

RESEARCH CENTRE

**Inria Saclay Center
at Institut Polytechnique de
Paris**

IN PARTNERSHIP WITH:

Institut Polytechnique de Paris

2022

ACTIVITY REPORT

Project-Team

XPOP

Statistical modelling for life sciences

IN COLLABORATION WITH: Centre de Mathématiques Appliquées
(CMAP)

DOMAIN

Digital Health, Biology and Earth

THEME

Modeling and Control for Life Sciences

Inria

Contents

Project-Team XPOP	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	2
2.1 Developing sound, useful and usable methods	2
2.2 Combining numerical, statistical and stochastic components of a model	2
2.3 Developing future standards	3
3 Research program	3
3.1 Scientific positioning	3
3.2 The mixed-effects models	3
3.3 Computational Statistical Methods	4
3.4 Markov Chain Monte Carlo algorithms	5
3.5 Parameter estimation	5
3.6 Model building	6
3.7 Model evaluation	7
3.8 Missing data	7
4 Application domains	8
4.1 Oncology	8
4.2 Anesthesiology	9
4.3 Modelling Parkinson's Disease	9
4.4 Population pharmacometrics	9
4.5 Mass spectrometry	10
5 Social and environmental responsibility	10
6 New software and platforms	10
6.1 New software	10
6.1.1 mlxR	10
6.1.2 Rsmx	10
6.1.3 SPIX	11
7 New results	11
7.1 Fast Automatic Model Building in Nonlinear Mixed-Effects Models	11
7.2 Modelling the COVID-19 dynamics	11
7.3 Variance reduction for additive functionals of Markov chains via martingale representations	12
7.4 Unifying mirror descent and dual averaging	12
8 Bilateral contracts and grants with industry	12
8.1 Bilateral contracts with industry	12
9 Dissemination	12
9.1 Promoting scientific activities	13
9.1.1 Journal	13
9.1.2 Leadership within the scientific community	13
9.1.3 Scientific expertise	13
9.2 Teaching - Supervision - Juries	13
9.2.1 Teaching	13
9.2.2 Supervision	13
9.3 Popularization	14
9.3.1 Platforms	14
9.3.2 Education	14

10 Scientific production	14
10.1 Major publications	14
10.2 Publications of the year	14

Project-Team XPOP

Creation of the Project-Team: 2017 July 01

Keywords

Computer sciences and digital sciences

- A3.1.1. – Modeling, representation
- A3.2.3. – Inference
- A3.3. – Data and knowledge analysis
 - A3.3.1. – On-line analytical processing
 - A3.3.2. – Data mining
 - A3.3.3. – Big data analysis
- A3.4.1. – Supervised learning
- A3.4.2. – Unsupervised learning
- A3.4.4. – Optimization and learning
- A3.4.5. – Bayesian methods
- A3.4.6. – Neural networks
- A3.4.7. – Kernel methods
- A3.4.8. – Deep learning
- A5.9.2. – Estimation, modeling
- A6.1.1. – Continuous Modeling (PDE, ODE)
- A6.1.2. – Stochastic Modeling
- A6.2.2. – Numerical probability
- A6.2.3. – Probabilistic methods
- A6.2.4. – Statistical methods
- A6.3.3. – Data processing
- A6.3.5. – Uncertainty Quantification

Other research topics and application domains

- B1.1.4. – Genetics and genomics
- B1.1.7. – Bioinformatics
- B1.1.10. – Systems and synthetic biology
- B2.2.3. – Cancer
- B2.2.4. – Infectious diseases, Virology
- B2.3. – Epidemiology
- B2.4.1. – Pharmacokinetics and dynamics
- B9.1.1. – E-learning, MOOC

1 Team members, visitors, external collaborators

Research Scientist

- Marc Lavielle [Team leader, Inria, Senior Researcher, HDR]

Faculty Members

- Erwan Le Pennec [Ecole polytechnique, Professor, HDR]
- Eric Moulines [Ecole polytechnique, Professor, HDR]

PhD Students

- Pablo Jimenez [Ecole polytechnique]
- Achille Thin [Ecole polytechnique]

Administrative Assistant

- Hanadi Dib [Inria]

2 Overall objectives

2.1 Developing sound, useful and usable methods

The main objective of XPOP is to develop new sound and rigorous methods for statistical modeling in the field of biology and life sciences. These methods for modeling include statistical methods of estimation, model diagnostics, model building and model selection as well as methods for numerical models (systems of ordinary and partial differential equations). Historically, the key area where these methods have been used is population pharmacokinetics. However, the framework is currently being extended to sophisticated numerical models in the contexts of viral dynamics, glucose-insulin processes, tumor growth, precision medicine, spectrometry, intracellular processes, etc.

Furthermore, an important aim of XPOP is to transfer the methods developed into software packages so that they can be used in everyday practice.

2.2 Combining numerical, statistical and stochastic components of a model

Mathematical models that characterize complex biological phenomena are defined by systems of ordinary differential equations when dealing with dynamical systems that evolve with respect to time, or by partial differential equations when there is a spatial component in the model. Also, it is sometimes useful to integrate a stochastic aspect into the dynamical system in order to model stochastic intra-individual variability.

In order to use such methods, we must deal with complex numerical difficulties, generally related to resolving the systems of differential equations. Furthermore, to be able to check the quality of a model (i.e. its descriptive and predictive performances), we require data. The statistical aspect of the model is thus critical in how it takes into account different sources of variability and uncertainty, especially when data come from several individuals and we are interested in characterizing the inter-subject variability. Here, the tools of reference are mixed-effects models.

Confronted with such complex modeling problems, one of the goals of XPOP is to show the importance of combining numerical, statistical and stochastic approaches.

2.3 Developing future standards

Linear mixed-effects models have been well-used in statistics for a long time. They are a classical approach, essentially relying on matrix calculations in Gaussian models. Whereas a solid theoretical base has been developed for such models, *nonlinear* mixed-effects models (NLMEM) have received much less attention in the statistics community, even though they have been applied to many domains of interest. It has thus been the users of these models, such as pharmacometricians, who have taken them and developed methods, without really looking to develop a clean theoretical framework or understand the mathematical properties of the methods. This is why a standard estimation method in NLMEM is to linearize the model, and few people have been interested in understanding the properties of estimators obtained in this way.

Statisticians and pharmacometricians frequently realize the need to create bridges between these two communities. We are entirely convinced that this requires the development of new standards for population modeling that can be widely accepted by these various communities. These standards include the language used for encoding a model, the approach for representing a model and the methods for using it:

- **The approach.** Our approach consists in seeing a model as hierarchical, represented by a joint probability distribution. This joint distribution can be decomposed into a product of conditional distributions, each associated with a submodel (model for observations, individual parameters, etc.). Tasks required of the modeler are thus related to these probability distributions.
- **The methods.** Many tests have shown that algorithms implemented in MONOLIX are the most reliable, all the while being extremely fast. In fact, these algorithms are precisely described and published in well known statistical journals. In particular, the SAEM algorithm, used for calculating the maximum likelihood estimation of population parameters, has shown its worth in numerous situations. Its mathematical convergence has also been proven under quite general hypotheses.
- **The language.** Mlxtran is used by MONOLIX and other modeling tools and is today by far the most advanced language for representing models. Initially developed for representing pharmacometric models, its syntax also allows it to easily code dynamical systems defined by a system of ODEs, and statistical models involving continuous, discrete and survival variables. This flexibility is a true advantage both for numerical modelers and statisticians.

3 Research program

3.1 Scientific positioning

"Interfaces" is the defining characteristic of XPOP:

The interface between statistics, probability and numerical methods. Mathematical modelling of complex biological phenomena require to combine numerical, stochastic and statistical approaches. The CMAP is therefore the right place to be for positioning the team at the interface between several mathematical disciplines.

The interface between mathematics and the life sciences. The goal of XPOP is to bring the right answers to the right questions. These answers are mathematical tools (statistics, numerical methods, etc.), whereas the questions come from the life sciences (pharmacology, medicine, biology, etc.). This is why the point of XPOP is not to take part in mathematical projects only, but also pluridisciplinary ones.

The interface between mathematics and software development. The development of new methods is the main activity of XPOP. However, new methods are only useful if they end up being implemented in a software tool. On one hand, a strong partnership with Lixoft (the spin-off company who continue developing MONOLIX) allows us to maintaining this positioning. On the other hand, several members of the team are very active in the R community and develop widely used packages.

3.2 The mixed-effects models

Mixed-effects models are statistical models with both fixed effects and random effects. They are well-adapted to situations where repeated measurements are made on the same individual/statistical unit.

Consider first a single subject i of the population. Let $y_i = (y_{ij}, 1 \leq j \leq n_i)$ be the vector of observations for this subject. The model that describes the observations y_i is assumed to be a parametric probabilistic model: let $p_Y(y_i; \psi_i)$ be the probability distribution of y_i , where ψ_i is a vector of parameters.

In a population framework, the vector of parameters ψ_i is assumed to be drawn from a population distribution $p_\Psi(\psi_i; \theta)$ where θ is a vector of population parameters.

Then, the probabilistic model is the joint probability distribution

$$p(y_i, \psi_i; \theta) = p_Y(y_i | \psi_i) p_\Psi(\psi_i; \theta) \quad (1)$$

To define a model thus consists in defining precisely these two terms.

In most applications, the observed data y_i are continuous longitudinal data. We then assume the following representation for y_i :

$$y_{ij} = f(t_{ij}, \psi_i) + g(t_{ij}, \psi_i) \varepsilon_{ij}, \quad 1 \leq i \leq N, \quad 1 \leq j \leq n_i. \quad (2)$$

Here, y_{ij} is the observation obtained from subject i at time t_{ij} . The residual errors (ε_{ij}) are assumed to be standardized random variables (mean zero and variance 1). The residual error model is represented by function g in model (2).

Function f is usually the solution to a system of ordinary differential equations (pharmacokinetic/pharmacodynamic models, etc.) or a system of partial differential equations (tumor growth, respiratory system, etc.). This component is a fundamental component of the model since it defines the prediction of the observed kinetics for a given set of parameters.

The vector of individual parameters ψ_i is usually function of a vector of population parameters ψ_{pop} , a vector of random effects $\eta_i \sim \mathcal{N}(0, \Omega)$, a vector of individual covariates c_i (weight, age, gender, ...) and some fixed effects β .

The joint model of y and ψ depends then on a vector of parameters $\theta = (\psi_{\text{pop}}, \beta, \Omega)$.

3.3 Computational Statistical Methods

Central to modern statistics is the use of probabilistic models. To relate these models to data requires the ability to calculate the probability of the observed data: the likelihood function, which is central to most statistical methods and provides a principled framework to handle uncertainty.

The emergence of computational statistics as a collection of powerful and general methodologies for carrying out likelihood-based inference made complex models with non-standard data accessible to likelihood, including hierarchical models, models with intricate latent structure, and missing data.

In particular, algorithms previously developed by POPIX for mixed effects models, and today implemented in several software tools (especially MONOLIX) are part of these methods:

- the adaptive Metropolis-Hastings algorithm allows one to sample from the conditional distribution of the individual parameters $p(\psi_i | y_i; c_i, \theta)$,
- the SAEM algorithm is used to maximize the observed likelihood $\mathcal{L}(\theta; y) = p(y; \theta)$,
- Importance Sampling Monte Carlo simulations provide an accurate estimation of the observed log-likelihood $\log(\mathcal{L}(\theta; y))$.

Computational statistics is an area which remains extremely active today. Recently, one can notice that the incentive for further improvements and innovation comes mainly from three broad directions: the high dimensional challenge, the quest for adaptive procedures that can eliminate the cumbersome process of tuning "by hand" the settings of the algorithms and the need for flexible theoretical support, arguably required by all recent developments as well as many of the traditional MCMC algorithms that are widely used in practice.

Working in these three directions is a clear objective for XPOP.

3.4 Markov Chain Monte Carlo algorithms

While these Monte Carlo algorithms have turned into standard tools over the past decade, they still face difficulties in handling less regular problems such as those involved in deriving inference for high-dimensional models. One of the main problems encountered when using MCMC in this challenging settings is that it is difficult to design a Markov chain that efficiently samples the state space of interest.

The Metropolis-adjusted Langevin algorithm (MALA) is a Markov chain Monte Carlo (MCMC) method for obtaining random samples from a probability distribution for which direct sampling is difficult. As the name suggests, MALA uses a combination of two mechanisms to generate the states of a random walk that has the target probability distribution as an invariant measure:

1. new states are proposed using Langevin dynamics, which use evaluations of the gradient of the target probability density function;
2. these proposals are accepted or rejected using the Metropolis-Hastings algorithm, which uses evaluations of the target probability density (but not its gradient).

Informally, the Langevin dynamics drives the random walk towards regions of high probability in the manner of a gradient flow, while the Metropolis-Hastings accept/reject mechanism improves the mixing and convergence properties of this random walk.

Several extensions of MALA have been proposed recently by several authors, including fMALA (fast MALA), AMALA (anisotropic MALA), MMALA (manifold MALA), position-dependent MALA (PMALA), ...

MALA and these extensions have demonstrated to represent very efficient alternative for sampling from high dimensional distributions. We therefore need to adapt these methods to general mixed effects models.

3.5 Parameter estimation

The Stochastic Approximation Expectation Maximization (SAEM) algorithm has shown to be extremely efficient for maximum likelihood estimation in incomplete data models, and particularly in mixed effects models for estimating the population parameters. However, there are several practical situations for which extensions of SAEM are still needed:

High dimensional model: a complex physiological model may have a large number of parameters (in the order of 100). Then several problems arise:

- when most of these parameters are associated with random effects, the MCMC algorithm should be able to sample, for each of the N individuals, parameters from a high dimensional distribution. Efficient MCMC methods for high dimensions are then required.
- Practical identifiability of the model is not ensured with a limited amount of data. In other words, we cannot expect to be able to properly estimate all the parameters of the model, including the fixed effects and the variance-covariance matrix of the random effects. Then, some random effects should be removed, assuming that some parameters do not vary in the population. It may also be necessary to fix the value of some parameters (using values from the literature for instance). The strategy to decide which parameters should be fixed and which random effects should be removed remains totally empirical. XPOP aims to develop a procedure that will help the modeller to take such decisions.

Large number of covariates: the covariate model aims to explain part of the inter-patient variability of some parameters. Classical methods for covariate model building are based on comparisons with respect to some criteria, usually derived from the likelihood (AIC, BIC), or some statistical test (Wald test, LRT, etc.). In other words, the modelling procedure requires two steps: first, all possible models are fitted using some estimation procedure (e.g. the SAEM algorithm) and the likelihood of each model is computed using a numerical integration procedure (e.g. Monte Carlo Importance Sampling); then, a

model selection procedure chooses the "best" covariate model. Such a strategy is only possible with a reduced number of covariates, i.e., with a "small" number of models to fit and compare.

As an alternative, we are thinking about a Bayesian approach which consists of estimating simultaneously the covariate model and the parameters of the model in a single run. An (informative or uninformative) prior is defined for each model by defining a prior probability for each covariate to be included in the model. In other words, we extend the probabilistic model by introducing binary variables that indicate the presence or absence of each covariate in the model. Then, the model selection procedure consists of estimating and maximizing the conditional distribution of this sequence of binary variables. Furthermore, a probability can be associated to any of the possible covariate models.

This conditional distribution can be estimated using an MCMC procedure combined with the SAEM algorithm for estimating the population parameters of the model. In practice, such an approach can only deal with a limited number of covariates since the dimension of the probability space to explore increases exponentially with the number of covariates. Consequently, we would like to have methods able to find a small number of variables (from a large starting set) that influence certain parameters in populations of individuals. That means that, instead of estimating the conditional distribution of all the covariate models as described above, the algorithm should focus on the most likely ones.

Fixed parameters: it is quite frequent that some individual parameters of the model have no random component and are purely fixed effects. Then, the model may not belong to the exponential family anymore and the original version of SAEM cannot be used as it is. Several extensions exist:

- introduce random effects with decreasing variances for these parameters,
- introduce a prior distribution for these fixed effects,
- apply the stochastic approximation directly on the sequence of estimated parameters, instead of the sufficient statistics of the model.

None of these methods always work correctly. Furthermore, what are the pros and cons of these methods is not clear at all. Then, developing a robust methodology for such model is necessary.

Convergence toward the global maximum of the likelihood: convergence of SAEM can strongly depend on this initial guess when the observed likelihood has several local maxima. A kind of simulated annealing version of SAEM was previously developed and implemented in MONOLIX. The method works quite well in most situations but there is no theoretical justification and choosing the settings of this algorithm (i.e. how the temperature decreases during the iterations) remains empirical. A precise analysis of the algorithm could be very useful to better understand why it "works" in practice and how to optimize it.

Convergence diagnostic: Convergence of SAEM was theoretically demonstrated under very general hypothesis. Such result is important but of little interest in practice at the time to use SAEM in a finite amount of time, i.e. in a finite number of iterations. Some qualitative and quantitative criteria should be defined in order to both optimize the settings of the algorithm, detect a poor convergence of SAEM and evaluate the quality of the results in order to avoid using them unwisely.

3.6 Model building

Defining an optimal strategy for model building is far from easy because a model is the assembled product of numerous components that need to be evaluated and perhaps improved: the structural model, residual error model, covariate model, covariance model, etc.

How to proceed so as to obtain the best possible combination of these components? There is no magic recipe but an effort will be made to provide some qualitative and quantitative criteria in order to help the modeller for building his model.

The strategy to take will mainly depend on the time we can dedicate to building the model and the time required for running it. For relatively simple models for which parameter estimation is fast, it is

possible to fit many models and compare them. This can also be done if we have powerful computing facilities available (e.g., a cluster) allowing large numbers of simultaneous runs.

However, if we are working on a standard laptop or desktop computer, model building is a sequential process in which a new model is tested at each step. If the model is complex and requires significant computation time (e.g., when involving systems of ODEs), we are constrained to limit the number of models we can test in a reasonable time period. In this context, it also becomes important to carefully choose the tasks to run at each step.

3.7 Model evaluation

Diagnostic tools are recognized as an essential method for model assessment in the process of model building. Indeed, the modeler needs to confront "his" model with the experimental data before concluding that this model is able to reproduce the data and before using it for any purpose, such as prediction or simulation for instance.

The objective of a diagnostic tool is twofold: first we want to check if the assumptions made on the model are valid or not ; then, if some assumptions are rejected, we want to get some guidance on how to improve the model.

As is the usual case in statistics, it is not because this "final" model has not been rejected that it is necessarily the "true" one. All that we can say is that the experimental data does not allow us to reject it. It is merely one of perhaps many models that cannot be rejected.

Model diagnostic tools are for the most part graphical, i.e., visual; we "see" when something is not right between a chosen model and the data it is hypothesized to describe. These diagnostic plots are usually based on the empirical Bayes estimates (EBEs) of the individual parameters and EBEs of the random effects: scatterplots of individual parameters versus covariates to detect some possible relationship, scatterplots of pairs of random effects to detect some possible correlation between random effects, plot of the empirical distribution of the random effects (boxplot, histogram, . . .) to check if they are normally distributed, . . .

The use of EBEs for diagnostic plots and statistical tests is efficient with rich data, i.e. when a significant amount of information is available in the data for recovering accurately all the individual parameters. On the contrary, tests and plots can be misleading when the estimates of the individual parameters are greatly shrunk.

We propose to develop new approaches for diagnosing mixed effects models in a general context and derive formal and unbiased statistical tests for testing separately each feature of the model.

3.8 Missing data

The ability to easily collect and gather a large amount of data from different sources can be seen as an opportunity to better understand many processes. It has already led to breakthroughs in several application areas. However, due to the wide heterogeneity of measurements and objectives, these large databases often exhibit an extraordinary high number of missing values. Hence, in addition to scientific questions, such data also present some important methodological and technical challenges for data analyst.

Missing values occur for a variety of reasons: machines that fail, survey participants who do not answer certain questions, destroyed or lost data, dead animals, damaged plants, etc. Missing values are problematic since most statistical methods can not be applied directly on a incomplete data. Many progress have been made to properly handle missing values. However, there are still many challenges that need to be addressed in the future, that are crucial for the users.

- State of arts methods often consider the case of continuous or categorical data whereas real data are very often mixed. The idea is to develop a multiple imputation method based on a specific principal component analysis (PCA) for mixed data. Indeed, PCA has been used with success to predict (impute) the missing values. A very appealing property is the ability of the method to handle very large matrices with large amount of missing entries.
- The asymptotic regime underlying modern data is not any more to consider that the sample size increases but that both number of observations and number of variables are very large. In practice

first experiments showed that the coverage properties of confidence areas based on the classical methods to estimate variance with missing values varied widely. The asymptotic method and the bootstrap do well in low-noise setting, but can fail when the noise level gets high or when the number of variables is much greater than the number of rows. On the other hand, the jackknife has good coverage properties for large noisy examples but requires a minimum number of variables to be stable enough.

- Inference with missing values is usually performed under the assumption of "Missing at Random" (MAR) values which means that the probability that a value is missing may depend on the observed data but does not depend on the missing value itself. In real data and in particular in data coming from clinical studies, both "Missing Non at Random" (MNAR) and MAR values occur. Taking into account in a proper way both types of missing values is extremely challenging but is worth investigating since the applications are extremely broad.

It is important to stress that missing data models are part of the general incomplete data models addressed by XPOP. Indeed, models with latent variables (i.e. non observed variables such as random effects in a mixed effects model), models with censored data (e.g. data below some limit of quantification) or models with dropout mechanism (e.g. when a subject in a clinical trial fails to continue in the study) can be seen as missing data models.

4 Application domains

4.1 Oncology

(joint project with ARELIS and COMPO Inria team)

The current methods of diagnosis of subjects presenting a physiological disorder or a pathology are done in vivo by clinical examination and/or imaging, and in vitro by acquisition of biological and molecular data. The diagnosis is ultimately established from these data by following the recommendations of experts by medical specialty. The analysis methodology is therefore long and costly for the healthcare system. There is still a need to be able to easily and inexpensively identify, within a population of subjects, those likely to present a physiological and/or pathological condition of interest in order to conduct further studies on these subjects alone.

Circulating free DNA (cfDNA) is a universal biomarker of disease states; research in the field of liquid biopsy is a major medical and industrial challenge in the field of oncology, Non-Invasive Prenatal Screening (NIPS) and transplant rejection, thanks to the identification of specific genomic signatures, respectively somatic, fetal, or donor. Although our understanding of the mechanisms contributing to cfsDNA production is far from being fully elucidated, these phenomena can be differentiated according to the size of the cfsDNA that will be detected in the circulation. Apoptosis is the essential element in the physiological maintenance of cellular homeostasis that removes unwanted and damaged cells. Apoptotic DNA, is a fragmented double-stranded DNA, with an average size of ~150-200 bp that corresponds to nucleosomal DNA. Whole genome sequencing (WGS) has identified that cfsDNA from plasma in patients with tumor pathology (ctDNA), or DNA circulating in maternal blood from the fetus, as well as donor cfsDNA in the recipient show size shortening compared with a healthy individual. The characterization of the size profile in cfDNA is a very interesting way to distinguish these states from a "normal" situation but current approaches are either too complex and expensive, or insufficiently sensitive for a profiling of the whole population, whatever the stage of the disease, including healthy donors.

ADELIS Technologies (Labège, France), propose an analysis process protected by three patents relating to the exclusive exploitation of μ LAS technology (CNRS-LAAS), in the form of the BiaBooster™ test which allows concentration-by-size characterization of plasma cfDNA, without extraction of nucleic acids even in healthy individuals for whom the concentration is low compared to pathological profiles. This instrument uses an electric field and hydrodynamic countercurrent flow of viscoelastic liquids, in which the occurrence of transverse forces directed towards the walls is shown. These forces increase with the molecular weight of the DNA and thus induce a progressive reduction of the migration speed of the DNA which triggers the size separation of DNA fragments. A fluorescent molecule is used to obtain the electrophoretic profiles via a very sensitive detector.

The objective of this project is to develop a mathematical model to improve the sensitivity of the biomarker.

4.2 Anesthesiology

(joint project with AP-HP Lariboisière and M3DISIM)

Two hundred million general anaesthetics are performed worldwide every year. Low blood pressure during anaesthesia is common and has been identified as a major factor in morbidity and mortality. These events require great reactivity in order to correct them as quickly as possible and impose constraints of reliability and reactivity to monitoring and treatment.

Recently, studies have demonstrated the usefulness of noradrenalin in preventing and treating intraoperative hypotension. The handling of this drug requires great vigilance with regard to the correct dosage. Currently, these drugs are administered manually by the healthcare staff in bolus and/or continuous infusion. This represents a heavy workload and suffers from a great deal of variability in order to find the right dosage for the desired effect on blood pressure.

The objective of this project is to automate the administration of noradrenalin with a closed-loop system that makes it possible to control the treatment in real time to an instantaneous blood pressure measurement.

4.3 Modelling Parkinson's Disease

This project is part of the Predstim protocol, which follows parkinsonian patients who have undergone deep brain stimulation. More precisely, it is an ancillary study, with a loading dose of levodopa per os followed by measurements of the plasma concentration of L-dopa and a metabolite (3-OMD) in 97 parkinsonian patients. At the same time, a clinical evaluation (UPDRS score) of the motor symptoms of Parkinson's disease was performed repeatedly.

The objective of the project is first to develop a pharmacokinetic-pharmacodynamic (PKPD) model using the data collected for the 97 patients in the study. The second objective is to integrate genetic data (in particular mutations of enzymes involved in the metabolism of L-dopa) in order to better explain the observed variability of response to treatment among patients.

4.4 Population pharmacometrics

(joint project with Lixoft)

Pharmacometrics involves the analysis and interpretation of data produced in pre-clinical and clinical trials. Population pharmacokinetics studies the variability in drug exposure for clinically safe and effective doses by focusing on identification of patient characteristics which significantly affect or are highly correlated with this variability. Disease progress modeling uses mathematical models to describe, explain, investigate and predict the changes in disease status as a function of time. A disease progress model incorporates functions describing natural disease progression and drug action.

The model based drug development (MBDD) approach establishes quantitative targets for each development step and optimizes the design of each study to meet the target. Optimizing study design requires simulations, which in turn require models. In order to arrive at a meaningful design, mechanisms need to be understood and correctly represented in the mathematical model. Furthermore, the model has to be predictive for future studies. This requirement precludes all purely empirical modeling; instead, models have to be mechanistic.

In particular, physiologically based pharmacokinetic models attempt to mathematically transcribe anatomical, physiological, physical, and chemical descriptions of phenomena involved in the ADME (Absorption - Distribution - Metabolism - Elimination) processes. A system of ordinary differential equations for the quantity of substance in each compartment involves parameters representing blood flow, pulmonary ventilation rate, organ volume, etc.

The ability to describe variability in pharmacometrics model is essential. The nonlinear mixed-effects modeling approach does this by combining the structural model component (the ODE system) with a statistical model, describing the distribution of the parameters between subjects and within subjects, as well as quantifying the unexplained or residual variability within subjects.

The objective of XPOP is to develop new methods for models defined by a very large ODE system, a large number of parameters and a large number of covariates. Contributions of XPOP in this domain are mainly methodological and there is no privileged therapeutic application at this stage.

However, it is expected that these new methods will be implemented in software tools, including MONOLIX and Rpackages for practical use.

4.5 Mass spectrometry

(joint project with the Molecular Chemistry Laboratory, LCM, of Ecole Polytechnique)

One of the main recent developments in analytical chemistry is the rapid democratization of high-resolution mass spectrometers. These instruments produce extremely complex mass spectra, which can include several hundred thousand ions when analyzing complex samples. The analysis of complex matrices (biological, agri-food, cosmetic, pharmaceutical, environmental, etc.) is precisely one of the major analytical challenges of this new century. Academic and industrial researchers are particularly interested in trying to quickly and effectively establish the chemical consequences of an event on a complex matrix. This may include, for example, searching for pesticide degradation products and metabolites in fruits and vegetables, photoproducts of active ingredients in a cosmetic emulsion exposed to UV rays or chlorination products of biocides in hospital effluents. The main difficulty of this type of analysis is based on the high spatial and temporal variability of the samples, which is in addition to the experimental uncertainties inherent in any measurement and requires a large number of samples and analyses to be carried out and computerized data processing (up to 16 million per mass spectrum).

A collaboration between XPOP and the Molecular Chemistry Laboratory (LCM) of the Ecole Polytechnique began in 2018. Our objective is to develop new methods for the statistical analysis of mass spectrometry data.

These methods are implemented in the SPIX software.

-

5 Social and environmental responsibility

Marc Lavielle is member of the Scientific Committee of *RESPIRE*, a French organization for the improvement of air quality.

6 New software and platforms

6.1 New software

6.1.1 mlxR

Keywords: Simulation, Data visualization, Clinical trial simulator

Functional Description: The models are encoded using the model coding language 'Mlxtran', automatically converted into C++ codes, compiled on the fly and linked to R using the 'Rcpp' package. That allows one to implement very easily complex ODE-based models and complex statistical models, including mixed effects models, for continuous, count, categorical, and time-to-event data.

URL: <http://simulx.webpopix.org/>

Contact: Marc Lavielle

6.1.2 Rsmlx

Name: R speaks Monolix

Keywords: Data modeling, Nonlinear mixed effects models, Statistical modeling

Functional Description: Among other tasks, 'Rsmxlx' provides a powerful tool for automatic PK model building, performs statistical tests for model assessment, bootstrap simulation and likelihood profiling for computing confidence intervals. 'Rsmxlx' also proposes several automatic covariate search methods for mixed effects models.

URL: <http://rsmlx.webpopix.org/>

Contact: Marc Lavielle

Partner: Lixoft

6.1.3 SPIX

Keywords: Data modeling, Mass spectrometry, Chemistry

Functional Description: SPIX allows you to - automatically identify, on the basis of statistical approaches, small but significant differences in spectra measured under different conditions, - model the kinetics of entities that evolve over time

URL: <http://spix.webpopix.org/>

Authors: Marc Lavielle, Yao Xu

Contact: Marc Lavielle

Partner: Laboratoire de Chimie Moléculaire - Ecole Polytechnique

7 New results

7.1 Fast Automatic Model Building in Nonlinear Mixed-Effects Models

Participants: Marc Lavielle.

The SAMBA (Stochastic Approximation for Model Building Algorithm) procedure was developed specifically for the construction of mixed-effects models. We have shown in [4] how this algorithm can be used to speed up this process of model building by identifying at each step how best to improve some of the model components. The principle of this algorithm basically consists in learning something about the best model, even when a poor model is used to fit the data. A comparison study of the SAMBA procedure with SCM and COSSAC show similar performances on several real data examples but with a much-reduced computing time. This algorithm is now implemented in Monolix and in the R package Rsmxlx.

7.2 Modelling the COVID-19 dynamics

Participants: Marc Lavielle.

Short-term forecasting of the COVID-19 pandemic is required to facilitate the planning of COVID-19 healthcare demand in hospitals. We have shown in [3] how daily hospital data can be used to track the evolution of the COVID-19 epidemic in France. A piecewise defined dynamic model allows a very good fit of the available data on hospital admissions, deaths and discharges. The change-points detected correspond to moments when the dynamics of the epidemic changed abruptly. Although the proposed model is relatively simple, it can serve several purposes: It is an analytical tool to better understand what has happened so far by relating observed changes to changes in health policy or the evolution of the

virus. It is also a surveillance tool that can be used effectively to warn of a resurgence of epidemic activity, and finally a short-term forecasting tool if conditions remain unchanged. The model, data and fits are implemented in an interactive web application.

In collaboration with Institut Pasteur (and other groups), we have evaluated in [5] the performance of 12 individual models and 19 predictors to anticipate French COVID-19 related healthcare needs from September 7th 2020 to March 6th 2021. We then built an ensemble model by combining the individual forecasts and tested this model from March 7th to July 6th 2021. We found that inclusion of early predictors (epidemiological, mobility and meteorological predictors) can halve the root mean square error for 14-day ahead forecasts, with epidemiological and mobility predictors contributing the most to the improvement.

7.3 Variance reduction for additive functionals of Markov chains via martingale representations

Participants: Eric Moulines.

We proposed in [1] an efficient variance reduction approach for additive functionals of Markov chains relying on a novel discrete-time martingale representation. Our approach is fully non-asymptotic and does not require the knowledge of the stationary distribution (and even any type of ergodicity) or specific structure of the underlying density. By rigorously analyzing the convergence properties of the proposed algorithm, we show that its cost-to-variance product is indeed smaller than one of the naive algorithms. The numerical performance of the new method is illustrated for the Langevin-type Markov chain Monte Carlo (MCMC) methods.

7.4 Unifying mirror descent and dual averaging

Participants: Eric Moulines.

We introduced and analyzed in [2] a new family of first-order optimization algorithms which generalizes and unifies both mirror descent and dual averaging. Within the framework of this family, we define new algorithms for constrained optimization that combines the advantages of mirror descent and dual averaging. Our preliminary simulation study shows that these new algorithms significantly outperform available methods in some situations.

8 Bilateral contracts and grants with industry

8.1 Bilateral contracts with industry

Participants: Marc Lavielle.

Xpop has a contract with ADELIS, a French analytical instrumentation company. The goal of this collaboration is to develop a method to analyze the physiological size profile of circulating DNA to detect different types of pathological abnormalities.

9 Dissemination

9.1 Promoting scientific activities

9.1.1 Journal

Member of the editorial boards

- Stochastic Processes and their Applications
- Journal of Statistical Planning and Inference
- Electronic Journal of Statistics
- Comptes Rendus de l'Académie des Sciences

Reviewer - reviewing activities Diverse and varied

9.1.2 Leadership within the scientific community

- Eric Moulines is associate researcher of the Alan Turing Institute
- Eric Moulines is scientific director Hi! Paris, Paris Center for Artificial Intelligence for Business and Society
- Eric Moulines is elected member of the French Académie des Sciences.

9.1.3 Scientific expertise

- Marc Lavielle is member of the evaluation committee of the International Center for Mathematics (CIMAT), Guanajuato, Mexico.
- Eric Moulines is member of the award committee of foundation "Charles Defforey".
- Marc Lavielle is member of the research working group Pharmacology of the Institut thématique multi-organismes (ITMO) Technologies pour la Santé.
- Eric Moulines is member of the Evaluation Committee of the Swiss nationale Science Foundations.
- Eric Moulines is member of the Evaluation Committee of IVADO and grant APOGEE, Canada.

9.2 Teaching - Supervision - Juries

9.2.1 Teaching

- Master : Eric Moulines, Regression models, 36, M2
- Engineering School : Eric Moulines, Statistics, 36, 2A, X
- Engineering School : Eric Moulines, Markov Chains, 36, 3A, X
- Engineering School : Erwan Le Pennec, Statistics, 36, 2A, X
- Engineering School : Erwan Le Pennec, Statistical Learning, 36, 3A, X

9.2.2 Supervision

PhD in progress:

- Achille Thin, October 2020
- Pablo Jimenez, October 2020

9.3 Popularization

9.3.1 Platforms

- Marc Lavielle developed and maintains the platform `covidix` for Covid-19 data visualization and modelling.
- Marc Lavielle developed and maintains the platform `respire` for air pollution data visualization.

9.3.2 Education

Marc Lavielle developed and maintains the learning platform `Statistics in Action`. The purpose of this online learning platform is to show how statistics (and biostatistics) may be efficiently used in practice using R. It is specifically geared towards teaching statistical modelling concepts and applications for self-study. Indeed, most of the available teaching material tends to be quite "static" while statistical modelling is very much subject to "learning by doing".

10 Scientific production

10.1 Major publications

- [1] D. Belomestny, E. Moulines and S. Samsonov. 'Variance reduction for additive functionals of Markov chains via martingale representations'. In: *Statistics and Computing* 32.1 (15th Feb. 2022), p. 16. DOI: [10.1007/s11222-021-10073-z](https://hal.inria.fr/hal-03921396). URL: <https://hal.inria.fr/hal-03921396>.
- [2] A. Juditsky, J. Kwon and É. Moulines. 'UNIFYING MIRROR DESCENT AND DUAL AVERAGING'. In: *Mathematical Programming, Series A* (10th June 2022). URL: <https://hal.inria.fr/hal-03921381>.
- [3] M. Lavielle. 'Using hospital data for monitoring the dynamics of COVID 19 in France'. In: *Journal of Data Science, Statistics, and Visualisation* (2022). URL: <https://hal.archives-ouvertes.fr/hal-03321804>.
- [4] M. Prague and M. Lavielle. 'SAMBA: a Novel Method for Fast Automatic Model Building in Nonlinear Mixed-Effects Models'. In: *CPT: Pharmacometrics and Systems Pharmacology* 11.2 (2022). DOI: [10.1002/psp4.12742](https://hal.inria.fr/hal-03410025). URL: <https://hal.inria.fr/hal-03410025>.

10.2 Publications of the year

International journals

- [5] J. Paireau, A. Andronico, N. El Hoz, M. Layan, P. Crepey, A. Roumagnac, M. Lavielle, P.-Y. Boëlle and S. Cauchemez. 'An ensemble model based on early predictors to forecast COVID-19 health care demand in France'. In: *Proceedings of the National Academy of Sciences of the United States of America* 119.18 (3rd May 2022), e2103302119. DOI: [10.1073/pnas.2103302119](https://hal-pasteur.archives-ouvertes.fr/pasteur-03690824). URL: <https://hal-pasteur.archives-ouvertes.fr/pasteur-03690824>.