*ECAI 2020*
*G.D. Giacomo et al. (Eds.)*


*doi:10.3233/FAIA200443*

# Hybrid Offline/Online Optimization Under Uncertainty

**Allegra De Filippo** and **Michele Lombardi** and **Michela Milano** [1]

**Abstract.** In this work we consider optimization problems that require to make *interdependent offline and online decisions* under uncertainty. We broadly refer to long-term strategic decisions as *offline* and to short-term operational decisions as *online*. For example, in Distributed Energy Management Systems we may need to define (offline) a daily production schedule for an industrial plant, and then manage (online) its power supply on a hour by hour basis. Traditionally offline and online phases are tackled in isolation, leading to some drawbacks: offline decisions are taken without regard for the capabilities of the downstream online solver; while the applicability of the best approaches for online decisions (e.g. anticipatory algorithms) is limited by the need to provide high responsiveness. Starting from a (literature-based) baseline, we define general methods for leading to significant quality improvements, at the cost of an increased computation effort either in the offline or the online phase. All our methods have broad applicability, and provide multiple options to balance the solution quality/time trade-off, suiting a variety of practical application scenarios with both offline and online decisions and featuring continuous and discrete decisions. An extensive analysis of the experimental results shows that offline/online integration may lead to substantial benefits.

## 1 Introduction

Optimization under uncertainty arises in many application areas[4], such as project scheduling, transportation systems, and energy management: fuel prices, activity durations, travel times, etc. are effectively stochastic in the real world. Optimization problems in this class can be seen as a sequence of multiple stages, such that at each stage part of the uncertainty is revealed and some decisions must be made. The need to account for multiple future developments makes stochastic optimization incredibly challenging[5, 6], which explains how approximate (sampling-based) methods and heuristics are the most popular solution techniques in practice. Due to such a complexity, the applicable approaches depend on the temporal granularity of the decisions to be made. Long-term "strategic" decisions (which are often very impactful) are typically solved via expensive, but more accurate, sampling-based approaches. Short-term "operational" decisions often need to be made over multiple steps, within a short time frame: they are commonly addressed via polynomial-time heuristics, while more advanced sampling-based methods (e.g. online anticipatory algorithms[3]) are applicable only if their computational cost is carefully managed. We will refer to the first class of problems (and solution approaches) as *offline* and to the second as *online*.

We move from the observation that many practical application scenarios combine *both an online and an offline phase*. The simplest approach to tackle such problems is to deal with the offline and online phase separately, resp. (e.g.) via a sampling-based method and a heuristic. However, we show that *substantial improvements can be obtained by treating the two phases in an integrated fashion*. This paper shows methods and results of our works [1, 2].

**Motivating Examples** Some real world use cases are typically solved via either offline or online models, while in fact they are *integrated offline/online problems*.

- *Energy Management Systems* (EMS) are key components of the electrical grid that maintain its stability both by shifting consumption (over time) and routing power flows from the available generators. In practice, the load shifts must be planned offline (the day ahead) and the power flow balance should be maintained online (e.g. hour by hour), so as to minimize the costs.
- In *transportation systems*, a central role is played by the Vehicle Routing Problem and its variants, which consists in establishing the paths for a set of vehicles to serve a set of customers. In a real world setting, many aspects (e.g. customer demands and travel times) are also subject to uncertainty. Several transportation companies focus on assigning *offline* customers to smaller scale operators, which are then in charge of choosing the routes *online*.

## 2 Literature-Based Baselines

Some of our techniques build on a (generic, *existing*) myopic online heuristic, other over a (generic, *existing*) online anticipatory algorithm. Then we propose the following key ideas: 1) Using (conditional) density estimation to produce samples that are more likely given the observed uncertainty (this yields dynamic sampling for the anticipatory algorithm); 2) Solving (offline) a large amount of past instances (with the anticipatory algorithm) to build a "contingency" table. In the online phase, we then use our conditional density estimation to find which pre-solved instance is the most compatible with the observed uncertainty and decisions, and we try to follow its prescribed decisions as well as possible (that is the job of the, rather fast, fixing heuristic); 3) Using (expensive, offline) parameter tuning to improve the online myopic heuristic (which stays very fast); 4) Many approaches for offline stochastic optimization with multiple online stages (e.g. offline scheduling or routing) rely on an approximation of the online problem obtained by collapsing all online stages into a single recourse stage. This eliminates non-anticipativity constraints, so that the decisions made when solving the approximate online problem have knowledge of the future. Unfortunately, the heuristic used to solve the *true* online problem is often myopic, which creates a discrepancy between the offline solution and the capabilities of the online solver. We enforce the myopic behavior in the online problem approximation (and obtain more realistic solutions)

[1] DISI, University of Bologna, Italy, email: {allegra.defilippo, michele.lombardi2, michela.milano}@unibo.it

by using a mathematical characterization of the online heuristic (its KKT conditions) as a constraint in the offline problem.

## 3 Focus on the Online Phase

The most natural way to improve online decisions consists in replacing the greedy heuristic with a sampling-based anticipatory algorithm.

**Online Anticipatory Algorithm** Sampling-based algorithms rely on scenarios to estimate future outcomes. Formally, a scenario $\omega$ specifies a value $\xi_i^\omega$ for all the random variables. Given a set $\Omega$ of scenarios, the system state $s_k$, and values for $\xi_O$ corresponding to the observed uncertainty, we assume that $A$ can compute the decisions for stage $k$: $x_k = A(s_k, \xi_O, \{\xi^\omega\}_{\omega\in\Omega})$
Once the decisions are computed, the next state can be determined. This is controlled via a state transition function *next* that, based on the current state, decisions, and observed uncertainty, computes: $s_{k+1} = next(s_k, x_k, \xi_O)$ , given the initial state $s_1$, a set of scenarios $\Omega$, and a set of values sampled from $\xi_k$ (which represent the online observations). The $O$ set is assumed to be initially empty. This generic anticipatory algorithm will be referred to as ANTICIPATE.

**Dynamic Scenario Sampling** Using a fixed set of scenarios (as in ANTICIPATE) is beneficial when the $\xi$ variables are statistically independent. When they are not, however, the set of scenarios may loose relevance as uncertainty is resolved. For example, a scenario based on a cloudy day forecast becomes less likely if fair weather is observed at the beginning of online execution. Defining a representative set of scenarios is critical for the approach effectiveness and it is usually done by exploiting the available offline information. Here, we assume that the such offline information is a collection of observed uncertain values. This definition captures many practical cases (e.g. forecasts or predictions, historical data), but we still focused on the improvement of the online phase with the ANTICIPATE algorithm. Formally, at stage $k$ we wish to sample scenarios that are likely to occur given the past observations, i.e. to sample the unobserved variables $\xi_{\bar{O}}$ according to the conditional distribution $P(\xi_{\bar{O}} \mid \xi_O)$. These probabilities can be approximated via any density estimation method, such as Kernel Density Estimation, or recent Deep Learning techniques such as Normalizing Flows. *Any such method can be trained on the offline information to obtain an estimator $\tilde{P}(\xi)$ for the joint distribution of the random variables*, yielding scenarios with a distribution that takes into account the observed values. This technique leads to the definition of our method ANTICIPATE-D that provides a modest advantage in terms of solution time, but can match and surpass ANTICIPATE in terms of quality.

## 4 Focus on the Offline Phase

As a second research direction we consider approaches that shift the computational load to the offline phase by exploiting the available information to manage the cost/quality tradeoff of online algorithms.

**Offline Online-Aware Phase** When stringent time constraints on the online decisions exist, it may be better to improve the greedy heuristic by simply adjusting its parameters. This is the main idea in the TUNING approach: this maintains the efficiency of the original greedy heuristic, at the price of a computationally expensive parameter tuning process, which is however performed offline. Shifting our attention to the offline decisions, we can also mitigate the discrepancy by translating the online greedy heuristic as a set of constraints, which can be injected in the offline (two-stage) model. This techniques leads to our ACKNOWLEDGE method. Interestingly, we show that the approach can be combined with parameter tuning to achieve even deeper integration: this idea is explored in our ACTIVE method.

**Offline Contingency Table** If a significant amount of time is available in the offline phase, we can exploit the offline information in an alternative fashion, by trying to prepare for each likely future development. Intuitively, we can treat each scenario as if it were an actual sequence of online observations, and process it via some anticipatory algorithm. By doing this, we build a pool of solutions that can then be used to guide an online method and we define our methods CONTINGENCY and CONTINGENCY-D (respectively without and with dynamic scenario sampling). We use the traces from $T$ to guide an efficient *fixing heuristic*, which tries to choose decisions having the largest chance of being optimal. Formally, the fixing heuristic solves: $\arg\max\{P^*(x_k \mid s_k\xi_O) : x_k \in X_k\}$, where $P^*$ is the probability that the chosen $x_k$ is optimal, given the state $s_k$ and the observed uncertainty. The $X_k$ set represents the feasible decision space, which is defined via problem-dependent constraints and auxiliary variables.

## 5 Results and Conclusions

To test our methods, we ground them on the two mentioned case studies. The EMS problem features a continuous (non-enumerable) decision space, while the second has pure discrete decisions. In our experiments, all the proposed methods improve over the baseline in terms of solution/quality trade-off. In the first case, we use an anticipatory algorithm both as a starting point and as a reference. Idea 1 (see Section 2) improves the anticipatory algorithm by a modest $\sim 1\%$ margin (same run time); idea 2 looses a small 1-2% margin in terms of quality, but with a massive reduction (two orders of magnitude) of the online run time. In the second setup, our starting point is an offline solver relying an a two stage approximation and a myopic online heuristic; we use an oracle (with perfect information) as a performance reference. Idea 3 can reduce the gap w.r.t. the oracle from $\sim 21\%$ to $\sim 7\%$; idea 4 goes even further, yielding a gap of $\sim 2\%$ w.r.t. the oracle. In both cases the online part is handle via the original, very fast, heuristic.

## ACKNOWLEDGEMENTS

This work has been partially supported by the H2020 Project AI4EU (g.a. 825619), the VIRTUS Project (CCSEB-00094).

## REFERENCES

[1] A. De Filippo, M. Lombardi, and M. Milano, 'Methods for off-line/online optimization under uncertainty.', in *IJCAI*, pp. 1270–1276, (2018).
[2] A. De Filippo, M. Lombardi, and M. Milano, 'How to tame your anticipatory algorithm', in *IJCAI*, pp. 1071–1077, (2019).
[3] L. Mercier and P. Van Hentenryck, 'Performance analysis of online anticipatory algorithms for large multistage stochastic integer programs.', in *IJCAI*, pp. 1979–1984, (2007).
[4] W. B. Powell, *A Unified Framework for Optimization Under Uncertainty*, chapter 3, 45–83.
[5] A. Shapiro and A. Philpott, 'A tutorial on stochastic programming', *Manuscript. Available at www2. isye. gatech. edu/ashapiro/publications. html*, **17**, (2007).
[6] P. Van Hentenryck and R. Bent, *Online stochastic combinatorial optimization*, The MIT Press, 2009.