

# Postdoctoral Research Proposal

## Online Estimation of Quantiles and Superquantiles via Variance-Reduced Stochastic Algorithms

### Supervision

#### Supervisors (alphabetical order):

- Bertrand Iooss and Jérôme Lonchamp (EDF R&D)
- Manon Costa (IMT)
- Sébastien Gadat (TSE)
- Thierry Klein (ENAC & IMT)

The project will be carried out under the supervision of the above-mentioned researchers, bringing together complementary expertise in stochastic approximation, risk measures, and applied industrial modeling.

## 1 Research Context and Motivation

Quantile-based risk measures, such as Value-at-Risk (VaR), and their coherent counterparts, superquantiles (Conditional Value-at-Risk, CVaR), are fundamental tools in risk management, finance, operations research, and machine learning. In modern applications, data is often generated sequentially or through complex simulations, which calls for *online* and *computationally efficient* estimation procedures.

Stochastic approximation methods provide a natural framework for such problems. However, their practical performance is often limited by the high variance of stochastic gradients, especially in tail-related estimation problems. Recent advances in stochastic optimization, particularly variance reduction techniques such as SPIDER, open new perspectives for improving the efficiency of these methods.

This project aims to bridge these areas by developing variance-reduced stochastic algorithms for online estimation of quantiles and superquantiles, with strong theoretical guarantees and practical efficiency.

## 2 Background: Quantiles and Superquantiles

Let  $X$  be a real-valued random variable with cumulative distribution function  $F$ . For  $\alpha \in (0, 1)$ , the  $\alpha$ -quantile is defined by

$$q_\alpha = \inf\{x \in \mathbb{R} : F(x) \geq \alpha\}.$$

Quantiles are widely used but suffer from important limitations, including lack of convexity and high sensitivity to estimation noise, which can lead to unstable optimization procedures. In contrast, the superquantile (Conditional Value-at-Risk, CVaR) provides a more robust and

coherent measure of risk, as it satisfies the axioms of coherent risk measures (in particular convexity and subadditivity) [13, 5]. The superquantile is defined as

$$\text{CVaR}_\alpha(X) = \mathbb{E}[X \mid X \geq q_\alpha],$$

and admits the following variational representation, which plays a central role in practice as it enables the formulation of superquantile estimation as a stochastic optimization problem:

$$\text{CVaR}_\alpha(X) = \min_{t \in \mathbb{R}} \left\{ t + \frac{1}{1-\alpha} \mathbb{E}[(X-t)_+] \right\}.$$

Unlike quantiles, CVaR is a coherent risk measure and enjoys convexity properties that facilitate optimization.

## 3 State of the Art

### 3.1 Stochastic Approximation for Quantile Estimation

The estimation of quantiles via stochastic approximation dates back to the seminal work of Robbins and Monro [1]. Extensions and refinements have been proposed in [2, 3, 4]. These methods provide recursive estimators but often suffer from slow convergence and sensitivity to noise.

The variational formulation of CVaR introduced in [5] has enabled the use of stochastic gradient methods for its estimation. Applications in risk-aware optimization and reinforcement learning have been developed in [6, 7].

Recent works have extended stochastic approximation techniques to the estimation of superquantiles. In particular, Bercu, Gadat, and Costa have proposed recursive procedures for the joint estimation of quantiles and superquantiles, combining Robbins–Monro type updates with additional averaging steps to capture tail expectations [14]. These approaches provide almost sure convergence results as well as asymptotic normality under suitable assumptions. A key contribution of this line of work is the analysis of the coupled dynamics between the quantile estimator and the superquantile estimator, which introduces additional challenges compared to standard stochastic approximation schemes. Moreover, these methods highlight the importance of carefully tuning step sizes and controlling the propagation of noise in the tail estimation process. In addition, a non-asymptotic analysis of recursive superquantile estimation has been developed in [15], providing finite-sample deviation bounds and shedding light on the convergence rates of these algorithms.

However, despite their strong theoretical guarantees, these methods may suffer from slow convergence in practice due to the intrinsic variance of the estimators, especially in rare-event regimes. This limitation motivates the investigation of variance reduction techniques, such as SPIDER, within this framework.

### 3.2 Variance Reduction Techniques

Variance reduction has become a central topic in stochastic optimization. Classical approaches include control variates and importance sampling [8]. More recent methods such as SVRG [9], SAGA [10], and SARAH [11] achieve improved convergence rates. SPIDER [12] introduces a recursive gradient estimator with near-optimal complexity for non-convex optimization. Its efficiency stems from a careful balance between estimator variance and computational cost.

**Gap in the Literature** While stochastic approximation for quantiles and variance reduction techniques have both been extensively studied, their combination remains largely unexplored, particularly in the context of *online* estimation of tail risk measures. Moreover, the non-smooth nature of quantile objectives poses additional challenges for integrating modern variance reduction schemes.

## 4 Research Objectives and Expected Contributions

The project will address the following objectives:

- Design novel online stochastic algorithms for quantile and CVaR estimation with adapted variance reduction techniques (notably SPIDER) to non-smooth and indicator-based objectives.
- Analyze convergence properties (almost sure convergence, rates, asymptotic distribution).
- Develop scalable implementations for high-dimensional and streaming settings.

The research will combine tools from stochastic approximation theory, variance reduction methods, Non-smooth optimization, among other. A key direction will be to construct hybrid algorithms where classical Robbins-Monro updates are enhanced with recursive variance-reduced estimators inspired by SPIDER. The project is rich in challenges and open problems.

### Theoretical challenges

- Extending SPIDER-type estimators to non-smooth objectives involving indicator functions.
- Controlling bias and stability in recursive variance-reduced schemes.
- Establishing non-asymptotic bounds for stochastically variance reduced quantile and CVaR estimation.
- Handling heavy-tailed distributions and rare-event regimes.

### Practical Challenges

- Designing algorithms with low memory footprint suitable for online settings.
- Ensuring robustness to model misspecification and dependent data.
- Applying the methods to real-world problems in particular industrial ones.

In particular, this project is strongly motivated by real-world applications arising in collaboration with EDF within the ANR project GATSBII. In this context, large-scale industrial data generated from energy systems (e.g., power production, network management, and maintenance processes) exhibit complex stochastic behaviors, including heavy tails, temporal dependence, and rare but critical events (see, e.g., [16]).

Accurate estimation of extreme quantiles and superquantiles is crucial for risk assessment tasks such as anticipating system failures, managing peak demand, and optimizing maintenance strategies under uncertainty. However, the volume and streaming nature of these data require online and computationally efficient methods, beyond classical batch approaches [17].

The proposed variance-reduced stochastic algorithms are particularly well-suited to this setting, as they enable real-time estimation with controlled computational cost and improved statistical efficiency. In particular, integrating SPIDER-type techniques into stochastic approximation schemes is expected to significantly enhance convergence speed and robustness when dealing with noisy industrial signals.

This collaboration will provide access to realistic datasets and use cases, allowing for thorough empirical validation of the proposed methods. It will also foster the development of algorithms tailored to the specific constraints of industrial environments, such as scalability, reliability, and interpretability. Ultimately, this work aims to contribute to the deployment of advanced risk-aware decision tools in energy systems.

## References

- [1] H. Robbins and S. Monro (1951). A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3), 400–407.
- [2] J. R. Blum (1954). Multidimensional stochastic approximation methods. *Annals of Mathematical Statistics*, 25(4), 737–744.
- [3] V. Fabian (1968). On asymptotic normality in stochastic approximation. *Annals of Mathematical Statistics*, 39(4), 1327–1332.
- [4] H.-F. Chen (2002). *Stochastic Approximation and Its Applications*. Kluwer Academic Publishers.
- [5] R. T. Rockafellar and S. Uryasev (2000). Optimization of Conditional Value-at-Risk. *Journal of Risk*, 2(3), 21–41.
- [6] A. Shapiro, D. Dentcheva, and A. Ruszczyński (2014). *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2nd edition.
- [7] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone (2015). Risk-sensitive and robust decision-making: A CVaR optimization approach. *Advances in Neural Information Processing Systems (NeurIPS)*, 28, 1522–1530.
- [8] S. Asmussen and P. W. Glynn (2007). *Stochastic Simulation: Algorithms and Analysis*. Springer.
- [9] R. Johnson and T. Zhang (2013). Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems (NeurIPS)*, 26, 315–323.
- [10] A. Defazio, F. Bach, and S. Lacoste-Julien (2014). SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in Neural Information Processing Systems (NeurIPS)*, 27, 1646–1654.
- [11] L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč (2017). SARAH: A novel method for machine learning problems using stochastic recursive gradient. *Proceedings of ICML*, 70, 2613–2621.
- [12] C. Fang, C. J. Li, Z. Lin, and T. Zhang (2018). SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 689–699.
- [13] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.
- [14] B. Bercu, M. Costa, and S. Gadat (2021). Stochastic approximation algorithms for superquantiles estimation. *Electronic Journal of Probability*, 26, paper no. 84, 1–26.
- [15] M. Costa and S. Gadat (2021). Non-asymptotic study of a recursive superquantile estimation algorithm. *Electronic Journal of Statistics*, 15(2), 4718–4769.
- [16] B. Iooss and J. Lonchamp, Robust tuning of Robbins-Monro algorithm for iterative uncertainty quantification - Application to wind-farm asset management. In: *B. Castanier, M. Cepin, D. Bigaud and C. Berenguer (Eds), Proceedings of the 31st European Safety and Reliability Conference*, Research Publishing, Singapore, 2021.
- [17] A. Ribés, T. Terraz, Y. Fournier, B. Iooss and B. Raffin, Unlocking large scale uncertainty quantification with in transit iterative statistics, In: *H. Childs, J. Bennet and C. Garth (Eds), In situ visualization for computational science*, Springer, 2022.