Numerical Methods for Optimal Control in High Dimensions

The American Institute of Mathematics

This is a hard–copy version of a web page available through http://www.aimath.org
Input on this material is welcomed and can be sent to workshops@aimath.org
Version: Fri Aug 19 14:09:18 2005

Table of Contents

A.1 Arisawa, Mariko

We would like to introduce a viscosity solution's approach to a class of second-order partial-integro-differential equations (PIDE, in short) arising in the mathematical finance.

The usual brownian models for the processes of the asset prices are replaced by the jump-diffusion processes. The derived PIDEs contain in the integral term, various Le'vy measures which correspond to the frequencies and the sizes of the jumps. In addition, according to the financial settings such as the American option, the look back option, the problem with transaction costs, a variety of nonlinear PIDEs are engendered. We shall introduce a new definition of the viscosity solution for such problems, and shall give the comparison, the existence, and the regularity of the solutions of viscosity solutions in this framework.

As far as the high-dimensional numerical analysis, a very few is known at this stage, while from the practical view point a lot of important problems should be explored, such as the modelization of the high-dimensional Le'vy measures arising in the mathematical finances. This is the problem which I would like to discuss with the participants of the workshop.

A.2 Capobianco, Enrico

Part 1

Systems Biology: On Dimensionality Reduction and Feature Selection in Genomics (work carried out at Boston University, Biomedical Engineering Department)

Gene networks offer a wealth of data; this is mainly due to the genomic dimensionality rather than the samples, as the latter usually come from measurements obtained under only a few experimental conditions or time points. It is therefore a challenging task to design suitable statistical models and to develop effective reverse engineering algorithms.

Network inference and reverse engineering deal basically with an inverse problem that is stimulating a great deal of biomedical research in systems biology: reconstructing gene-gene interactions from measurements obtained under specific experimental conditions. Some of the open problems calling for solutions are now listed.

First, the signature of noise is pervasive in genetic networks. For instance, in perturbation experiments only a few genes change expression value, while most genes show either noisy or constant patterns. The measurements are usually taken at an initial time and then regularly recorded until the genes reach a steady state. During this time interval, the gene temporal patterns are subject to several kinds of fluctuations, and only some of them are induced by the perturbations, directly or indirectly as a cascade effect.

Noise is thus a strong conditioning factor for the gene temporal patterns; the ones that appear smooth are usually of interest compared to those characterized by a random signature. Consequently, sharp time-localized fluctuations lead to discard many patterns and focus just on subsets of genes.

An important tool is thresholding, which deals with the separation of signal from noise. It identifies gene groups by selecting the genes that are differentially expressed with regard to

noise-dependent patterns. Furthermore, a gene can be considered outlying if its expression value is significantly different (over- or under-expressed) compared to the average expression value computed for all the genes.

Methodological aspects

Statistically speaking, if we are able to isolate a certain number of modes or components which represent in a compact way the systems dynamics underlying the gene network under examination, we might also try to compute statistics based on these modes through the gene sets that have been identified by them.

Due to the inherent biological characterization of the estimated components, the thresholding step can validate or contradict the hypothesis that the observed expression values of the genes are significantly different from specific test statistics.

The idea is that for each available condition (a sample point where a measurement is taken or an experiment is conducted), a genomic profile X is considered a mixture (via a certain matrix A) of unknown biological influences (modes or components) S that may have more or less impact over each specific condition. Then, feature sets are identified by linear combinations of genes which are delivered by each identified and estimated component.

More formally, consider at time t an $\epsilon-$ noise genomic system such as:

$$X = AS + \epsilon \tag{1}$$

The dimensions of the model (for $k = 1$) are $S \in R^m$, $X \in R^n$, and $A \in R^{n \times m}$, and the gene expression matrix X has rows representing the different genes involved in the regulatory network under exam, while the columns list the measurements at successive conditions.

The mixture signals in X represent $k-$dimensional vectors $x_i, i = 1, \ldots, n$, and the components S represent $k-$dimensional vectors $s_j, j = 1, \ldots, m$. The latter are mixed up linearly (or non-linearly) through the mixing matrix A. Both S and A are unknown variables and must be estimated from the only available information set X.

This gene selection goal implies that a genetic regulatory network is a redundant system because inherently noisy, but also owing to the high-dimensionality and dependence between genes, only in part justified by biological reasons. Equivalently, the underlying system of gene dynamics is sparse; in other words, the gene-gene interaction matrix is only partially active, as some of the supposed links among genes are not supported by significant evidence.

Exploiting sparsity when designing statistical inference procedures for gene network analysis as a key goal is to target just subsets of genes when measuring the effects of the experiments. And sparsity refers to dimensionality too, through the possibility of approximating the truly intrinsic dimension of the input space. In turn, this may bring some valuable experimental insight too, by suggesting for instance how to calibrate the perturbations and selecting the genes expected to become targets.

In order to explore these two aspects, redundancy and sparsity, clustering has been widely employed by computational biologists, but several limitations have appeared; therefore, independent and principal component analysis (ICA and PCA, respectively) can also

be proposed (and compared) as examples of flexible approximation tools targeted to dimensionality reduction and gene feature selection.

Part 2

Semi-Parametric Estimation in In Vivo MR Spectroscopy (ongoing joint work with: H. Ratiney, C. Cudalbu, S. Cavassila, D. Graveron-Demilly RMN, CNRS UMR 5012, Univ. Claude Bernard LYON I-CPE, France R. de Beer, D. van Ormondt, Applied Physics, TU Delft, The Netherlands).

Magnetic Resonance Spectroscopy (MRS) is a unique tool for non-invasive in vivo detection and quantitation of metabolites. This makes MRS an indispensable tool for combating major diseases. Although modern MRS-methods are increasingly capable of detecting and quantifying metabolites, perturbation by signals from so-called macromolecules and unwanted metabolites poses problems.

Three procedures for alleviating the problem exist:

1) Tuning of the scanner to a metabolite of interest – 'spectral editing' – can clean up the MRS signal significantly.

2) Separate (approximate) measurement of the macromolecule signal and subsequent subtraction from the MRS signal.

3) Semi-parametric estimation of the model parameters of interest from the MRS signal.

Our strategy is to process directly in the measurement domain; in MRS, this is the time domain. Alternatively, one can process in the frequency or spectral domain. A main disadvantage of the latter is that it starts with estimation of the spectrum which is not trivial in the case of missing samples or non-Cartesian sampling.

Methodological aspects

In semi-parametric estimation, the task is to disentangle the parametric and non-parametric parts of a signal. In MRS, the parametric part pertains to the metabolites, the non-parametric part to the macromolecules. In the measurement domain, the metabolite signal persists over time much longer than the macromolecule signal. On the other hand, during the initial period, the latter strongly dominates the former. This property is exploited to bring about the disentanglement.

Through simulations, we have successfully investigated:

1) Iteration of the disentanglement procedure.

2) Derivation of Cramer-Rao Bounds (CRB) for the metabolite concentrations, taking the non-parametric macromolecule signal into account.

This contributes to experimental design. Furthermore:

1) Our previous contributions did not report on iteration of the disentanglement procedure. Now, we include cases where such iteration clearly helps.

2) We have devised a novel way (in MRS, at least) to estimate CRBs for the parametric part, taking the non-parametric part into account.

3) Only information about the point of time where the non-parametric part 'decays into the noise' is needed.

We thus have shown that through semiparametric models it is possible to augment the applicability of MRS to Medicine, while in order to have the best possible impact we need to optimize the accuracy of the estimates for the parameters of interest.

In particular, one seeks an improved reliability of the lower bounds computed on parameters to be estimated with a planned scan, from which it depends in turn an improved reliability of the 'experimental design', or, in other words, a better prediction of feasibility of expensive MRS-scans in the clinic.

Then, one would like to address:

A- what is the predicted minimum detectable concentration of each of the 40 or so metabolites of interest with a given scanner measurement protocol?

B- Is the try justifiable?

A.3 de Farias, Daniela

In this workshop, I am interested in learning more about the main open issues and approaches across the different subfields related to high-dimensional optimal control. In addition, I also see the workshop as an opportunity to discuss some of the open issues and explore potential directions in my own subfield.

My current research interests lie in approximate methods for large-scale stochastic dynamic programming, with particular focus on value function approximation. I would hope to see at least some activities and discussions pertaining to that topic. Some of the directions in approximate dynamic programming that I consider worth exploring are as follows:

- Approaches for validation of approximate DP methods. Interest in approximate DP has been fuelled by successful application in a variety of areas ranging from finance to games and helicopter control. Yet, it is still difficult to know a priori when, how and why a given method will be successful in tackling a particular problem. In my view, there is a need for more systematic approaches to validation of algorithms beyond illustration via application to one or two specific problems. These systematic approaches could be of many different kinds, which could be discussed in more or less technical detail at the workshop. For instance, one could pursue the development theoretical guarantees of performance and errors, or a more systematic empirical comparison based on a benchmark of large-scale DP problems.
- Comparison and unification of different algorithms. Approximate DP is currently represented by a conglomeration of ideas and algorithms, and it is often difficult to know how they relate to one another and which is most appropriate for any class of applications. A streamlined comparison among the approaches and potentially unifying theory would be useful developments to the field.
- Systematic and/or adaptive design of approximation architectures. A central issue in approximate dynamic programming that remains largely open is how to choose the approximation architecture, which remains a problem-specific task, and approximate DP algorithms typically start with the premise that an approximation architecture has been pre-specified. More streamlined methods of approximation architecture selection together with theory characterizing optimality and convergence rates would be desirable. Alternatively, it may be worth trying to explore whether a canonical approximation architecture could be developed at least for some specific class of problems.
- Performance-Oriented Approximate DP. The design, analysis and comparison of value function approximation algorithms is often based on approximation or Bellman errors. Both situations entail the choice of a norm whereby to measure and trade off approximation or Bellman errors across states. Little work has been done to develop

a better understanding of how different criteria capture the quality of the policies being generated. Theoretical analysis is required to identify which criteria — Bellman or approximation error, which norms? — best reflect policy performance; at the same time, algorithms that take such criteria into account must be designed.

- Alternative paradigms and connections with stochastic programming. As a more open-ended direction, one could consider whether extensions of the Markov decision process paradigm, possibly combining it with stochastic programming, and/or a completely different approach could be more suitable for tackling certain classes of real-world problems without running into the curse of dimensionality.

A.4 Haykin, Simon

For my talk, I will do the following:

1. Provide a brief background on the classical Kalman filtering algorithm, emphasizing its virtues and limitations.

2. Describe particle filtering, rooted in Bayesian estimation and Monte Carlo simulation. Here again, I will highlight the virtues and limitations of this second approach.

3. The background would then be set for describing a new approach that has the potential for making a significant difference to the literature. In particular, I will pose a nonlinear recursive problem that is currently the stumbling block. As such, to progress further with this approach, we have to solve this problem.

It could be that the solution I am looking for will pop up in the supplemental technical breakout session following my talk on the second day.

A.5 Kurzhanski, Alexander

Among the central topics of modern control theory are problems of control synthesis for complex systems. These include problems of control under uncertainty and conflict as well as those under complex constraints. The solutions to such problems are well formalizable within Hamiltonian techniques and Dynamic Programming ideas, being as a rule reduced to the solution of HJB and HJBI partial differential equations or/and variational inequalities. An example where such methods are successfully applied are problems of forward and backward reachability for uncertain systems with further passage to control synthesis, safety verification , measurement feedback and related issues. The advent of HJB methods for continuous systems is due to the introduction of appropriate theories for generalized viscosity-type solutions and their equivalents. Another important topic is the solution to problems in dynamics and control for set-valued systems in terms of Hamiltonian formalism and Dynamic Programming.

At present there is a sharp necessity for a breakthrough in numerical methods for solving equations and variational inequalities of the HJB-HJBI types motivated by problems of the above. There is also an interest in numerical methods for the comparison principle in HJB theory, which allows to calculate upper and lower bounds to exact solutions of HJB equations. Another perspective is the calculation of set-valued solutions to problems in evolution dynamics, estimation and control. Together with my colleagues we have developed an ellipsoidal calculus aimed at such problems and closely connected to HJB theory. The calculus is applicable to systems with original linear structure, but its methods allow effective calculations in high dimensions as well as computer animation in high dimensions through computer windows. For nonlinear systems some comparison methods were indicated.

The discussion of such methods as well as those for stochastic dynamics are within my primary interests at the AIM workshop.

A.6 Kushner, Harold

I have been interested in numerical methods in stochastic (and deterministic) control for several decades, and have developed the current methods of choice for the stochastic problem (which is often the best method for determinstic problems as well, as seen in applications to complex variational problems in out book). These cover virtually all of the cost functions and systems of current interest. But, in applications, there is a serious dimensioanlity problem, beyond 4 dimensions.

Various approximation schemes (say, Q-learning or neural nets) have been hyped, but with llittle to show, outside of some very special cases. Additionally, solving the Bellman or HJ-Bellman equation is not enough. One wants the control, and to investigate the effects of implemetation of reasonable approximations.

The visualization problem, even in 3 dimensions, is formidable.

A.7 Mitchell, Ian

The solution to the Hamilton-Jacobi(-Bellman)(-Isaacs) partial differential equation (HJ PDE) provides the optimal value function (eg cost to go) for bounded input optimal control or zero-sum differential game problems respectively. These equations arise in many disparate fields, and my personal interest lies in how they can be used to approximate the reachable sets of continuous dynamic systems in order to verify and/or synthesize provably safe control policies.

For these purposes, it may be possible to relax the goal of a convergent approximation of the solution of the PDE. Instead, we seek guaranteed under or over approximations of the solutions; for example, an underapproximation of the solution might overapproximate the unsafe reachable set, which might thereby generate false negative verifications but no false positives.

With or without such relaxations, is it possible to take advantage of structure in the system dynamics to reduce the typically exponential growth in computational cost with system dimension? For example, is it possible to take advantage of weak or linear coupling between certain components to reduce the degrees of freedom needed by an approximation, perhaps through partial parameterization? Or another approach: Can sampling methods be applied to produce quantifiable probabilistic verifications?

In addition to these questions related to the application area of verification, I am interested in the class of algorithms designed to approximate solutions of HJ PDEs – often broadly called "level set methods." This interest includes both the time-dependent and stationary (or static) versions of the equation, as well as issues related to the use of such numerical approximations in the control of robotic systems.

A.8 Oberman, Adam

I am interested in finite difference methods for high dimensional control and stochastic control problems.

The types of general questions I am interested in are in 2 groups: - we know in theory that we can solve the PDE for the value function. How can we do this in practice for medium dimensional (3-4-5) problems? - Given an approximate numerical solution of the PDE for

the value function, it isn't trivial to implement a near-optimal control strategy. How do we best do this and can we understand how far from optimal these approximations are?

I am curious about how much of a gap there is between theory and practice for nonlinear control problems. Given a control problem, we can write down the PDE for the value function, and we know that this PDE has unique viscosity solutions. From solutions of the PDE, we can then obtain the control.

What I would like to know is the following: Suppose we have an approximate solution of the PDE, and we use this approximation to derive a control. How far from optimal will the resulting solution be?

Also: suppose the optimal solution swtiches strategy infinitely many times. (This happens even in well-known examples, like stabilizing the linear/nonliner pendulum). Even supposing we had the exact solution, how far off are we from optimal if we approximate the control by one which only switches finitely many times, (which is all that we can do in practice).

A.9 Osher, Stanley

I am interested in solving high dimensional Hamilton-Jacobi equations numerically, especially those arising in control theory and involving the level set method.

A.10 Ostrov, Daniel

My primary interest in optimal control in multidimensions has come from applications in financial mathematics. For example, I have recently been working (with Jonathan Goodman) on how to determine the best strategy for trading a number of different stocks so that the expected "utility" of the stocks is maximized. (The term "utility" refers to a specified balance between the desires for a higher return from the stocks and a low total variance.) It is known that in the presence of transactions costs (that is, a charge for buying or selling stocks), there is a "hold region" in the multidimensional space of stock prices, and within this region, one should not trade stocks, but on the boundary of the hold region, one trades so that the stocks do not leave this hold region. It is believed, but not known for certain, that in the asymptotic limit as the transaction costs get small, the hold region becomes a multidimensional parallelogram. However, even assuming that this is true, it is not clear what the orientation of the parallelogram should be.