- for undergraduate students: ideally, Statistics and Linear Algebra, but this is not required, we will recall the needed material anyway.

Main objectives: to learn advanced computational techniques used in solving economic and financial problems -- from the computational viewpoint, of course.

Why economics and finance are important: high-level perspective. The ultimate goal of science and engineering is to make the world a better place. Numerous innovations do make our lives better:

- cell phones and emails make it easier to communicate,
- commercial websites make buying easier.

However, in many cases, innovations come with negative side effects: e.g.,

- fracking makes energy cheaper but can lead to pollution,
- self-driving cars will probably make travel safer, but they may increase unemployment.

How to take into account everyone's interests? Modern economics and finance techniques enable us to formulate such questions in precise numerical terms -- after which we need to design and apply computational techniques to solve the resulting computational problems.

What we will do in the class in this regard. In this class, we will learn the basic computational ideas and techniques used in solving the corresponding problems.

Of course, we will only learn the basics. To really become a *quant* (a specialist in computational economics and finance), it is necessary to learn many technical details and tricks -- and it is not possible to cover all this in a one-semester course. However, what we will do is cover most basic ideas behind these tricks.

Why economics and finance are important: pragmatic perspective. In the real world, every business needs to ne profitable. The need to take economic and financial aspects into account influences how much effort we spend on a software: when we release it, how much we test it, how much efforts we can afford to spend on optimizing it.

It is not realistic to expect that every employee understand all the related economic and financial details, but having a basic understanding definitely helps one to become a more productive employee -- and improves the chances of moving up the ladder, to leadership positions.

What we will do in the class in this regard. Again, this class is not a substitute for real economics and business classes. However, some basic knowledge acquired in this class will hopefully help you better understand how companies function.

Important aspects of decision making in economics and finance. As we have mentioned earlier, one of our main objectives is to come up with strategies for group decision making, strategies that take into account interest of al the people involved.

In order to make these decisions, we need to have a good understanding of individual people's preferences and interests. Once we learn people's preferences, we can come up with algorithms that help people make decisions which best reflect these preferences. It is also important to take into account that when people actually make decisions, they often do not use complex optimization algorithms, they use their intuition which often leads only to sub-optimal decisions. It is therefore important to learn not only how people *should* make decisions, but also how they *actually* make decisions.

Whatever decisions we make, these decisions affect the future. Therefore, to make appropriate decisions, we must make reasonable predictions about the future state of economics.

- When we decide which company to work for, we are trying to predict which company is more probable to survive in the long run.
- When we invest money for retirement, we are trying to predict which companies' stocks is more probable to remain valuable.

To predict future values of corresponding quantities, we can use past values of this quantity and/or current (and past) values of related quantities.

These are all the problems that we will deal with in this class:

- prediction,
- individual decision making, and
- group decision making.

Traditional (basic) approach to prediction and decision making. The simplest predictions models are linear models, when the predicted value is estimated as a linear combination of the past and current values of one or several quantities. The coefficients of this linear combination must be determined based on the available data. The standard way of finding these coefficients is by minimizing the mean squared error. This *Least Squares* method will be the first thing we study in this class.

Once we can predict the values of different quantities, the next step is to make a decision that would maximize the corresponding objective function. The simplest objective functions are quadratic, so we will study how to optimize quadratic functions. Our first example will be on how to best invest money -- based on the 1950s portfolio optimization work of the Nobelist Harry Markowitz.

We will also discuss how Markowitz theory helps decrease medicines' side effects and speed up machine learning.

Need to go beyond traditional techniques. Traditional techniques assume:

- that the dependencies are linear,
- that we have full information about all the data, and
- that the optimization function is quadratic.

All these assumptions are simplifying: in real life,

- dependencies are often non-linear,
- we usually have only partial information, and
- objective functions are more complex.

To deal with real-life situations, we need to use advanced computational techniques. This is what we will study in this class.

In dealing with such complex problems,

- it is important not just to come up with an optimal solution;
- it is also important to have reasonably *simple* solutions, solution which can be easily used in the real company, and
- in view of the uncertainty, it is important to have *robust* solutions, solutions which work not just for some specific values of the corresponding coefficients, but also under possible deviations from these values.

Main ideas behind the advanced economic and financial techniques. Sometimes, to select a proper model or a proper algorithm, it is important to compare similar situations -- and/or similar representations of the same situation. For example, in physics, many fundamental equations can be derived from the natural requirement that the corresponding formulas not change if we simply change the measuring unit (e.g., from minutes to seconds). This *symmetry* approach is productive not only in physics, we will see that it is also productive in economics and finance.

Symmetry ideas can help to find the models if we already know the objective function. When we do not yet have a clear expression for the objective function, symmetry ideas can help to find such an expression.

In some cases, it helps to consider three or more different situations and to require *consistency*. A good example of such consistency is *additivity*: when several countries form a strong alliance -- like European Union (EU) -- then, e.g., the formulas for trade with EU should lead to similar results whether we consider EU as a single economic entity or as several different countries.

It also often helps to compare economic and financial situations with situations from other areas. For example, there are many similarities between physical and economic processes, so many, there there is a whole direction in economics, known as *econophysics*. Its latest ideas are to borrow ideas and techniques from quantum physics; this is known as *quantum econometrics*.

In the class, we will show how these ideas help us solve the problems related to prediction and decision making in economics and finance.

Specific topics covered in this class: general idea. Let us list specific topics covered in this class. Of course, this list is approximate. We may not have enough time to cover all of this, in which case we will follow the wise advise of one of my Russian colleagues: "It is better not to have time for everything than not to understand anything" ("Luchshe nichego ne uspet' chem nichego ne poniat'.")

In all these topics, the emphasis will be on the main ideas, but we will also write some code -- usually, for simplified situations and simplified techniques.

Specific topics covered in this class: prediction. How prediction works?

- we need to select a prediction model; such models usually comes with parameters that need to be determined from the data;
- based on the data, we need to find the values of these parameters which fit the data the best; for that, we need to describe fitness in precise terms, and then come up with efficient algorithms for finding the best-fit parameters.

Topics related to selecting a model:

- why linear models;
- which non-linear models should we choose: symmetry-motivated approach;
- first application: gravity model of trade;
- second application: how to predict production;
- state-of-the-art: main ideas behind quantum econometrics.

Topics related to selecting a probability distribution:

- symmetry-motivated distribution functions and their symmetry-motivated combinations;
- important case study: heavy-tailed Student distributions; see also;
- another case study: Matern covariance model;
- skew normal distributions;
- distributions of extreme values;
- *copulas*;
- symmetry-motivated objective functions: from Laplace indeterminacy principle to maximum entropy techniques and generalized entropy.

Algorithms. Depending on what information we have about the corresponding probability distributions, we need different algorithms:

- if we know the distribution -- or the family containing the actual distribution -- then we should use *maximum likelihood*;
- different techniques should be used if if for some error components, we only know upper bounds -- this case is called interval uncertainty;

- if we have no information about the probabilities, we should use *robust methods* like l^p-techniques or empirical likelihood methods; in general, robust predictive econometrics leads to more accurate predictions.

Robustness can also be used as a criterion for selecting a model.

For symmetry-motivated non-linear models, the corresponding symmetries help simplify the algorithms.

Specific topics covered in this class: ideal individual decision making. We will start with a brief overview of the traditional decision making theory, theory centered around the notion of utility. We will then show how symmetries help find the dependence of utility on several parameters.

We will then analyze how to make decisions under (interval) uncertainty. The main idea is Nobelist Leo Hurwicz's *optimism-pessimism criterion*. As an example, we will show how Markowitz's portfolio selection problem needs to be modified when we have no information about correlations.

Specific topics covered in this class: how people actually make decisions. According to the traditional decision theory, ideally, people should:

- take into account all available information,
- make adequate estimates of the corresponding probabilities, and
- select the alternative for which the expected utility is the largest.

In practice, due to the limited ability of human information processing, we:

- take only some information into account,
- use approximate estimates of probabilities, and
- instead of always selecting the best alternative, often select close-to-optimal ones with some probability.

In this class, we will consider, explain, and analyze three example of such behavior:

- peak-end rule when people only take into account the peak and the end experiences; this rule is related to Dow Peak-and-Trough Theory of stock market behavior;
- probability-related empirical weights discovered by Nobelish Daniel Kahneman, and
- Nobelist Daniel McFadden's description of probabilistic choice.

Specific topics covered in this class: group decision making. We start with the traditional approach to group decision making: Nash's bargaining solution. To illustrate this idea, we will use two examples:

- gauging the state of a country's economy; for this, the formula coming from Nash's bargaining solution is more adequate than the usual Gross Domestic Product (GDP) -- since this formula takes inequality into account;
- the bankruptcy problem; in this example, we will follow the work of Nobelist Robert Aumann.

What we do not cover at all. Conflict situations and related *game-theoretic* techniques are a whole separate topic, requiring a special class.

Main Source: there are may books on computational methods in economics and finance, but they are either too heavy on economics, or too heavy on mathematics. Instead, we will use handouts.

Projects: An important part of the class is a project. There are three possible types of projects:

- An ideal class project is if you do something related to the class which is useful for your future thesis or dissertation. Please check with your advisor about it, maybe he or she wants you to read and report on some economics-related paper, maybe you need to do some economics-related research, whatever your advisor recommends will be a very good project for this class, just let the class instructor know what exactly you plan to do.
- If you have not yet selected an advisor, but you already know what research area you want to work in, come talks to the class instructor, we will try to find some appropriate topic -- and if you have any proposals already, great.
- If you do not have a research topic or you have a one but your advisor cannot find anything economics-related that will be helpful for your future thesis or dissertation, come talk to the instructor too.
- Maybe you like economics and want to start doing a related project, then come and talk to the instructor, we will try to find something that will be of interest to you.

A project can be:

- reviewing and reporting on a related paper, or
- doing some independent research (not research as in high school, but research as in graduate school, i.e., trying to come up with something new), or
- programming something economics-related.

The most important aspect of the project is that it should be useful and/or interesting to *you*. The instructor can assign a project to you, there are plenty of potential projects, but if each student selects a project that he or she likes, this will be much much better for everyone.

Tests: There will be three tests and the final exam.

Grades: Each topic means home assignments (mainly on the sheets of paper, but some on the real computer). Some of them may be graded. Maximum number of points:

- first test: 10
- second test: 10
- third test: 15
- home assignments: 10
- final exam: 35
- project: 20

(smart projects with ideas that can turn into a serious scientific publication get up to 40 points).

A good project can help but it cannot completely cover possible deficiencies of knowledge as shown on the test and on the homeworks. In general, up to 80 points come from tests and home assignments. So:

- to get an A, you must gain, on all the tests and home assignments, at least 90% of the possible amount of points (i.e., at least 72), and also at least 90 points overall;

- to get a B, you must gain, on all the tests and home assignments, at least 80% of the possible amount of points (i.e., at least 64), and also at least 80 points overall;
- to get a C, you must gain, on all the tests and home assignments, at least 70% of the possible amount of points (i.e., at least 56), and also at least 70 points overall.

Standards of Conduct: You are expected to conduct yourself in a professional and courteous manner, as prescribed by the UTEP Standards of Conduct.

Graded work, e.g., homework and tests, is to be completed independently and should be unmistakably your own work (or, in the case of group work, your team's work), although you may discuss your project with other students in a general way. You may not represent as your own work material that is transcribed or copied from another person, book, or any other source, e.g., a web page.

Academic dishonesty includes but is not limited to cheating, plagiarism and collusion.

- Cheating may involve copying from or providing information to another student, possessing unauthorized materials during a test, or falsifying data (for example program outputs) in laboratory reports.
- Plagiarism occurs when someone represents the work or ideas of another person as his/her own.
- Collusion involves collaborating with another person to commit an academically dishonest act.

Professors are required to -- and will -- report academic dishonesty and any other violation of the Standards of Conduct to the Dean of Students.

Disabilities: If you feel you may have a disability that requires accommodation, contact the The Center for Accommodations and Support Services (CASS) at 747-5148, go to Room 106 E. Union, or e-mail to cass@utep.edu. For additional information, please visit the CASS website.

See You All in the Class!