Fondation Mathématique Jacques Hadamard

Université Paris Saclay

Research initiative on recommendation methods (IRMER)

Call for projects – March 2018

Overview and scientific scope

The Research initiative on recommendation methods (IRMER - *Initiative de Recherche autour des MEthodes de Recommendation*) is a corporate patronage funded by Criteo and operated by the Jacques Hadamard Mathematical Foundation (FMJH - *Fondation Mathématique Jacques Hadamard*).

It is part of the Gaspard Monge Program for Optimization, operational research and their interactions with data science (PGMO - *Programme Gaspard Monge pour l'Optimisation, la recherche opérationnelle et leurs interactions avec la science des données*), launched by EDF and the FMJH, and with Thales, Orange, and Criteo as other associate members. The focus of this IRMER initiative is recommendation methods, both from a theoretical and a practical viewpoints (developing new methods and evaluating their performance on real data).

Mathematicians and computer scientists from both the academic and industrial worlds can benefit from it. Projects are open to all academic researchers with no restrictions due to administrative or geographic location. Nevertheless, small teams with few but committed researchers will be favored.

Projects to be funded should be relevant to the field of data science (including machine learning, statistics, and computer science in relation to data analytics) and should be focused on providing new results concerning recommendation systems. A detailed list of specific problems to be addressed and technologies to consider is provided in appendix.

Objectives

The objective is to support research projects through collaborative actions between academic researchers and industrial researchers or practitioners, focused on improving recommendation systems. These projects are encouraged to be a kick-off for a future partnership between academic and industrial researchers.

Each proposal is thus formed by a pair given by an academic team and a partner company. The academic team must clearly identify a scientific leader, whose lab will manage the funding for the rest of the team.

The partner company must identify a corresponding member and will have to write a support letter describing the industrial challenges to be addressed, the data sets to be studied and the expected benefits of the collaborative research to be undertaken. *The true research (and not only development) nature of the project should be underlined in this letter*.

The partner company does not necessarily need to be Criteo, though Criteo is extremely willing to build partnerships through this sub-program of PGMO.

Call for projects: schedule

- March 2018: publication of this call for projects
- Before final submission: prospective candidates are encouraged to get in touch with the PGMO board to get some pre-submission feedback on the proposal (via Vianney Perchet, see contact details later in this document)
- May 2018: deadline for submission of the projects (link indicated below)
- July 2018: notifications of acceptance or rejection to the project leaders (after recommendations issued by the scientific committee and final decisions made by the executive committee of the PGMO)

Submission of projects: via EasyChair, at the URL
https://easychair.org/conferences/?conf=pgmo2018

Template: a submission template is provided on EasyChair as well at
https://www.fondation-hadamard.fr/fr/pgmo-calls-projects/2018-call-project

(Note that a single PDF file describing the scientific content of the project as well as all required administrative information is expected; it can be written in French or in English.)

Call for projects: rules

What follows is only a summary of the general PGMO submission rules, fully detailed at
https://www.fondation-hadamard.fr/fr/pgmo-calls-projects/2018-call-project

Necessity of an industrial partner; note on data and on the codes

An IRMER project consists of a pair composed by an academic team and an industrial partner.

Projects must emphasize the link with real data. Projects based on public/open data or on the creation of public/open data resembling to industrial or confidential data will be particularly welcome. If the data sets to be studied need to be collected and created first, the project leaders must describe the methodology to be followed and provide a timeline.
Codes are encouraged to be made available publicly.

Funding expectations / Budget rules

Projects duration should typically last 1 year (typically, October 2018 – November 2019).
Funding per project will be typically from 10 to 15 kEuros. We expect to fund about 2 or 3 projects.

All typical research expenses such as travels, computers, internships, invitations of researchers, purchase of data, etc., can be covered. Upon funding, an agreement will be signed between the main lab for the project and FMJH, and the lab will handle the obtained money. This lab or research institution or teaching institution must be from the academic world.

Commitment by funded project teams

PGMO / IRMER being a program of the FMJH, which is part of the Paris Saclay University, all teams of funded projects will be asked to participate to at least one research event in the Saclay area, e.g., at the end of the project: typically, the annual 2-day-long PGMO workshop in Fall 2019.

Support by PGMO / IRMER will have to be acknowledged in publications relative to funded projects.

A follow-up committee composed of representatives of the funding company Criteo may visit the project teams during the 2018-19 year.

Contacts

From the PGMO executive board:

- Vianney Perchet [ENS Paris Saclay], vianney.perchet@normalesup.org

PGMO / IRMER industrial sponsor
(May help to build projects based on the list of suggestions provided in appendix)

- Clément Calauzenes [Criteo], c.calauzenes@criteo.com

Scientific committee: Its composition can be found at the bottom of the page
https://www.fondation-hadamard.fr/fr/pgmo

List of suggested topics and technologies

Several topics are welcome, including but not limited to the following ones (listed with no order of priority). They were suggested by the industrial sponsor Criteo.

Axis 1: Low Dimensional Representations, Embeddings and Feature Selection

In recommendation systems, but also in many other classification problems, data are gathered in very high dimensional spaces. Indeed, one product can be simultaneously described by many different means: images, text files describing some of its characteristics, PDF files such as manuals, and even sometimes videos.

There already exist some black-box methods that can be used to embed these documents in a smaller dimensional space (for instance, recently based on deep learning). Although their performance is individually promising, combining them efficiently is much more challenging.

Similarly, users entering a recommendation system can be described by vectors of (almost) arbitrarily length given the amount of data available to these systems. However, these features are often highly correlated and categorical, which makes their selection more difficult.

The first axis of this call for project focuses on the development of any technique that can allow an **efficient learning based on data of very high dimension**, with an application to recommendation systems. We mentioned deep learning, feature selections, etc., but there are just illustrative. Any other approach is welcome.

Axis 2: Efficient and Fast Optimization and Decision Making

Recommendation systems have several particularities compared to traditional optimization systems. They need to be very reactive and provide answers in a short time frame.

Reactivity is crucial as the decision of buying products by users can be influenced by many external stimuli, the frequency of some of them being rather long-term. For instance, consider selling umbrellas; if the weather of the upcoming weekend is rainy, they will be more sold than barbecues. If the forecast is sunny, it is actually the other way around. As a consequence, the learning (based on more or less recent data) has to be updated frequently -- the more frequent the better.

However, because of the size of the data, the optimization of the learning engine can take up to several hours. It is therefore crucial to find new ways to optimize it sequentially (e.g., based on some current solution) and to reduce this learning time.

Another particularity of recommendation systems is that they have to provide their answer in no time: no one wants to wait for several seconds for a webpage to load. Finding the good compromise between the complexity of a model and its speed is crucial. For instance, it might be easy to build a sophisticated deep network improving marginally the performance of an existing system; yet, if it does not output its decision immediately, this network is useless.

This second axis therefore concerns **fast optimization techniques and algorithms** than can be used to handle to the high frequency constraints of a real-world recommendation system.

Axis 3: Repeated, Real Time Decision Making

Interestingly, recommendation systems are used not only once by a specific user, but several times. This allows a system to learn from its past mistakes to improve itself for the future. But to do that, it needs information on the different possible outcomes that it can gather by exploring its decision space. Online learning theory contains typically the main theoretical tools to that end, and one of the main promising approach to solve this is called the exploration-exploitation dilemma.

However, real-life systems suffer from many drawbacks and particularities that are not covered by the classical assumptions of this theory. Data are cyclical, change slowly over time, are strongly correlated with the first appearance of a user, etc.

Moreover, data are usually gathered after interactions with one or many agents (such as in repeated auctions). Hence, the data generating process is certainly not stochastic, but might come from strategical behaviours.

The third axis focuses on **real life, practical repeated decision making** and how the actual constraints violate the usual assumptions and prevents the use of classical algorithms.

Axis 4: Long Term Interactions, Time Dependencies

In recommendation systems, interactions with other agents, or any user, are repeated over time -- often many times. As a consequence, there are strong time-dependencies in the data. For instance, a user might get increasingly annoyed each time a new ad is displayed, until the failure point where she/he quits the system (with an ad-blocker, by removing cookies, etc.) or buys an item.

Modeling data as a one-shot interaction might be the worst idea possible. The states of the world evolve continuously and are influenced by the auctions of the system (these states being visible or hidden).

Markov Decision Processes [MDP] and their variants (partially observable, hidden, etc., MDPs) could be an interesting way to model this kind of repeated interaction. The associated learning techniques to solve them could range from reinforcement learning, recurrent neural networks, to stochastic games. Obviously, another possibility comes from time series and their possible integration in the above systems.

The fourth proposed axis relies on **taking into account time dependencies to devise fast, efficient algorithms** for the optimization of the system and/or repeated decision making.

List of suggested methods

We also suggest relevant methods; please keep in mind that this list is not exhaustive:

- Stochastic, Parallel and Distributed optimization
- Probabilistic forecasts
- Deep learning, convolution, recurrent networks
- Cryptographic machine learning (supervised, unsupervised, semi-supervised)
- Bandit methods, exploration/exploitation dilemma
- Reinforcement learning
- Hadoop and/or scaling environments for data science
- Typical and atypical patterns; breakpoint detection; weak signals
- Multivariate time series, including Hidden Markov Models (HMM): analysis of causality, graph methods, functional links between components
- “Hybrid” models for time series, e.g. Piecewise Deterministic Markov Models
- Design of experiments
- Visual analytics
- Image statistical models, synthesis, classification
- Text, image, file embeddings
- Model selection in high dimension
- Functional data analysis