





Contextual Robust Optimisation with Uncertainty Quantification

Egon Peršak^(✉)  and Miguel F. Anjos^(D) 

University of Edinburgh, Edinburgh, UK
E.Persak@sms.ed.ac.uk

Abstract. We propose two pipelines for convex optimisation problems with uncertain parameters that aim to improve decision robustness by addressing the sensitivity of optimisation to parameter estimation. This is achieved by integrating uncertainty quantification (UQ) methods for supervised learning into the ambiguity sets for distributionally robust optimisation (DRO). The pipelines leverage learning to produce contextual/conditional ambiguity sets from side-information. The two pipelines correspond to different UQ approaches: i) explicitly predicting the conditional covariance matrix using deep ensembles (DEs) and Gaussian processes (GPs), and ii) sampling using Monte Carlo dropout, DEs, and GPs. We use i) to construct an ambiguity set by defining an uncertainty around the estimated moments to achieve robustness with respect to the prediction model. UQ ii) is used as an empirical reference distribution of a Wasserstein ball to enhance out of sample performance. DRO problems constrained with either ambiguity set are tractable for a range of convex optimisation problems. We propose data-driven ways of setting DRO robustness parameters motivated by either coverage or out of sample performance. These parameters provide a useful yardstick in comparing the quality of UQ between prediction models. The pipelines are computationally evaluated and compared with deterministic and unconditional approaches on simulated and real-world portfolio optimisation problems.

Keywords: Prediction and Optimisation · Prescriptive Analytics · Uncertainty Quantification · Distributionally Robust Optimisation

1 Introduction

Real world decision problems are seldom deterministic. The perennial operational risk of mathematical programming is the sensitivity of the problem to its parameterisation. Small differences in the parameters governing the objective or the constraints can render solutions highly suboptimal or infeasible. Optimisation under uncertainty is a mature field which has devised a number of tractable approaches that derive robust or expectation optimal solutions, thus managing the uncertainty. Since the true underlying distributions are never known in practice, the most successful approaches exploit available samples of parameters

with statistically valid constructs of uncertainty. When contextual information exists, using what amounts to an unsupervised approach should lead to overly conservative decisions. Contexts of the problem for which existing samples are information-poor may result in overly confident decisions. Problem parameters can be estimated based on available contextual information using supervised learning. This is commonly referred to as predict-then-optimize and is how prescriptive analytics is often performed. This is a form of contextual optimization, but prediction models tend to be overly confident and in their vanilla form do not quantify the certainty of their predictions. Using point parameter estimates preserves the operational risk, which may be exacerbated due to inconsistent out-of-sample performance. Given the ubiquity of predict-then-optimize decision making, improving its reliability and out-of-sample performance will result in tangible impact.

In this work we propose an approach to adapt predictive models with uncertainty quantification (UQ) to a robust optimization setup. We mainly focus on distributionally robust models (DRO), forming the pipeline UQ-DRO. Similar logic can be applied in robust ways for safety-critical situations. Section 2 introduces the concept of robust prediction and optimization. Section 3 presents implemented predictive methods with UQ. The ambiguity sets used with UQ are defined in Sect. 4. Section 5 sets out data-driven algorithms for robustness parameter specification and the objectives of the two pipelines. Section 6 computationally evaluates the approach on a simulated and a real data portfolio optimization problem.

2 Robust Predict-then-Optimize

Prediction models provide an estimate of the conditional expected value of the target. Predictive uncertainty can be decomposed into epistemic, and aleatoric uncertainty. Epistemic uncertainty is a result of a lack of information about the true data generating process (DGP). Aleatoric uncertainty is the underlying stochasticity of the process and is irreducible. Even in the best-case supervised scenario, the persisting aleatoric uncertainty presents an operational risk, thus motivating the use of robust approaches. A well-tuned robust predict-then-optimize approach would provide a meaningful scoring criterion for predictive models used in optimization, reflecting their worst-case outcome. The key driver of decision quality in such a system would be the level of epistemic uncertainty, which would determine the needed level of conservativeness. As such, the predictive model should be highly expressive, trained delicately, as phenomena such as overfitting may increase out-of-sample epistemic uncertainty, and capable of reasonably capturing the uncertainty of its predictions. The prediction task in this case is more difficult as we are typically interested in the values of many parameters, encouraging the use of multiple-output models to better capture interdependence.

Existing approaches for conditional optimization have utilised local nonparametric regression methods such as K-nearest neighbours with DRO [2,3,15] to

provide a conditional sample for a variety uncertain optimisation problems. Building on this [8] propose a method based on distribution trimmings and cast it as a partial mass transportation problem to hedge against the limitations of inferring conditional distributions with limited samples, providing an outer layer of robustness. On the other hand, a growing body of literature has looked at fused approaches wherein the prediction function is optimised with respect to decision loss. Often referred to as Predict-and-Optimise, it involves differentiating across a solver which has been achieved either with surrogate gradients [1, 17] or differentiating the optimality conditions [7, 16] of a potentially relaxed problem. These models implicitly learn how to deal with conditional uncertainty.

In contrast our work uses global supervised learning methods. This is motivated by the idea that global models have the potential to cross-learn about different contexts through shared patterns in the data. This improves the model's ability to infer about contexts which are information-poor, a key case of which are out-of-sample contexts. We do not make any assumptions on the DGP, instead relying on a cross-validation type approach to determine robustness parameters.

3 Predictive Models with Uncertainty Quantification

We propose the use of three predictive approaches which have high expressive power and capacity for UQ. The approaches cover both main directions in UQ, namely ensemble, and Bayesian techniques. The ensemble approach is a deep ensemble (DE) [11], an ensembling technique which treats ensemble members as mixture model components. Constituent models are neural networks designed to predict both a mean vector and a covariance matrix. They are trained using a form of gradient descent to minimise the negative log likelihood given a parametric assumption about the DGP's uncertainty, usually heteroskedastic Gaussian, though Laplacian likelihood may be more appropriate for heavy tails. Denote the available contextual information as $x \in \mathbb{R}^n$, and the model parameters as θ . The model maps $\mathbb{R}^n \mapsto \mathbb{R}^p \times \mathbb{R}^{p \times p}$ or from contextual information x to a mean vector $\mu(x) \in \mathbb{R}^p$ and covariance matrix $\Sigma(x) \in \mathbb{R}^{p \times p}$. The optimisation problem with a Gaussian maximum likelihood objective is:

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \frac{1}{2} (\mathbf{y} - \mu(\mathbf{x}; \theta))^T \Sigma(\mathbf{x}; \theta)^{-1} (\mathbf{y} - \mu(\mathbf{x}; \theta)) + \frac{1}{2} \ln(|\Sigma(\mathbf{x}; \theta)|) \quad (1)$$

Note that the covariance matrix is symmetric, so only $p + \frac{p(p+1)}{2}$ outputs need to be predicted. The structural concern is that Σ should be positive semidefinite (PSD). We follow [14] who encourage a PSD estimate by using an exp activation function for variance terms and a tanh activation function for predicting correlation coefficients from which they construct covariances. The size of the estimated matrix grows quadratically, so this approach is unlikely to scale to large problems. In practice, the exponential activation and subsequent matrix construction can lead to numerical difficulties with the determinant. Since we are interested in the log of the determinant we can reformulate this part of the

loss function into a sum of the log of its eigenvalues. If an odd number of eigenvalues are negative, we clip the value at a small positive ϵ . While this means that the resulting matrix may not be PSD in intermediate steps, it enables more informative gradients and tends to predict PSD matrices after training. This procedure is fully differentiable and was key to stabilising training in addition to standardising variables.

The second approach is Monte-Carlo dropout (MCD), which is an ensemble technique that became popular after a Bayesian analysis showed that it can be cast as approximate inference in deep Gaussian processes [9]. MCD relies on a regularisation technique for deep learning called dropout. Dropout deactivates neurons in the network randomly according to some prior parameterisation (usually Bernoulli), thus limiting the gradient information during that pass to the active units. Dropout regularisation can be thought of as training an implicit ensemble of models within the network, but is deactivated at test time. MCD retains dropout at test time and uses it as an empirical sampling technique to estimate the posterior uncertainty of predictions given contextual information x . This approach should scale better, but tends to be worse at UQ.

Finally, we propose the use of a Bayesian non-parametric regression approach. The most commonly used such approach is a Gaussian Process (GP) which is assumed to be the distribution across functions. Any set of observations about the function value is assumed to have a multivariate Gaussian joint distribution parameterised by a mean, and kernel function which measures the similarity of contextual information. Predictions are made by marginalising the probability distribution of a new point given its contextual information. The choice of kernel is key for modelling (scale and Matern in our case) and kernels are often parametric, thus allowing for some optimisation, typically by optimising the marginal likelihood. Given that we are interested in multi-task prediction, we follow the setup of [5], which models interdependence with a task-similarity kernel.

4 Conditional Ambiguity Sets

We incorporate UQ in various forms of DRO. DRO seeks to obtain a solution that has the least worst expected value across all distributions in an ambiguity set. To exemplify, say the uncertain parameter $\xi \sim \mathcal{P}$ is only in the objective $h(x, \xi)$. The DRO formulation is:

$$\min_{x \in \mathcal{S}} \sup_{\mathcal{P} \in \mathcal{D}} \mathbb{E}_{\mathcal{P}}(h(x, \xi)) \quad (2)$$

It is less conservative than robust optimisation approaches and does not suffer from the optimiser’s curse (overly optimistic out-of-sample) like stochastic programming. The two prevailing ways of defining ambiguity sets are by using moments or disturbance metrics. Moment-based sets are typically convex sets constructed in reference to a stated or estimated moment, usually using conic formulations. Disturbance sets are defined as all distributions within a certain disturbance metric. Even though these problems are semi-infinite, they often admit tractable reformulations by exploiting duality.

We propose the use of ambiguity sets that incorporate the output of UQ. DE and GP quantify their uncertainty with predicted covariances. We employ the approach from [6], which defines ambiguity sets in terms of moment-uncertainty. The ambiguity set $\mathcal{D}(\hat{\mu}, \hat{\Sigma}, \gamma_1, \gamma_2)$ is defined as:

$$\mathcal{D}(\hat{\mu}, \hat{\Sigma}, \gamma_1, \gamma_2) = \left\{ \mathbb{P} \in \mathcal{P}(\mathbb{S}) \left| \begin{array}{l} \xi \sim \mathbb{P} \\ (\mathbb{E}(\xi) - \hat{\mu})^T \hat{\Sigma}^{-1} (\mathbb{E}(\xi) - \hat{\mu}) \leq \gamma_1 \\ \mathbb{E}((\xi - \hat{\mu})^T (\xi - \hat{\mu})) \preceq \gamma_2 \hat{\Sigma} \end{array} \right. \right\}, \quad (3)$$

where $\mathbb{S} \subseteq \mathbb{R}^p$ is the support of the set. The set defines all distributions for which the expected value lies within a scaled ellipsoid uncertainty set centred on the model prediction and shaped by the UQ, and for which the true covariance lies within a PSD cone defined by scaled UQ. This ambiguity set does not assume that the model is correct or that it captures its conditional uncertainty well. It enables us to parametrically define a space of distributions around our model's predictions within which we can guarantee a worst case expectation. Under mild convexity assumptions about the objective function [6], this ambiguity set has a tractable semidefinite programming (SDP) robust counterpart.

For sampling-based UQ we propose the use of ambiguity sets defined by the Wasserstein metric [12]. The ambiguity set is defined as the set of distributions that are within a ball from the empirical reference distribution, which in our case is the n conditionally generated samples $\hat{\mathbb{P}}_n(x_i)$. The ambiguity set is defined as $\mathcal{D}_w(\hat{\mathbb{P}}_n, \phi) = \{\mathbb{Q} \in \mathcal{P}(\mathbb{S}) | d_{W,p}(\mathbb{Q}, \hat{\mathbb{P}}_n) \leq \phi\}$, where $d_{W,p}$ is the p -norm Wasserstein metric $d_{W,p}(\hat{\mathbb{P}}_n, \mathbb{Q}) = \inf_{\Pi} \{\int_{\mathbb{S} \times \mathbb{S}} \|\hat{p} - q\|^p d\Pi(p, q)\}$, and Π denotes the joint distribution of p, q whose marginals are $\hat{\mathbb{P}}_n, \mathbb{Q}$ respectively. Wasserstein ambiguity sets are generally less tractable, but robust counterparts exist in a number of settings. In the context of predict-then-optimize, this allows us to account for sampling error and bias. Posterior sampling from a GP provides a conditional reference distribution based on our structural beliefs about the DGP. MCD is a black box, but provides a more centred form of sampling.

5 Data-driven Robustness Parameter Specification

We see two ways of leveraging the UQ-DRO pipeline depending on how robustness parameters are set. We can either set them to probabilistically cover potential outcomes, or induce limited robustness. The drawback of the former is that robustness tends to come with a cost to performance on average as it is overly conservative. In turn, limited robustness may increase average performance by reducing the sensitivity of decision quality to mild parameter uncertainty.

We aim to achieve coverage with the UQ and uncertain moments pipeline. We achieve this by finding the smallest values γ_1, γ_2 such that the defined uncertainty sets are likely to contain the true moments. We use a holdout set \mathbb{V} as a proxy for the problem. Lower values of these parameters indicate that a model is better tuned for estimating its uncertainty in the context of the optimisation model.

We cast the setting of these parameters as optimisation problems. We want to find the smallest $\hat{\gamma}_1$ such that the (unknown) conditional mean is within the ellipsoidal uncertainty set defined by the predicted mean and variance:

$$\arg \min_{\gamma_1} \{ \gamma_1 | \mathbb{E}[(\xi - \hat{\mu}(x))^T \hat{\Sigma}^{-1}(x)(\xi - \hat{\mu}(x)) \leq \gamma_1 | x] \}. \quad (4)$$

We use $(x, \xi) \in \mathbb{V}$ as a proxy for this problem, by calculating the distance for each point and then picking the median. This is motivated by an assumption that the true conditional distributions are symmetric on average, so realisations are more distant from the predicted mean than the true mean approximately half of the time. Setting γ_2 is slightly more difficult as the associated constraint is defined using a Loewner order. We solve the following SDP problem for each point in the holdout set for each $(x_i, \xi_i) \in \mathbb{V}$:

$$\hat{\gamma}_{2,i} = \arg \min_{\gamma_2} \{ \gamma_2 | \gamma_2 \hat{\Sigma}(x_i) - Z \succeq 0, \gamma_2 \geq 0 \}, \quad (5)$$

where $Z = \frac{1}{|\mathbb{V}|} \sum_{i \in \mathbb{V}} ((\xi_i - \hat{\mu}(x_i))^T (\xi_i - \hat{\mu}(x_i)))$ and set $\hat{\gamma}_2(\alpha_2)$ as the $1 - \alpha_2$ quantile of the obtained $\hat{\gamma}_{2,i}$. The linear matrix inequality constraint in this problem is simple so it should not be a computational bottleneck. We use $\alpha_2 = 0.1$ to encourage a 90% coverage of covariances, but this can be tinkered with.

We use the sampling-Wasserstein pipeline to achieve limited robustness. We want to obtain a reference holdout ϕ and then scale it by multiplying it with some constant $k \leq 1$. We obtain the reference ϕ by computing the p -Wasserstein distance between every distribution realisation ξ_i and the predicted sample $\hat{\mathbb{P}}_n(x_i) = \frac{1}{n} \sum_{j=1}^n \delta_{\hat{\xi}_{i,j}}$ which we treat as a mixture of Dirac delta distributions. Since both are discrete, this is equivalent to calculating the earth mover's distance (EMD) between the two. The p -Wasserstein distance between the two can be cast as a linear optimisation problem:

$$d_{W,p}(\hat{\mathbb{P}}_n(x_i), \delta_{\xi_i}) = \min_T \{ \langle T, M \rangle | T \mathbf{1} = \mathbf{p}_p, T^T \mathbf{1} = \mathbf{p}_\xi \}, \quad (6)$$

where T is the optimal transport matrix, $M \in \mathbb{R}^{n \times 1}$ is the moving cost, which is calculated as the point-wise p -norm between the sample $\hat{\mathbb{P}}_n(x_i)$ and ξ_i , and $\mathbf{p}_p, \mathbf{p}_\xi$ are the discrete densities (in this case equally weighted). Since T is only a vector, the optimisation is trivial: the optimal transport plan is $t_k = \frac{1}{n}$ for all k . The EMD for holdout entry i is therefore $\frac{1}{n} \sum_{j=1}^n |\hat{\xi}_{i,j} - \xi_i|_p^p$ and we set the reference ϕ as the 0.9 quantile of these values. However, the true distribution of ξ is very unlikely to be a discrete point and such a large ambiguity set will likely lead to overly conservative solutions, which is why we scale it down with k .

6 Computational Evaluation and Discussion

Our approach was evaluated on a common prediction and optimisation problem, namely portfolio optimisation. The code is publicly available¹. The uncertain

¹ <https://github.com/EgoPer/Contextual-Robust-Optimisation-with-UQ>.

parameters are the asset returns \mathbf{p} . Asset returns are famously difficult to predict due to a high noise to signal ratio and concept drift. We followed the problem setup from [4, 8], which uses a linear reformulation of CVAR [13] in the objective. The DRO optimisation problem for ϵ -CVAR is:

$$\begin{aligned} \min_{\mathbf{x}, \beta} \inf_{\mathbb{P}} \quad & \mathbb{E}_{\mathbb{P}}[\beta + \frac{1}{\epsilon}(-\mathbf{p}^T \mathbf{x} - \beta)^+ - \lambda \mathbf{p}^T \mathbf{x}] \\ \text{s.t.} \quad & \mathbf{e}^T \mathbf{x} = 1, \mathbf{x} \geq 0 \end{aligned} \quad (7)$$

where λ governs the trade off between tail risk and returns, and $(a)^+ = \max(0, a)$. We set $\epsilon = 0.1$ (expected value of the 10% worst cases), and λ at 1. We use 25 samples for each conditional Wasserstein approach, set $k = 0.1$ on the simulated problem, and $k = 0.02$ on the real data problem.

6.1 Simulated Problem

The simulated version of this problem is based on a problem used in two existing papers [4, 8] about DRO with side/contextual information, but we introduce significant non-linearity and heteroskedasticity. Three independent inputs are simulated as standard normal variables $x_{1,2,3} \sim \mathcal{N}(0, 1)$. The conditional joint distribution of the simulated returns is $\mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$, where $\mu(\mathbf{x}) = \bar{\mu} + y(\mathbf{x})^2$, $\Sigma(\mathbf{x}) = [(\frac{1}{5}\tanh(x_1) + 1) \cdot \bar{\Sigma}^{\frac{1}{2}}]^2$ ($\bar{\mu}, \bar{\Sigma}$ as in [4]).

We generate five datasets at five training set sizes (20% holdout) and train models five times due to the stochastic nature of their training. The test set is the same across experiments and is deliberately generated out-of-sample (100 samples of $x_{1,3} \sim \mathcal{N}(2, 1), x_2 \sim \mathcal{N}(-2, 1)$, same DGP). Given that we have access to the true DGP, we can approximately evaluate the performance of each solution (in our case using a 10^4 sample Monte Carlo simulation). We construct a deterministic equivalent, which gives us the true optimal value, allowing us to measure regret. We run three unconditioned models that derive their uncertainty inputs from the whole training set, an uncertain moments (UM) model, a Wasserstein (WASS) model, and a sample average approximation (SAA) model. We run a conditional SAA model using the moment outputs of DE to sample a normal distribution. We run five of our contextual models: DE-UM, GP-UM, DE-Wasserstein (DE-WASS), GP-WASS, and MCD-WASS.

Figure 1 displays boxplots of the mean out-of-sample regret for each approach. The DE-WASS approach dominates across data sizes improving in performance with a growing size of the dataset, outdoing a robustly performing unconditional Wasserstein. MCD-WASS is close in performance to the unconditional case, while the GP version lags slightly. The prediction models were trained in an out-of-the-box manner, so it is conceivable that they would outperform using hyperparameter optimisation or in richer data environments. The conditional UM methods outperformed the unconditional case, though the gap

² The outputs of $y(\mathbf{x})$ are defined as:

$$\begin{aligned} y_1 &= 30 \tanh(x_2 \exp(\frac{1}{2}x_1 - 2)), y_2 = 50 \tanh(x_1) \sin(3x_3), y_3 = 10 \ln(|x_1 x_2 x_3|), \\ y_4 &= \sin(x_2) + x_1^2 - x_1 x_2, y_5 = 20(\sin(x_1) + \sin(\frac{x_2}{10x_3})), y_6 = y_1 - y_2. \end{aligned}$$

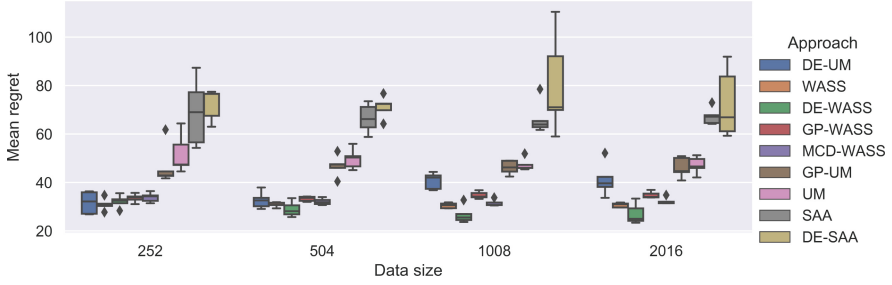


Fig. 1. Mean performance from simulated experiments.

between DE and the unconditional UM got smaller with an increase in the size of data, possibly reflecting overfitting as the models were trained for the same number of epochs in each case. The results support the use of UQ for constructing contextual DRO approaches.

For the real problem we use a slightly reduced version of the dataset from [10], which contains returns on five large US indices alongside 106 contextual covariates such as technical and economic indicators. We cannot directly evaluate the objective function in 7 so we approximate the CVAR and the expected return using a two year testing set training each model five times.

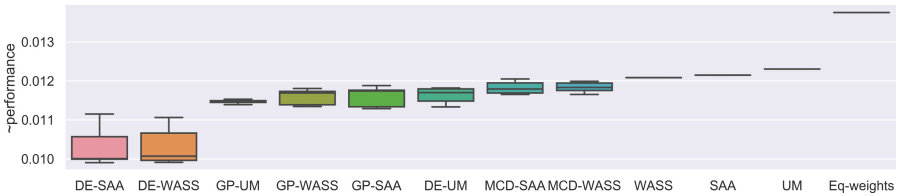


Fig. 2. Approximate performance from real-data experiment.

Figure 2 presents the approximate performance of these methods. The contextual approaches outperform all unconditioned approaches. The Wasserstein approaches are very competitive with non-robust contextual SAA equivalents and are less sensitive to model training, illustrating the positive trade-off of limited robustness. Notably, the DE-WASS approach is best performing in both experiments.

The pipelines offer an effective way of introducing robustness to prediction and optimisation problems by leveraging established methods for UQ. They can also be used as a means of achieving contextual DRO, though analysis is needed to establish the desired convergence properties that established DRO techniques have. We see much potential for refining these pipelines with regularisation, predictive model architecture, contextual setting of robustness parameters, and end-to-end learning.

References

1. Amos, B., Kolter, J.Z.: Optnet: differentiable optimization as a layer in neural networks. In: International Conference on Machine Learning, pp. 136–145. PMLR (2017)
2. Bertsimas, D., Kallus, N.: From predictive to prescriptive analytics. *Manag. Sci.* **66**(3), 1025–1044 (2020)
3. Bertsimas, D., McCord, C., Sturt, B.: Dynamic optimization with side information. *Eur. J. Oper. Res.* **304**(2), 634–651 (2023)
4. Bertsimas, D., Van Parys, B.: Bootstrap robust prescriptive analytics (2017). <https://arxiv.org/abs/1711.09974v1>
5. Bonilla, E.V., Chai, K., Williams, C.: Multi-task gaussian process prediction. *Adv. Neural Inf. Process. Syst.* **20** (2007)
6. Delage, E., Ye, Y.: Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* **58**(3), 595–612 (2010)
7. Elmachtoub, A.N., Grigas, P.: Smart predict, then optimize. *Manag. Sci.* **68**(1), 9–26 (2022)
8. Esteban-Pérez, A., Morales, J.M.: Distributionally robust stochastic programs with side information based on trimmings. *Math. Program.* 1–37 (2021). <https://doi.org/10.1007/s10107-021-01724-0>
9. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: International Conference on Machine Learning, pp. 1050–1059. PMLR (2016)
10. Hoseinzade, E., Haratizadeh, S.: CNNpred: CNN-based stock market prediction using a diverse set of variables. *Expert Syst. Appl.* **129**, 273–285 (2019)
11. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. *Adv. Neural Inf. Process. Syst.* **30** (2017)
12. Mohajerin Esfahani, P., Kuhn, D.: Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. *Math. Program.* **171**(1), 115–166 (2018)
13. Rockafellar, R.T., Uryasev, S.: Conditional value-at-risk for general loss distributions. *J. Bank. Financ.* **26**(7), 1443–1471 (2002)
14. Russell, R.L., Reale, C.: Multivariate uncertainty in deep learning. *IEEE Trans. Neural Netw. Learn. Syst.* **33**, 7937–7943 (2021)
15. Srivastava, P.R., Wang, Y., Hanasusanto, G.A., Ho, C.P.: On data-driven prescriptive analytics with side information: a regularized nadaraya-watson approach (2021). <https://arxiv.org/abs/2110.04855>
16. Vlastelica, M., Paulus, A., Musil, V., Martius, G., Rolínek, M.: Differentiation of blackbox combinatorial solvers. arXiv preprint [arXiv:1912.02175](https://arxiv.org/abs/1912.02175) (2019)
17. Wilder, B., Dilkina, B., Tambe, M.: Melding the data-decisions pipeline: decision-focused learning for combinatorial optimization. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 1658–1665 (2019)