ATLAS project – The team (B1)
"from Applications to Theory in Learning and Adaptive Statistics"

The team is composed of three main participants, Patricia Reynaud-Bouret, Magalie Fromont, and Gilles Stoltz (devoting 90 % of their research time to the ATLAS project), plus 11 members, whose participation rates range from 15 % to 40 %. The equivalent of the size of the team in terms of full-time positions is 5 of these.

Place and originality. Before desribing how the team was constituted, we present its place in the DMA, the mathematics department of Ecole normale supérieure, Paris. It relies on a core of researchers affiliated there, namely, Patricia Reynaud-Bouret and Gilles Stoltz. Most importantly, if the project is funded, this will correspond to the creation of a statistics (sub-)team. The DMA is a small laboratory–about only 30 permanent positions–whose mathematical range should cover the whole spectrum of (French) mathematics. (This is because Ecole normale supérieure is a small college.) Consequently, none topic may be represented by a large and active team within the DMA. This explains why our team is not physically concentrated in one single laboratory. On the other hand, its autonomy will be very high, since it already gathers all statisticians of the DMA, and there is a tradition of independence and autonomy between the subteams of the DMA. As an example of the emerging autonomy, we mention that we have been organizing our own bi-monthly held seminar (on learning theory) for almost three years. We now need financial support for the ATLAS team to exist–for instance, we need to be able to pay for our internal scientific trips.

The constitution of the team. The project gathers problems and applications of both learning theory and adaptive statistics. Oracle inequalities show that the techniques used and the results stated in both fields may be exactly the same in some cases, see, for instance, Magalie Fromont's paper *Model selection by bootstrap penalization for classification*. She is the perfect link between the interests of Patricia Reynaud-Bouret (more in the sphere of influence of functional adaptive statistics) and those of Gilles Stoltz (mostly interested, until now, in learning theory). They met in University Paris-Sud, Orsay, where they all got their PhD under the supervision or co-supervision of Pascal Massart[1]. They also all used or developed concentration-of-measure inequalities during and after their PhD thesis .

The other participants to the project may be divided into two groups. The first group gathers those whose expertise is needed for the theoretical advances in the fields. Thus, Philippe Marchal and Emmanuel Roy are probabilists, competent, respectively, in the concentration-of-measure phenomenon for infinitely divisible laws and in the ergodic theory of point processes ; Vincent Rivoirard (thresholding techniques), Béatrice Laurent (adaptive tests), and Marie Sauvé (regression function estimation) work in the field of adaptive statistics ; Jérôme Renault and Tristan Tomala are specialists of game theory (and more particulary, of repeated games, possibly with imperfect monitoring).

The second group concerns the four applications we have in mind. Sophie Schbath provides expertise in DNA analysis. Djalil Chafaï has connections with pharmacologists ; in addition, we underline that he defended a PhD on the concentration-of-measure phenomenon, so that he will contribute to the theoretical part of the project as well. Vivien Mallet defended a PhD in the construction and statistical evaluation of a diffusion model for air-quality forecasting. Finally, Christine Tuleau worked two years on problems originating from industry (with a car producer company named Renault), and is now to perform functional classification on data from food processing industry.

In conclusion, the team we formed has a statistical core concentrated on learning theory and adaptive statistics, and also contains researchers covering a large spectrum of applications and of other theoretical fields of interest, like probability and game theory.

[1] Some other participants of the project are also former student of his.

ATLAS project – Scientific content (B2) "from Applications to Theory in Learning and Adaptive Statistics"

The theory of adaptive statistics is at the core of our project since it provides links between concentration of measure, game theory, and applications.
Adaptive statistics consists in designing optimal inference methods without knowing the regularity of the target object. A typical example is the estimation by model selection of the intensity of a Poisson process, where we seek estimators achieving the minimax rate of convergence over Besov balls without knowing the parameter and the radius of the Besov ball where the intensity lies. A closely related question may be formulated within a field of learning theory called individual sequences ; there, the problem of adaptation is due to an on-line constraint. The estimations have to be formed day after day, and we want to perform almost as well as the best of a set of reference predictors (which may only be determined in retrospect). In a game-theoretic formulation, this is just playing a repeated game of prediction.
In all these settings, the main technical problem is that the number of observations is fixed and usually rather small. To obtain the convergence rates, we thus need to bound precisely the gap betwen the statistics and their expectations. Concentration of measure is the perfect tool to do that since it often provides non-asymptotic exponential inequalities for the deviation of these statistics above their means.
In this project, applications are not only viewed as a possible opening but also as a source of interesting theoretical issues. We present in detail below the applications that consist in studying the statistical distribution of words over DNA, testing the homogeneity of the distribution of white cells, classifying functional data, and predicting ozone peaks.

In view of the wide range of mathematical methods of interest of the various applications, we provide an overview of the project as a map (see next page). The existing links are represented in **black**, the links in progress are in **green**, the emerging issues are in **red**.

Scientific context. Mathematical methods at hand.

Before describing formally the purposes of this project, we present briefly the mathematical content of each box of the graphical overview.

Concentration of measure. An interesting aspect of concentration of measure consists in exhibiting exponential inequalities for the deviations of a statistic from its mean. Some of the most efficient concentration results for sums of random variables are Hoeffding–Azuma and Bernstein's inequalities. Other inequalities may be obtained via first, entropy or covariance formulas, and second, Herbst style arguments.

Minimax theory, game theory[2]. In zero-sum games, the target quantities are given by minimax quantities (since the payoff for the first player is the opposite of the second player's payoff, and, as we recalled, both players want to maximize their payoffs). More generally, in any estimation problem, this minimax quantity is a minimax error given by the infimum over all estimation procedures (all, not only the adaptive ones) of their worst case errors. The minimax error (one for each given ensemble of regularity parameters) is a lower bound on the error of any adaptive procedure, and efficient adaptive procedures are those whose average errors match this minimax error, simultaneously within large sets of ensemble parameters (possibly up to small logarithmic factors).

Functional adaptation. The purpose is to infer a function on which very few is known. This problem has been extensively investigated for fifteen years, providing efficient procedures from theoretical and practical

[2] In game theory, payoffs are considered instead of errors ; a payoff is the opposite of an error, and players want to maximize their payoffs whereas statisticians aim at minimizing their errors.

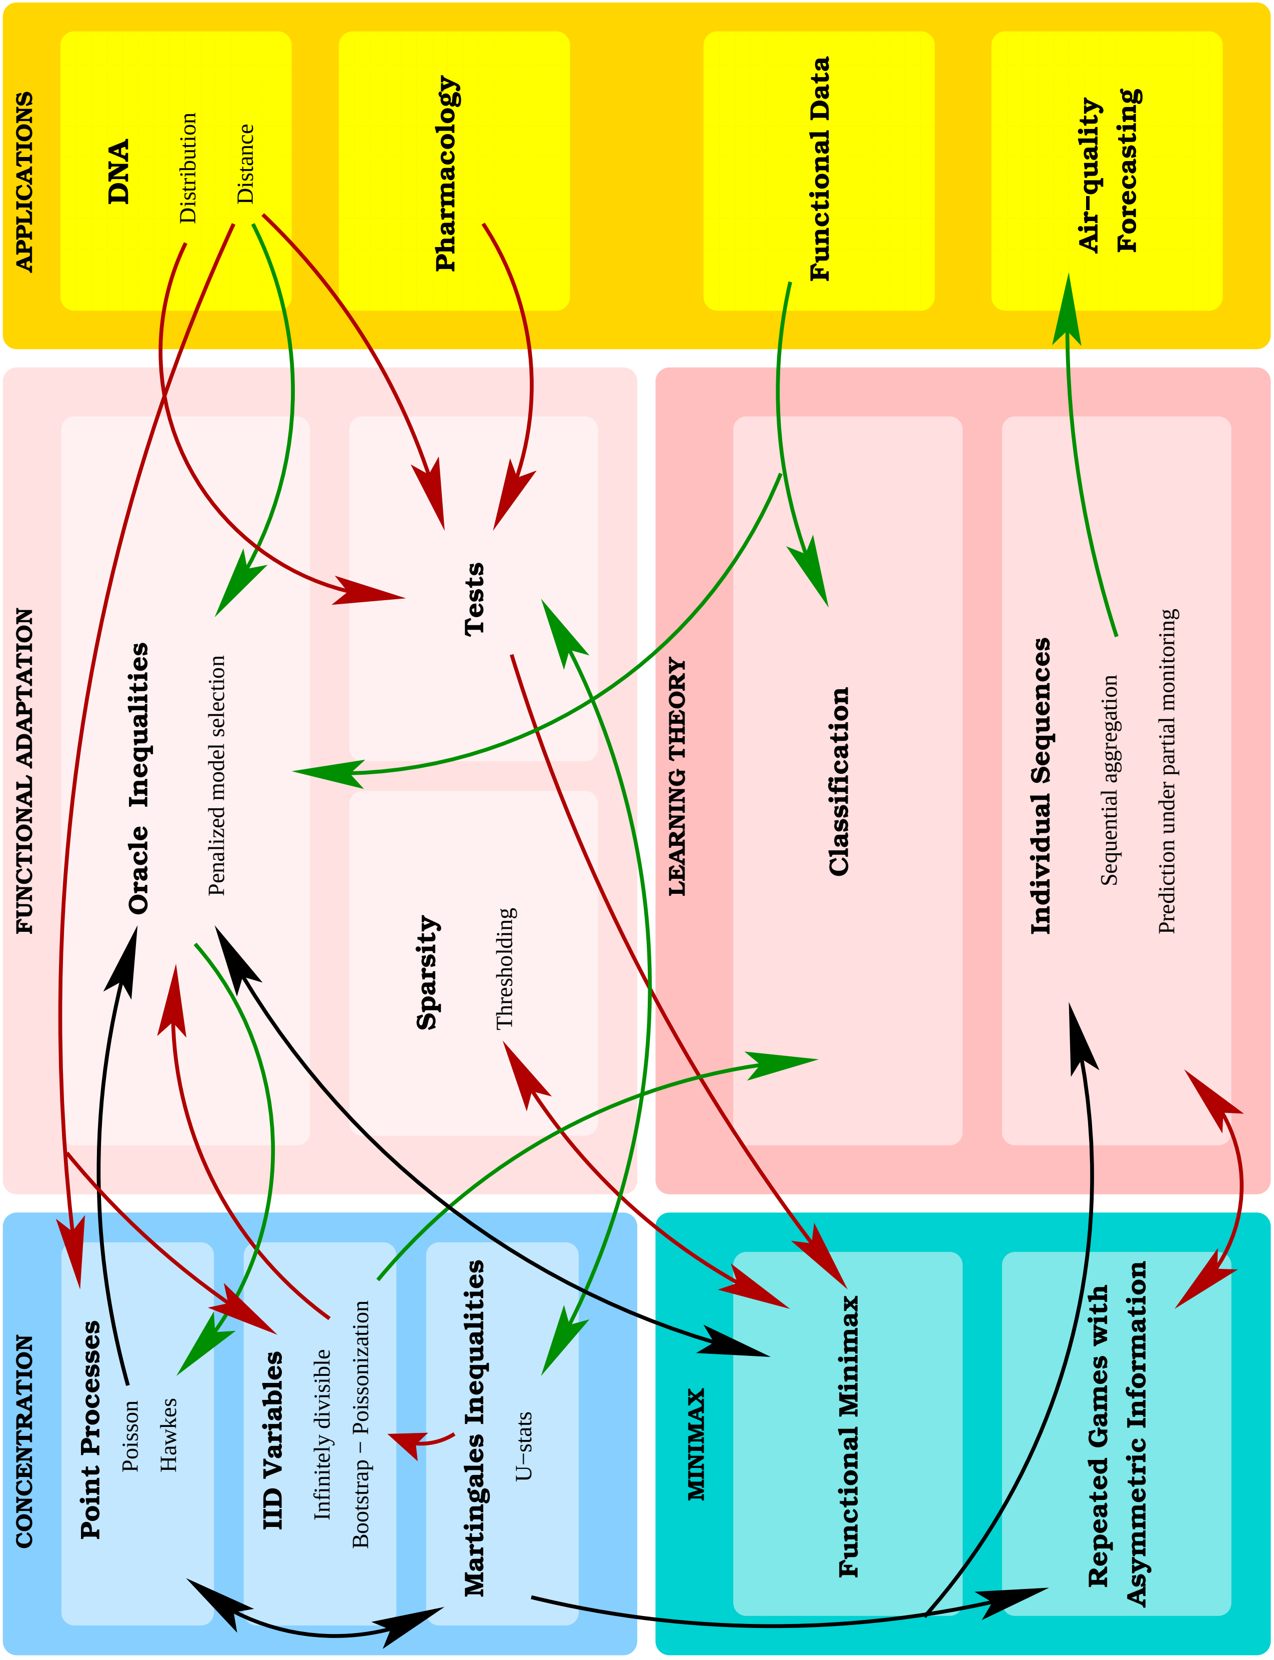

Figure 1: A graphical overview of the ATLAS projet. Existing links in **black**, links in progress in **green**, emerging issues in **red**.

points of view, such as Lepski's method [13], model selection [4], and wavelet thresholding [7]. In particular, wavelet bases that can adapt to the local smoothness at fine scales and to the sparsity of signals are the key ingredients to construct some fast algorithms, which are efficient for data compression, image denoising, or analysis of nonstationary signals. We underline the genuinely multidisciplinary nature of wavelets, since these are very popular tools not also in mathematics, but also in other scientific fields, like geophysics, astrophysics, and biology.

Learning theory. In this projet we consider two subfields of learning theory, classification and prediction of individual sequences.
Classification is concerned with (binary) labelling of observations. A classification rule, that is, a function that associates to each future observation a label, is to be built with the help of n observation–label pairs. Then, its statistical properties are studied. The reference (optimal) predictor is the Bayes classifier, built on the (unobserved) regression function of the data. The excess error between the error of the classifier constructed on the n observations and the Bayes classifier depends on various parameters of the underlying law of the data ; and we want to adapt to them.
Prediction of individual sequences is concerned with on-line learning with the help of a finite number of reference predictors, called experts. At all rounds, they all provide an advice for the next outcome, and based on these advise and on their performances in the past, the statistician has to form its own prediction.

Scientific context. Pursued applications.

We indicate here the applications we have in mind.

Schbath *Distribution of words over* DNA. DNA is represented by a long text written in the alphabet $\{A, C, T, G\}$. To be "read" by the organism, enzymes generally stick to small locations of the DNA (that may be interpreted as words written with the alphabet $\{A, C, T, G\}$). The size of a DNA molecule is typically of the order of $10^6 - 10^9$ bases whereas the typical length of a word is 4 to 9 bases.
A possible modeling for these positions is given by a point process. The case of a Poisson process means that the positions of the words are independent, see [21]. One may then wonder whether there are special locations on the DNA where the points are accumulating.
But it may be more interesting to model interactions between words. Another example of point processes used for DNA (see [10]) is the Hawkes process. The Hawkes processes are used here to model the idea that the biological function of a word w can be efficient only if another word m lies in a certain range around w. This correlation is given for a Hawkes processes through a range function, usually called "reproduction" function. This range function can be positive or negative (in the latter case, this means that the word m must not be present in a given range). The special (and easier) case is given by $m = w$ and a non-negative range function.

Chafaï *Pharmacology.*
Counting the different white blood cells (leukocytes) is a commonly used biological test. Leukocytes play an important role in the immune system, and are of different types. Several techniques are used by the biologists to perform the counting. The advantage of visual counting by a medical technologist is that blood cells misidentified by an automated counter can be identified visually. To identify the numbers of the different leukocytes, a blood film is made, and a certain number (usually 100) of leukocytes are counted by the examination of some part of the film. This yields an approximation of the distribution among leukocytes types. This heuristical method is not satisfactory, and is not mathematically driven. However, it turns out that the spatial distribution of the leukocytes on the film can be modelled by a spatial point process. We believe that the recent development of adaptative estimation and tests techniques for spatial point processes can help to improve the performance and the practice of leukocyte countings, and make this statement precise in the next part of this project.

Fromont, Tuleau *Functional data.* In binary classification, the case when the observations take their values in $\mathbb{R}^d$ has been extensively studied. But in many real-world applications coming from food processing industry or speech recognition for example, the observations to be classified are more accurately represented by discretized functions than by standard vectors. A simple way of formalizing this framework is to assume that the space $\mathcal{X}$ in which the instance variable X takes its values is a functional space.

Functional data → Classification

The most simple and popular classifiers are probably the kernel and k-nearest neighbor (k-NN) rules. They have been investigated in such functional space cases for a few years now (see, for instance, Abraham, Biau, and Cadre, 2003, or Cerou and Guyader, 2005). However, such direct approaches suffer from a drawback commonly referred to as the curse of dimensionality. To overcome it, most of the traditional effective methods for $\mathbb{R}^d$-valued data analysis have been adapted to handle functional data (the "functional data analysis" methods). In a recent paper, Biau *et al.* [3] thus apply the usual k-NN rule to the coordinates of the projections of the data on a suitably chosen vector space of finite dimension d. The choice of both the dimension d and number of neighbors k is made automatically by a minimization of a penalized empirical classification error performed after some data-splitting device. The resulting classifier satisfies an oracle-type inequality.

Two main issues remain however unsolved. From a theoretical point of view, it is not clear why one should prefer implementing the procedure based on the minimization of the empirical classification error without any penalization. Furthermore, Biau *et al.* [3] point out that the data-splitting device is unstable, and suggest therefore to consider many random splits of the data and to combine the corresponding classifiers. Whereas this combination technique is commonly used in practice, it has no theoretical foundation in the present classification context.

In a work in progress, Fromont and Tuleau address both issues, and note that the penalty of order $n^{-1/2}$ proposed by Biau *et al.* is actually too heavy when some margin-type assumptions are satisfied. This suggests to take a penalty either equal to zero or of order smaller than $n^{-1/2}$. By using a recent result of Boucheron, Bousquet, and Massart, they have already obtained a new oracle inequality that justifies such non-penalized or slightly penalized versions of the procedure. They have also shown via an experimental study based on speech recognition or food processing industry data that the introduction of such a small-order penalization stabilizes the data-splitting device while preserving good performances.

Classification → Oracle inequalities

This work now requires a deeper theoretical analysis to calibrate precisely the order of the (slight) penalties to use, and hence to prove (from a theoretical point of view) the relevance of the introduction of these penalties as a stabilization tool. This analysis involves a careful study of the minimax risk and of the rate of convergence of the k-NN rule under some margin-type assumptions in the present functional context. Some significant results in this direction have been obtained recently by Audibert and Tsybakov, and Györfi in finite dimension, and the purpose here would be to explore these paths in a functional data framework.

Mallet, Stoltz

Air-quality forecasting (prediction of ozone peaks). Very few is known with respect to the practical behavior of estimation procedures provided by prediction of individual sequences. These are robust procedures that do not rely on any stochastic assumption on the observations and that aggregate sequentially the predictions of a finite set of given experts. They are guaranteed to perform as well as the best of the experts, even if the latter can only be determined in hindsight.

Individual sequences → Air-quality forecasting

In air-quality forecasting, the experts will be different prediction models and we will use individual sequences to combine them in a smart way. Every model is defined by its physical and chemical assumptions, by its numerical schemes and by its forcing fields.

A very preliminary empirical study is provided by Mallet and Sportisse [14], for one particular algorithm of individual sequences, in the case where the range of the ozone peaks and the length of the prediction period are known. First, a totally adaptive version of such on-line aggregating algorithms will have to be designed.

Second, as they noted, there are missing data (at each prediction step, several monitoring stations are unavailable), which prevents from evaluating the performance of the algorithm at some time steps in some parts of the field of study. In addition, since it is time-consuming to compute the prediction of all models, we may wish not to evaluate all of them, but only the best ones plus a (rather small) random fraction of the other ones. This is a setting of partial monitoring of the prediction, which we need first to study theoretically (see the section on game theory below).

Constructing new paths on the graphical overview.

In this section we present in some more detail most of the **black**, **green**, and **red** arrows of the map (recall that they respectively correspond to completed work, work in progress, and new research issues).

Chafaï, Fromont, Laurent, Reynaud, Rivoirard

Poisson processes. In this setting, we observe a Poisson process on a space $\mathbb{X}$, where $\mathbb{X}$ is typically $[0,1]$, $\mathbb{R}$, $[0,1]^2$ or $\mathbb{R}^2$. We assume that the Poisson process has an intensity s with respect to some measure μ (typically the Lebesgue measure), and we want either to estimate s or to test the homogeneity of the process (i.e., whether s is constant or not).

Point processes → Oracle inequalities

Oracle inequalities ↔ Functional minimax

Intensity estimation (easy part, done). This preliminary work was a part of Patricia Reynaud-Bouret's PhD thesis. She obtained in [15] exponential inequalities for suprema of integrals with respect to the centered Poisson process. These inequalities were the key tool to prove oracle inequalities for the penalized projection estimators (a particular case of model selection). These estimators [15] are adaptive in the minimax sense for Besov balls, as well as for some sets of the type "elements having D non-zero coefficients among N" (she computed the minimax risks for the latter).

DNA, Pharmacology → Tests

Homogeneity tests (work in progress). Adaptive tests are required for applications. We have been asked twice already the same question, on testing the homogeneity for Poisson processes–the first time, in order to test the homogeneity of the distribution of the words along DNA for a composed Poisson process [?], and the second time, to test the homogeneity of the distribution of white cells. On the one hand, an easy answer could be that such tests already exist (at least for classical Poisson process), as variants of Kolmogorov–Smirnov tests, for instance. On the other hand, the crucial question is the kind of alternatives one wants to separate from the null hypothesis. In both applications we consider in this projet, the alternatives are not smooth, but rather of "cluster" type (are there one or more spots where the points seem to aggregate themselves?). The corresponding intensities are truly irregular.

Tests ↔ Martingales inequalities

Magalie Fromont and Béatrice Laurent [9] have recently built adaptive tests clearly outperforming Kolmogorov-Smirnov tests for those kind of alternatives, in the setting of testing a density from a n-sample. We aim at reproducing this result in a Poisson setting. The main concentration tool in [9] is given by exponential inequalities for U-statistics ([11], see also Giné *et al.*, 2000), and it turns out that [11] also provides the material needed to get adaptive tests in the Poisson framework–except one thing. These martingale techniques only deal, in a first approach, with Poisson processes in one dimension spaces. Since the white cells applications need Poisson processes in two dimensions, we will have to extend all these results to d–dimensional settings (at least, for $d = 2$).

Tests → Functional minimax

Another issue is to get minimax bounds for functional classes which reflect the complexity of having "cluster" alternatives. In the same time, we need to consider a more complex procedure incorporating a priori weights on the models. This method has been used for estimation but not for tests as far as we know.

Sparsity ↔ Functional minimax

Sparsity (intensity estimation, second part – new research issue). The problem above, as well as the work on Hawkes processes, see below, reveal that the classical functional classes, on which the minimax risk is computed and to which the procedures are adaptive, are really too smooth for the encountered practical problems. The target functions that the applications require are of the form a "very smooth" function (null, constant, etc.) plus a very small number of highly localized disruptions. The complexity of these target functions lies in the fact that we do not know the number and the locations of the disruptions. These functions are referred to as sparse since most of their coefficients are null once written on the proper wavelet basis. Unlike classical Besov spaces, the weak Besov spaces include those targets [20]. We are currently concentrating on building threshold estimators in the Poisson framework that do not depend on the (possibly infinite) support, and hope that the procedure will be adaptive to weak Besov balls.

Reynaud, Roy, Schbath

Hawkes processes.

Case of a single word, i.e., $m = w$ (work in progress).

DNA → Oracle inequalities

Non-parametric estimators of the range functions are known, some of them may be obtained by spline methods, see [10]. To choose the number of knots, Gaëlle Gusto and Sophie Schbath use Akaike's criterion, but the latter fails as soon as too many estimators with the same number of knots are in competition. On the other hand, the shape of the range functions that the biologists have the intuition of has a special form.

They expect them to be non-zero on very small intervals centered on the required range of interactions. Even in the easy case of piecewise constant estimators, the number of estimators required to catch the good one is large–and Akaike's criterion is just useless. On the other hand, functional adaptation, and more precisely penalized model selection, may give a way to select the right estimator in this case.

But as we have already seen it for Poisson processes, model selection requires corresponding concentration inequalities. Here, two types of concentration are needed. The first comes from martingales inequalities [16] and was already developed for the penalized projection estimators of Aalen multiplicative intensity [17]. The second one take its roots in ergodic theory. For the statistical problem, the ergodic theorem ensures that the random norm we need to work with converges toward a deterministic norm depending of the parameters of the problem. However, as we already mentioned, asymptotic results prevent from adaptive estimation, and thus, exponential inequalities specifying this convergence are needed. This is the main subject of [18]. Several limitations occur, for instance, we are only able to deal with non-negative range functions. Moreover, the resulting penalties depend on unknown parameters. However, this theoretical study indicates the shape of the penalty. A possible solution is then to use heuristic results on the slope [12], that Birgé and Massart [4] proved theoretically valid in the Gaussian framework. We are currently not able to confirm theoretically that the slope method is reasonable, and focus in a first time on an extensive study by simulations, some of them lasting over a month on our computers. Later, we would need heavier concentration results to justify the slope heuristic in this context.

Oracle inequalities → Point processes, Martingales inequalities

DNA → Concentration

The case of two words (new research issue). We have no idea as well how to find concentration inequalities for the random norms when two words are involved (even for a nonnegative range function). Our ultimate goal is to test, given two words w and m, whether their range function is null or not, i.e., whether there is a biological link or not. Since the alternative hypotheses will be quite irregular, we will need adaptive tests procedures, and consequently also the concentration inequalities for corresponding U-statistics–which is completely out of reach right now.

DNA → Tests

Infinitely divisible variables. This work was initiated by the observation that Hawkes processes are infinitely divisible. So, perhaps concentration inequalities for norms of infinitely divisible i.i.d. random vectors would lead to proper exponential inequalities. A first study was not sharp enough to solve the problems raised by the Hawkes processes. However it improved already in certain cases the concentration formulas given in the Poisson setting [15]. Christian Houdré, Philippe Marchal, and Patricia Reynaud-Bouret are currently improving these concentration formulas, and are already on the track of an exponential inequality for the norm of a projection of an infinitely divisible i.i.d. random vector having an exponential moment.

It turns out that this is exactly the quantity appearing in model selection in a regression framework. Marie Sauvé has already built penalized projection estimators in regression with i.i.d. errors having an exponential moment, and this framework seems to be more general. But in practice, the lack of general concentration formulas prevents her, for instance, to use wavelet basis. If the study is restricted to infinitely divisible errors, it is likely to give a broader cast of estimators.

Marchal, Reynaud, Sauvé, (Houdré)

DNA → Infinitely divisible variables

Infinitely divisible variables → Oracle inequalities

Bootstrap – Poissonization. We consider again the setting of binary classification, and assume that we have access to a set S of candidate classifiers. Vapnik and Chervonenkis (1971) suggested to pick the classifier ϕ_n with smallest classification error on the data (recall that we observe a n-sample of observation–label pairs). The next issue is how to choose S. Vapnik (1982) introduced structural risk minimization, which consists in considering a collection of such sets S, associating to each of these a penalty term depending on its complexity, and choosing the one with smallest sum of best empirical error plus penalty term. A popular measure of the complexity of the set is provided by the Vapnik-Chervonenkis dimension. Unfortunately, it does not depend on the distribution of the data, only on the sets themselves– and hence may overestimate the complexity for specific distributions. Koltchinskii (2001) and Bartlett, Boucheron, and Lugosi [1] thus introduce the so-called Rademacher penalties, which are completely determined from the data. They prove oracle type inequalities showing that such random penalties provide optimal classifiers in a global minimax sense. Whereas these Rademacher penalization techniques are now of common use in statistical learning theory, they are not so popular yet in the applied statistics community. Actually, statisticians often stick to resampling tools such as bootstrap or jacknife, in practice.

Fromont

Bootstrap - Poissonization → Classification

The global minimax point of view (done). Magalie Fromont [8] connects the two approaches by introducing a new family of penalties based on classical bootstrap processes such as Efron's or i.i.d. weighted bootstrap ones. She derives non-asymptotic properties for the corresponding estimators from concentration inequalities (exponential inequalities on the deviations of these bootstrap penalties). In particular, by using some Poissonization arguments, she proves that these estimators achieve the global minimax risk over sets of functions built from Vapnik-Chervonenkis classes. The obtained results generalize those of Koltchinskii (2001) and of Bartlett, Boucheron, and Lugosi [1], in the case of Rademacher penalties, which can thus be seen as special cases of bootstrap-type penalties.

Localized versions of bootstrap penalties (under construction). Rademacher penalties are also used in the statistical learning theory not only to construct classifiers, but also to provide sharp error bounds. Following the ideas initially introduced by Koltchinskii and Panchenko (1999), Bartlett, Bousquet and Mendelson [2] and Bartlett, Mendelson and Philips (2004) propose some localized versions of Rademacher averages as tight data-dependent measures of complexity. In the same spirit, we may introduce localized versions of bootstrap penalties to obtain sharp error bounds. This requires improving the exponential inequalities obtained by Magalie Fromont [8], for instance, under margin assumptions and by using a combination of the concentration inequality by Boucheron, Lugosi, and Massart (2000) with some refined Poissonization techniques. In parallel, new concentration inequalities, based on some martingale techniques and that could directly be applied to the bootstrap quantities, are expected.

Bootstrap – Poissonization ← Martingale inequalities

Bootstrap – Poissonization → Oracle inequalities

The localized Rademacher averages have been proved to provide estimators that achieve the minimax risk under margin assumptions (those introduced in Tsybakov [22], and in Massart and Nédélec, 2005). The above exponential inequalities for localized versions of bootstrap penalties would probably lead straight to sharp oracle inequalities, under these margin assumptions. Thus classifiers based on bootstrap samples of the data and that are adaptive to the margin could be obtained.

Renault, Stoltz, Tomala

Prediction of individual sequences under partial monitoring. An equivalent name for this problem is playing a repeated game with asymmetric information (given by feedback signals for the first player and perfect monitoring of the actions for the second player). Recall that a special case was provided by air-quality forecasting problems, and we now address the general issue of prediction under partial monitoring. Formally, given a finite class of actions–corresponding to constant experts–, the first player aims at playing in a way such that, in retrospect, he has obtained an average payoff close to a certain minimax quantity. When both players, after taking their actions, are allowed to see the other player's action (perfect monitoring), then this minimax quantity is the payoff obtained by the best constant action for the given sequence of the second player's actions. (It might not be obvious yet why this is a minimax quantity. This will be explained below.) Strategies achieving this minimax quantity are called Hannan-consistent, and there exist many of them in settings of perfect monitoring, see Blackwell (1956), Hannan (1957) for early references.

Individual sequences ↔ Repeated games with asymmetric information

Hannan-consistency under partial monitoring (done). They also exist in some situations of partial monitoring, where the first player only gets to see a deterministic feedback depending on the pair of actions chosen by both players, whereas the second player still has perfect monitoring of the game. Of particular interest is the case where the feedback is the obtained payoff– this is the bandit setting. In this particular case, and more generally, in all cases where (in some sense made precise by Piccolboni and Schindelhauer) the payoffs may be reconstructed from the feedbacks, Hannan-consistency may still be achieved, at a rate at least $n^{-1/3}$, where n is the number of game rounds, see Cesa-Bianchi, Lugosi, and Stoltz [5]. This is still the case (under a natural condition of reconstruction) when the feedbacks are given by signals drawn at random according to a probability distribution depending on the actions profiles.

Martingale inequalities → Individual sequences

The general minimax quantity, under partial monitoring (new research issue). Now, Rusticchini (1999) indicates the minimax quantity in the general case of random signals depending in an arbitrary way on the actions profiles. In good cases, of course, this is the payoff of the best expert. In general, it equals the minimum (over all sequences of the second player's actions that yield on average the same signals as the played sequence) of the maximum (over all mixed actions of the first player) of the corresponding average payoff. He also shows, by means of a generalization of Blackwell's approachability theorem (provided by Mertens, Sorin, and Zamir, 1994) that this minimax quantity is achievable. (See similar techniques in Renault and Tomala [19].) This, in principle, yields a strategy for the first player ; it is however not easy to exhibit it from the proof. Rates of convergence towards the minimax quantity are not underlined by Rusticchini. Our intuition is that these should be of the order of $n^{-1/3}$ as well. In addition to that, we aim at exhibiting an explicit strategy for the first player, comparable to those of Cesa-Bianchi, Lugosi, and Stoltz [5] and Mannor and Shimkin (2003). The latter indicate a strategy in the particular case when the feedback is independant of the second player's actions, and a closer look at their proof yields an $n^{-1/3}$ rate. A combination of the proofs of both papers, in addition to refined martingale and concentration inequalities is our first lead. The refined inequalities might be of the same flavour as those used in Cesa-Bianchi, Mansour, and Stoltz [6], where the rates of convergence toward the payoff of the best expert in perfect monitoring, bandit setting, and general partial monitoring, are obtained via a statistical approach in terms of conditional variances of some estimators of the payoffs.

Martingale inequalities → Individual sequences

References

[1] P. Bartlett, S. Boucheron, and G. Lugosi. Model selection and error estimation. *Machine Learning*, 48:85–113, 2002.

[2] P. Bartlett, O. Bousquet, and S. Mendelson. Localized Rademacher complexities. *Annals of Statistics*, 33(4):1497–1537, 2005.

[3] G. Biau, F. Bunea, and M. Wegkamp. Functional classification in Hilbert spaces. *IEEE Transactions on Information Theory*, 51(6):2163–2172, 2005.

[4] L. Birgé and P. Massart. Generalized Cp criterion for gaussian model selection. Technical Report 647, Université Paris VI, Avril 2001.

[5] N. Cesa-Bianchi, L. Lugosi, and G. Stoltz. Regret minimization under partial monitoring. Technical report, Ecole normale supérieure.

[6] N. Cesa-Bianchi, Y. Mansour, and G. Stoltz. Improved second-order inequalities for prediction under expert advice. *Proceedings of COLT'05*, 217–232, 2005.

[7] D. L. Donoho, I. M. Johnstone, G. Kerkyacharian, and D. Picard. Wavelet shrinkage: asymptopia ? *Journal of the Royal Statistical Society*, B–57(2):301–369, 1995.

[8] M. Fromont. Model selection by bootstrap penalization for classification. *Machine Learning*, 2006. To appear.

[9] M. Fromont and B. Laurent. Adaptive goodness-of-fit tests in a density model. *The Annals of Statistics*, 2005. To appear.

[10] G. Gusto and S. Schbath. F.A.D.O.: a statistical method to detect favored or avoided distances between occurrences of motifs using the Hawkes model. Preprint, submitted, 2005.

[11] C. Houdré and P. Reynaud-Bouret. Exponential inequalities, with constants, for U-statistics of order two. In *Stochastic inequalities and applications, Progress in Probability*, 56:55–69, 2003.

[12] E. Lebarbier. *Quelques approches pour la détection de ruptures à horizon fini*. PhD thesis, Université Paris XI, 2002.

[13] O. V. Lepski. Asymptotically minimax adaptive estimation (I). Upper bounds. Optimally adaptive estimates. *Teor. Veroyatnost. i Primenen.*, 36(4):645–659, 1991.

[14] V. Mallet and B. Sportisse. Toward ensemble-based air-quality forecasts. In revision for *Journal of Geophysical Research*.

[15] P. Reynaud-Bouret. Adaptive estimation of the intensity of inhomogeneous Poisson processes via concentration inequalities. Probability Theory and Related Fields, 126(1):103–153, 2003.

[16] P. Reynaud-Bouret. Compensator and exponential inequalities for some suprema of counting processes. *Statistics and Probability Letters*, 2006. To appear.

[17] P. Reynaud-Bouret. Penalized projection estimators of the Aalen multiplicative intensity. *Bernoulli*, 2006. To appear.

[18] P. Reynaud-Bouret and E. Roy. Some non asymptotic tail estimates for Hawkes processes. *Bulletin of Belgian Mathematical Society*, special edition on BeNeLuxFra Conference in Gent, 2005. To appear.

[19] J. Renault et T. Tomala. *Communication equilibrium payoffs in repeated games with complete information and imperfect monitoring*. Games and Economic Behavior, 49, 124–156, 2004.

[20] V. Rivoirard. Non linear estimation over weak Besov spaces and minimax Bayes method. *Bernoulli*, 2006. To appear.

[21] S. Robin, F. Rodolphe, and S. Schbath. *DNA, Words and Models*. Cambridge University Press, 2005.

[22] A. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.

ATLAS project – Organization chart of the laboratory (B3-1) "from Applications to Theory in Learning and Adaptive Statistics"