How to Tame Your Anticipatory Algorithm

Allegra De Filippo, Michele Lombardi, and Michela Milano

DISI, University of Bologna, Bologna, Italy
{allegra.defilippo,michele.lombardi2,michela.milano}@unibo.it

Abstract. Sampling-based anticipatory algorithms can be very effective at solving online optimization problems under uncertainty, but their computational cost may be prohibitive in some cases. Given an arbitrary anticipatory algorithm, we present three methods that allow to retain its solution quality at a fraction of the online computational cost, via a substantial degree of *offline* preparation. Our approaches are obtained by combining: 1) a simple technique to identify likely future outcomes based on past observations; 2) the (expensive) offline computation of a "contingency table"; and 3) an efficient solution-fixing heuristic. We ground our techniques on two case studies: an energy management system with uncertain renewable generation and load demand, and a traveling salesman problem with uncertain travel times. In both cases, our techniques achieve high solution quality, while substantially reducing the online computation time.

1 Introduction

Optimization problems under uncertainty arise in many important application domains, such as production scheduling or energy system management. These problems often benefit from making all or part of their decisions online, reacting and adapting to external events. In this context, stochastic online anticipatory algorithms have proved particularly effective (see e.g. [11]). However, many of such algorithms have a considerable computational cost, which may be problematic if (as it is often the case) online decisions must be taken within a short time frame.

In most practical settings, however, a substantial amount of time and information is available before the online problem is solved, in an *offline phase*. We refer to this sort of data as *offline information*. Usually, it is employed to characterize the uncertain elements and for sampling likely outcomes (i.e. scenarios). We will show how to exploit this information at a much deeper level.

This paper reports and summarizes the methods and the results presented in [7] and accepted for publication at IJCAI 2019.

We propose three hybrid offline/online methods that build over a given, sampling-based, anticipatory algorithm, and allow to match its solution quality at a fraction of the online computational cost: 1) a technique to estimate the probability of future outcomes, given past observations; 2) a scheme for building a "contingency table", with precomputed solutions to guide the online choices; and 3) an efficient fixing heuristic for adapting the precomputed solutions to run-time conditions.

We ground our approaches on a (numeric) energy management problem with uncertain loads and generation from Renewable Energy Sources (RES), and on a (discrete)

Traveling Salesman Problem with uncertain travel times. We show how our methods reach a solution quality comparable with the anticipatory algorithm, with lower (or dramatically lower) online computational cost.

2 Background and Related Work

Historically, the literature on optimization under uncertainty has focused on offline problems [17, 19]. These methods usually rely on sampling (yielding a number of *scenarios*) to obtain a statistical model of future uncertainty.

More recently, improvements in the solution techniques and computational power have enabled the application of similar methods to online optimization. This lead to so-called *online anticipatory algorithms*, which proved very effective at finding robust, high quality, solutions as uncertainty slowly reveals itself [11]. Notable algorithms and results are described in [5, 2, 3, 15, 6].

Online anticipatory algorithm typically rely on scenario sampling to estimate the possible developments for a fixed number of future steps, known as *look-ahead horizon*. Larger sample sizes result in higher accuracy, but also in more and bigger (possibly NP-hard) problems to be solved. Considerable research effort has therefore focused on improving the efficiency of these algorithms [2, 4, 14]. The approaches from [12, 16, 13] attempt instead to reduce the number of scenarios by increasing their relevance by taking into account past observations while sampling.

3 "Taming" an Anticipatory Algorithm

Our goal is reducing the online computational cost of a given sampling-based anticipatory algorithm, referred to as A, by exploiting the existence of an offline phase. Such A algorithm is the main input for all our methods.

Similarly to [11], we view online optimization under uncertainty as a stochastic n-stage problem. At each stage some uncertainty is resolved, and some decision must be made. A stage k is associated to a decision variable x_k (e.g. the power flows between loads and generators) and a state variable s_k (summarizing the effect of past decisions). All variables may be vector-valued.

We assume that uncertainty is *exogenous*, i.e. not affected by the decisions (e.g. the RES generation does not depend on how we choose to route it), and modeled via a set of random variables ξ_i. Which variables are observed at each stage depends on the state, and is controlled by a *peek* function: $O = peek(s_k)$ which returns a set O with the indices of the observed variables. We will use the notation ξ_O to denote the observed ξ variables, and $\xi_{\bar{O}}$ for the unobserved ones.

Base Anticipatory Algorithm Sampling-based algorithms rely on scenarios to estimate future outcomes. Formally, a scenario ω specifies a value ξ_i^ω for all the random variables. Given a set Ω of scenarios, the system state s_k, and values for ξ_O corresponding to the observed uncertainty, we assume that A can compute the decisions for stage k:

$$x_k = A(s_k, \xi_O, \{\xi^\omega\}_{\omega \in \Omega}) \tag{1}$$

Algorithm 1 ANTICIPATE (s_1, ξ)

Sample a set of scenarios Ω
for $k = 1 \ldots n$ **do**
    $O = O \cup peek(s_k)$ *# observe uncertainty*
    $x_k = A(s_k, \xi_O, \{\xi^\omega\}_{\omega \in \Omega})$ *# obtain decisions*
    $s_{k+1} = next(s_k, x_k, \xi_O)$ *# determine next state*
return s, x

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:

$$s_{k+1} = next(s_k, x_k, \xi_O) \tag{2}$$

Given the initial state s_1, a set of scenarios Ω, and a set of values sampled from ξ_k (which represent the online observations), we can simulate the execution of the method via Algorithm 1. The O set is assumed to be initially empty. This generic anticipatory algorithm will be referred to as ANTICIPATE in the remainder of the paper.

OFFLINE INFORMATION Defining a representative set of scenarios Ω 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, data from test runs). More importantly, this means that *the offline information is in fact a collection of scenarios*. We will denote the offline information as I, index its element with ω, and assume (as it is usual) that I is representative of the true probability distribution of the random variables. A set Ω of scenarios for ANTICIPATE can be obtained by sampling a number of elements from I.

3.1 Basic Contributions

All our methods rely on three techniques, which will be described in this section.

PROBABILITY ESTIMATION FOR SCENARIO SAMPLING: Using a fixed set of scenarios (as in ANTICIPATE) is beneficial when the ξ_i 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. 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 $P(\xi_{\bar{O}} \mid \xi_O)$. If we draw the scenarios from the offline information (which guarantees physically meaningfulness), then sampling requires to estimate the conditional probabilities *of the elements in I*. From basic probability theory, this is given by the ratio of two joint probabilities:

$$\forall \omega \in I, \quad P(\xi_{\bar{O}}^\omega \mid \xi_O) = \frac{P(\xi_O \xi_{\bar{O}}^\omega)}{P(\xi_O)} \tag{3}$$

where $P(\xi_O \xi_{\bar{O}}^\omega)$ is the probability that observed values occur together with the remaining predictions from the scenario, and $P(\xi_O)$ is the probability that the values are observed. The joint probability at the numerator can be approximated via any density

---

Algorithm 2 BUILDTABLE (s_1, AA)

---

for $\omega \in I$ **do**
    $s^\omega, x^\omega = AA(s_1, \xi^\omega)$
return $\{\xi^\omega, s^\omega, x^\omega\}_{\omega \in T}$

---

estimation method, such as Kernel Density Estimation [20], Gaussian Mixture Models [10], or recent Deep Learning techniques such as Normalizing Flows [18] and Real NVP [8]. *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.* An estimator for the distribution $P(\xi_O)$ at the denominator can then be derived from $\tilde{P}(\xi)$ via marginalization, i.e. by averaging the contribution of the unobserved variables. We perform this step over all possible completions of the observed values *in the offline information.* Overall, we have:

$$\forall \omega \in I, \quad \tilde{P}(\xi_{\bar{O}}^\omega \mid \xi_O) = \frac{\tilde{P}(\xi_O \xi_{\bar{O}}^\omega)}{\sum_{\omega' \in T} \tilde{P}(\xi_O \xi_{\bar{O}}^{\omega'})} \tag{4}$$

This estimator defines a discrete distribution over the offline information I. The chosen marginalization technique guarantees an estimate that is approximately proportional (not approximately equal) to the true $P(\xi_O)$. Hence, we have that:

$$\forall \omega \in I, \quad P(\xi_{\bar{O}}^\omega \mid \xi_O) \propto \tilde{P}(\xi_{\bar{O}}^\omega \mid \xi_O) \tag{5}$$

Sampling from I according to Equation (4) yields scenarios with a distribution that takes into account the observed values.

BUILDING A CONTINGENCY TABLE: If a significant amount of time is available in the offline phase, we can exploit the offline information more aggressively, by trying to prepare for each likely future development. Intuitively, we can *treat each scenario $\omega \in I$ 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. The process is outlined in Algorithm 2, which requires as input the initial state s_1 of the system, and a solution algorithm AA, accepting the same parameters as ANTICIPATE. The result is an augmented version of the offline information, where each scenario ω is additionally associated to the sequence of states s^ω visited by the algorithm and its sequence of decisions x^ω. We refer to this data structure as *contingency table*, and to its elements as *traces*. We denote the table as T.

FIXING HEURISTIC: 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, it solves:

$$\arg\max\{P^*(x_k \mid s_k \xi_O) : x_k \in X_k\} \tag{6}$$

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. Closed-forms for P^* can be obtained separately for discrete and numeric problems, based on the contingency table.

Algorithm 3 FIXING (s_1, ξ, T)

---

for $k = 1 \ldots n$ **do**
    $O = O \cup peek(s_k)$
    $\Omega =$ top elements in T by descending Equation (9)
    Compute p_{jv} and/or p_ω based on Ω
    Solve Equation (7)/(10) to obtain x_k
    $s_{k+1} = next(s_k, x_k, \xi_O)$
return s, x

Overall, *in case of discrete decisions, the problem from Equation* (6) *translates into*:

$$\arg\min \left\{ -\sum_{j=1}^{m} \sum_{v \in D_j} \log p_{jv} [\![x_{kj} = v]\!] : x_k \in X_k \right\} \tag{7}$$

where $[\![\cdot]\!]$ denotes the truth value of a predicate, D_j is the domain of x_{kj}, and:

$$p_{jv} = \frac{\sum_{\omega \in T, x_{kj}^\omega = v} P(\omega)}{\sum_{\omega \in T} P(\omega)} \tag{8}$$

Here, $P(\omega)$ is a compact notation for the probability that *we reach the same state as trace ω, and then everything goes according to plan*. It can be approximated using:

$$\forall \omega \in T, \quad P(\omega) \propto \tilde{P}(s_{sk+1}^\omega \mid s_k) \tilde{P}(\xi_{\bar{O}}^\omega \mid \xi_O) \tag{9}$$

where $\tilde{P}(\xi_{\bar{O}} \mid \xi_O)$ is the estimator from Equation (4), and $\tilde{P}(s_{sk+1} \mid s_k)$ is a second estimator obtained via similar means. The cost function in Equation (7) is linear if a one-hot encoding is adopted for x_{kj}, and the size of T affects only the computation of the p_{jv} values. Overall, the problem is efficient to solve. *In case of numeric decisions, we have instead*:

$$\arg\min \left\{ \sum_{j=1}^{m} \sum_{\omega \in T} p_\omega \frac{1}{2\sigma_j} (x_{kj} - x_{kj}^\omega)^2 : x_k \in X_k \right\} \tag{10}$$

with:

$$p_\omega = \frac{P(\omega)}{\sum_{\omega' \in T} P(\omega')} \tag{11}$$

The cost function is quadratic and convex, and the problem size is small due to the same arguments as Equation (7).

Intuitively, the discrete version of the heuristic is related to minimizing weighted discrepancies w.r.t. the traces in T, i.e. to weighted Hamming distances. The numeric version is instead related to weighted Euclidean distances. The pseudo-code for the heuristic is provided in Algorithm 3. The only difference with the process described so far is that the p_{jv} and p_ω probabilities may be computed based on a subset Ω of the full contingency table. This may be useful to bias the choice of the online decision according to the most relevant traces.

Algorithm 4 ANTICIPATE-D (s_1, ξ)

Train the $\tilde{P}(\xi)$ estimator on I
for $k = 1 \dots n$ **do**
    $O = O \cup peek(s_k)$
    Sample Ω from T, according to Equation (4)
    $x_k = A(s_k, \xi_O, \{\xi^\omega\}_{\omega \in \Omega})$
    $s_{k+1} = next(s_k, x_k, \xi_O)$
return s, x

Algorithm 5 CONTINGENCY (s_1, ξ)

Train the $\tilde{P}(\xi)$ estimator on I
$T = $ BUILDTABLE$(s_1, $ ANTICIPATE$)$
Train the $\tilde{P}(s_k s_{k+1})$ estimators on T, for all k
$s, x = $ FIXING(s_1, ξ, T)
return s, x

3.2   Hybrid Offline/Online Methods

Our three solution methods can now be defined with relative ease, by combining the techniques just described.

ANTICIPATE-D Our first hybrid method is obtained from ANTICIPATE by simply replacing the static set of samples with a dynamically adjusted one. The dynamic set can be populated according to the estimated probabilities from Equation (4), so as not to loose relevance: this may enable to reach similar solution qualities with fewer scenarios, at the cost of training an estimator offline. We refer to this approach as ANTICIPATE-D, and its pseudo-code is in Algorithm 4

CONTINGENCY The second method is based on the idea of computing robust solutions for the scenarios in the offline information, and then use them as guidance for the FIXING heuristic. Robust solutions are obtained by using ANTICIPATE, so that hopefully the (fast) fixing heuristic will be able to match their quality: the price to pay is a hefty offline computational effort. We refer to this approach as CONTINGENCY, and its pseudo-code is reported in Algorithm 5.

CONTINGENCY-D Our final method is similar to the previous one, except that the contingency table is populated with *non-robust* solutions. This is done by using ANTICIPATE *with a single scenario, given by the values of ξ^ω (i.e. the pretend online observations)*. This technique (referred to as ANTICIPATE$_1$) provides perfect information about the future, so that achieving robustness is entirely delegated to the FIXING heuristic. The approach is likely to loose reliability, but has two important advantages: 1) lower offline computational costs; and 2) while ANTICIPATE is a stochastic algorithm, ANTICIPATE$_1$ is deterministic. So, this method may provide anticipatory-like results even when no anticipatory algorithm is available. We refer to this method as CONTINGENCY-D, and its pseudo-code is reported in Algorithm 6.

4   Grounding the Approaches

Grounding our approaches requires to specify: 1) the x, s and ξ variables, 2) the $peek$ and $next$ functions, 3) the sampling-based algorithm A, and 4) the feasible space X_k

---

Algorithm 6 CONTINGENCY-D (s_1, ξ)

Train the $\tilde{P}(\xi)$ estimator on I
$T = \text{BUILDTABLE}(s_1, \text{ANTICIPATE}_1)$
Train the $\tilde{P}(s_k s_{k+1})$ estimators on T, for all k
$s, x = \text{FIXING}(s_1, \xi, T)$
return s, x

---

for the FIXING heuristic. Additionally, evaluating the solution quality requires to define 5) a cost metric.

    We show how this can be done in two case studies: 1) a Virtual Power Plant energy management problem with numerical decisions; and 2) a combinatorial Traveling Salesman Problem with uncertain travel times. In both cases, the input anticipatory algorithm A is given by a Mathematical Programming model, based on the Sample Average Approximation. The models are slight improvements over those by [6], whose work brought to attention the interplay between offline and online phases. Both approaches are serviceable, but not necessarily representative of the state-of-the-art (especially for the TSP). We focus on the minimal information needed.

5   Experimentation

We empirically evaluated the three hybrid offline/online methods on realistic instances for the two case studies. The baseline in both cases are myopic heuristics, obtained by running ANTICIPATE with an empty set of scenarios.

    Our methods are evaluated over different uncertainty realizations, obtained by sampling the random variables for the loads and RES generation in the VPP, and for the travel times in the TSP. In both cases, we use models of uncertainty that ensure realistic statistical dependence between the variables. These models are used to obtain the offline information I and the sequences of observations for the actual experiments. For the VPP, grid electricity prices change every 15 minutes, which is also the duration of our online stages. New offline information (e.g. market prices) becomes available every day, hence our horizon corresponds to $24 \times 4 = 96$ stages. We use (real) physical bounds for power generation from [1, 9]. The initial battery state, efficiency, and power flow limit are also based on real data [1, 9]. The same goes for the data, which is also from [9]. Different instances have then been obtained by manually scaling load and RES generation. For the TSP we use classical benchmarks[1], by including problems from 10 to 40 nodes. In the TSP each stage represents a visit, hence our horizon corresponds to the total number of nodes. We use Kernel Density Estimation (KDE with Gaussian Kernels) to obtain all approximate distributions. As an underlying solver we use Gurobi [2], which can handle both MILPs and Quadratic Programs. Each evaluated algorithm (in each considered configuration) is run 50 times, with the same 50 sequences of realizations. We use a time limit of 300 seconds for both the VPP and the TSP. For each run we record both the time required by each approach and the corresponding solution cost, and we report their average values over the 50 realizations. In

---

[1] http://myweb.uiowa.edu/bthoa/TSPTWBenchmarkDataSets.htm
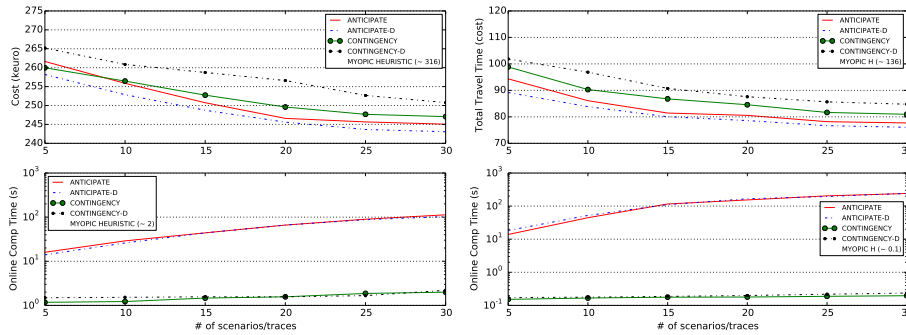
[2] Available at http://www.gurobi.com

Fig. 1: (left) Methods solution/quality comparison for the VPP. (right) For the TSP.

all cases, $|I| = |T| = 100$, and for the CONTINGENCY method, the contingency table is built using ANTICIPATE with 20 scenarios. Due to space constraints, we report results only for a few representative instances.

The offline training times of the KDE models are roughly the same for all the three hybrid methods (~ 65 sec for the VPP and ~ 32 sec for the TSP). Building the contingency tables for CONTINGENCY takes $\sim 6,000$ sec in the VPP and $\sim 15,000$ sec in the TSP, but only ~ 400 and $\sim 2,000$ sec for CONTINGENCY-D. In Section 5 we show the cost/quality tradeoff of the proposed methods and of ANTICIPATE for the VPP (base instance) and for the TSP (a representative 20 customers instance). The use of a dynamic Ω set of scenarios allow ANTICIPATE-D to work better than ANTICIPATE. The CONTINGENCY method is surprisingly close in terms of quality to the original anticipatory algorithm, especially considered that its online computational cost is dramatically smaller (one to two orders of magnitude). CONTINGENCY-D performs (as expected) slightly worse than CONTINGENCY, but it still much better than the myopic heuristic. Increasing the number of guiding traces is beneficial for both methods, and in particular for CONTINGENCY-D.

6 Conclusion

We have presented three methods that can be applied *to a generic anticipatory algorithm* to reduce its online computational effort by exploiting offline information. In particular, both CONTINGENCY and CONTINGENCY-D are dramatically faster than ANTICIPATE during online operation. Between the two of them CONTINGENCY is significantly more reliable in terms of quality, but may require a substantial amount of offline computation. The ANTICIPATE-D technique provides a modest advantage in terms of solution time, but can match and even surpass ANTICIPATE in terms of quality.

The ability to shift part of the computational cost to an offline stage provides a significant degree of flexibility to stochastic anticipatory algorithm, and likely to increase their applicability. We believe there is room for improving the scalability and efficiency of our methods, and achieving this goal is part of our current research directions.

References

1. Bai, H., Miao, S., Ran, X., Ye, C.: Optimal dispatch strategy of a virtual power plant containing battery switch stations in a unified electricity market. Energies **8**(3), 2268–2289 (2015)
2. Bent, R., Van Hentenryck, P.: Regrets only! online stochastic optimization under time constraints. In: AAAI. vol. 4, pp. 501–506 (2004)
3. Bent, R., Van Hentenryck, P.: Online stochastic optimization without distributions. In: ICAPS. vol. 5, pp. 171–180 (2005)
4. Borodin, A., El-Yaniv, R.: Online computation and competitive analysis. cambridge university press (2005)
5. Chang, H.S., Givan, R., Chong, E.K.: On-line scheduling via sampling. In: AIPS. pp. 62–71 (2000)
6. De Filippo, A., Lombardi, M., Milano, M.: Methods for off-line/on-line optimization under uncertainty. In: IJCAI. pp. 1270–1276 (2018)
7. De Filippo, A., Lombardi, M., Milano, M.: How to tame your anticipatory algorithm. In: *to be published at* IJCAI (2019)
8. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real nvp. arXiv preprint arXiv:1605.08803 (2016)
9. Espinosa, A.N., Ochoa, L.N.: Dissemination document "low voltage networks models and low carbon technology profiles". Tech. rep., University of Manchester (June 2015)
10. Gauvain, J.L., Lee, C.H.: Maximum a posteriori estimation for multivariate gaussian mixture observations of markov chains. IEEE transactions on speech and audio processing **2**(2), 291–298 (1994)
11. Hentenryck, P.V., Bent, R.: Online stochastic combinatorial optimization. The MIT Press (2009)
12. John, G.H., Langley, P.: Static versus dynamic sampling for data mining. In: KDD. vol. 96, pp. 367–370 (1996)
13. Lin, M., Tang, K., Yao, X.: Dynamic sampling approach to training neural networks for multiclass imbalance classification. IEEE Transactions on Neural Networks and Learning Systems **24**(4), 647–660 (2013)
14. Mercier, L., Van Hentenryck, P.: Performance analysis of online anticipatory algorithms for large multistage stochastic integer programs. In: IJCAI. pp. 1979–1984 (2007)
15. Mercier, L., Van Hentenryck, P.: Amsaa: A multistep anticipatory algorithm for online stochastic combinatorial optimization. Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems pp. 173–187 (2008)
16. Philpott, A.B., De Matos, V.L.: Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. European Journal of Operational Research **218**(2), 470–483 (2012)
17. Powell, W.B.: A Unified Framework for Optimization Under Uncertainty, chap. 3, pp. 45–83. INFORMS (2016)
18. Rezende, D.J., Mohamed, S.: Variational inference with normalizing flows. arXiv preprint arXiv:1505.05770 (2015)
19. Sahinidis, N.V.: Optimization under uncertainty: state-of-the-art and opportunities. Computers & Chemical Engineering **28**(6), 971 – 983 (2004), http://www.sciencedirect.com/science/article/pii/S0098135403002369, fOCAPO 2003 Special issue
20. Silverman, B.W.: Density estimation for statistics and data analysis. Routledge (2018)