Optimization under Severe Uncertainty: a Generalized Minimax Regret Approach for Problems with Linear Objectives

Tuan-Anh Vu[1], Sohaib Afifi[1], Éric Lefèvre[1], and Frédéric Pichon[1]

Univ. Artois, UR 3926, Laboratoire de Genie Informatique et d'Automatique de l'Artois (LGI2A), F-62400 Bethune, France
{tanh.vu, firstname.lastname}@univ-artois.fr

Abstract. We study a general optimization problem with an uncertain linear objective. We address the uncertainty using two models: belief functions and, more generally, capacities. In the former model, we use the generalized minimax regret criterion introduced by Yager, while in the latter one, we extend this criterion, to find optimal solutions. This paper identifies some tractable cases for the resulting problem. Furthermore, when focal sets of the considered belief functions are Cartesian products of intervals, we develop a 2-approximation method that mirrors the well-known midpoint scenario method used for minimax regret optimization problems with interval data.

Keywords: Minimax regret · Belief functions · Capacities · Linear programming.

1 Introduction

Uncertainty is ubiquitous in optimization problems, leading to numerous frameworks for handling it. This paper revisits the famous minimax regret criterion within robust optimization. In essence, this criterion arises from two key motivations in decision-making under uncertainty: (i) the common human tendency to regret choices especially if a better option is discovered later; and (ii) the desire for an option with the best worst-case performance. The minimax regret criterion, widely studied for optimization problems with uncertain coefficients in the objectives [8], assumes a classical setting where only a so-called scenario set of possible realizations of coefficients, is available. Under this limited information, the criterion aims to find a solution that minimizes the maximum regret across all scenarios.

However, some partial information is usually at hand in real-life situations. For example, knowing the scenario set allows us to consult experts who can assess the likelihood of each scenario occurring. In such cases, refining the minimax regret criterion to account for partial information becomes necessary to better reflect real-world situations. Interestingly, under *evidential* uncertainty, i.e., when the uncertainty is modeled by belief functions [10], a notion of generalized minimax regret has already been introduced by Yager [13]. Similarly, a recent work by Adam and Destercke [1] discussed a related notion within the possibilistic framework.

Following our recent paper on a general optimization problem with uncertain linear objective [11], we also study the same problem. Unlike [11] which considered five other criteria, this paper uses generalized minimax regret criteria to find optimal solutions. Furthermore, we address situations where the uncertainty about the objective coefficients is *severe*, i.e., we cannot identify a single probability measure to represent it. For this reason, we use more general frameworks, namely belief functions and capacities, to model this uncertainty.

The paper is organized as follows. Section 2 presents some elements about belief functions. In Section 3, we incorporate Yager's criterion [13] to the considered problem. For Yager's criterion, two types of belief functions, where (i) their frames are finite and (ii) their frames are infinite but their focal sets take a special form, are addressed in Section 4 and Section 5, respectively. We then extend Yager's criterion to a more general setting where uncertainty is modeled by capacities or lower probabilities [6,4] in Section 6. The paper ends with a conclusion.

2 Belief function theory

Let Ω be the set of all possible values of a variable of interest ω. In this paper, we assume that Ω is a closed subset of $\mathbb{R}^n$. In belief function theory [10], adapting the presentation of [12], partial information about the true (unknown) value of ω is given by a mapping $m : \mathcal{C} \mapsto [0,1]$ called mass function, where $\mathcal{C}$ is a finite collection of closed subsets of Ω, such that $\sum_{A\in\mathcal{C}} m(A) = 1$ and $m(\emptyset) = 0$. If Ω is finite, we usually take $\mathcal{C} = 2^\Omega$.

Mass $m(A)$ quantifies the amount of belief allocated to the fact of knowing only that $\omega \in A$. A focal set of m is a subset $A \subseteq \Omega$ such that $m(A) > 0$. Let $\mathcal{F} = \{F_1, \dots, F_K\}$ be the set of all focal sets of m. The mass function m induces a belief function Bel defined on $\mathcal{B}(\Omega)$ the Borel subsets of Ω where $Bel(A) = \sum_{B\in\mathcal{F}:B\subseteq A} m(B)$.

3 Problem Formulation

In this paper, we focus on a general problem with a linear objective:

$$\{\min c^T x : x \in \mathcal{X} \subseteq \mathbb{R}^n_{\geq 0}\} \tag{P}$$

where $\mathcal{X}$ is a compact set and $c \in \mathbb{R}^n$ is the coefficient vector of the objective. (P) is a linear programming problem if $\mathcal{X} = \{x \in \mathbb{R}^n_{\geq 0} : Mx \leq b\}$ where M is a $q \times n$ matrix and b is a q-vector. If $\mathcal{X} \subseteq \{0,1\}^n$, then (P) is a combinatorial problem. In this paper, we assume that the coefficient vector c is uncertain and let $\Omega \subseteq \mathbb{R}^n$ be the set of possible values of c. Each $c \in \Omega$ is then called a scenario.

3.1 The minimax regret criterion

The regret $R(x,c)$ of a solution x under a scenario $c \in \Omega$ is defined as $R(x,c) := c^T x - val^*(c)$ where $val^*(c) := \min_{x\in\mathcal{X}} c^T x$ the optimal value of (P) under c. The maximum regret $R(x)$ of x is defined as $R(x) := \max_{c\in\Omega} R(x,c)$: it represents the regret of x in the worst case scenario in Ω. The goal is to find a solution x having minimum $R(x)$ by solving the problem:

$$\min_{x\in\mathcal{X}} R(x) = \min_{x\in\mathcal{X}} \max_{c\in\Omega}(c^T x - val^*(c)) \tag{MR}$$

3.2 The generalized minimax regret criterion under evidential uncertainty

If some partial knowledge about c is given by a mass function m, we generalize the minimax regret criterion as follows [13]. For each focal set F of m, the maximum regret of x is $R^F(x) := \max_{c\in F}(c^T x - val^*(c))$. The expected maximal regret of x with respect to m is then defined as

$$\overline{R}(x) := \sum_{r=1}^{K} m(F_r) R^{F_r}(x). \tag{1}$$

In this paper, we focus on addressing Problem (GMR):

$$\min_{x\in\mathcal{X}} \overline{R}(x) = \min_{x\in\mathcal{X}} \sum_{r=1}^{K} m(F) \max_{c\in F_r}(c^T x - val^*(c)). \tag{GMR}$$

Note that if m is a vacuous mass function, i.e., Ω is the only focal set of m, then (GMR) becomes (MR).

Remark 1. The information about the true scenario can be given by a *possibility distribution* $\pi : \Omega \to [0,1]$ with values of π representing possibility degrees of elements in Ω, among which there exists a c such that $\pi(c) = 1$. This representation of uncertainty is practical, as π can, for instance, be constructed from expert assessments. Assume that $1 = \alpha_1 > \ldots > \alpha_K > \alpha_{K+1} = 0$ are the distinct values of π. For each α_i, the associated α_i cut of π is defined as: $F_{\alpha_i} = \{c \in \Omega : \pi(c) \geq \alpha_i\}$. Obviously, $F_{\alpha_1} \subset \ldots \subset F_{\alpha_K}$. If we construct a mass function on Ω with focal sets F_{α_i} and $m(F_{\alpha_i}) = \alpha_i - \alpha_{i+1}$ $\forall i \in \{1, \ldots K\}$, we return to the version of generalized minimax regret criterion under possibilistic framework, introduced in [1].

4 When Ω is finite

In this case, we have a mass function m on a finite set of l elements $\Omega = \{c^1, \ldots, c^l\} \subset \mathbb{R}^n$. For combinatorial optimization problems, the intractability of (MR) has been well-documented, see e.g., [8], thereby implying the intractability of (GMR) for such problems as well. The main result in this section, therefore, concerns a case where (GMR) is tractable.

Proposition 1. *Assume that (P) is a linear programming problem. Then (GMR) can be solved efficiently provided $|\mathcal{F}|$ is not large. In particular, if $|\mathcal{F}|$ is polynomially bounded in l then (GMR) can be solved in polynomial time.*

Proof. We reformulate (GMR) as:

$$\begin{aligned} &\min \sum_{F \in \mathcal{F}} m(F) z_F \\ &z_F \geq c^T x - val^*(c) \; \forall F \in \mathcal{F}, c \in F \\ &Mx \leq b, \; x \in \mathbb{R}^n_{\geq 0}. \end{aligned} \tag{2}$$

Note that (2) is a linear programming problem. Moreover, for each $c \in \Omega$, $val^*(c) = \min\{c^T x : x \in \mathbb{R}^n_{\geq 0}, Mx \leq b\}$ can be computed efficiently by standard linear programming solvers. Therefore,(2) can be solved efficiently provided the number of focal sets is not large and be solvable in polynomial time if $|\mathcal{F}|$ is polynomially bounded. □

5 When Ω is infinite and focal sets of m are Cartesian products of intervals

In this section, we assume that each focal set F_r of m is a Cartesian product of intervals, i.e.,

$$F_r = \times_1^n [l_i^r, u_i^r] \; \forall r.$$

When m has a unique focal set of such type, we get back to the famous interval uncertainty representation in robust optimization. We remark that under interval representation, the classical minimax regret Problem (MR) is intractable in both cases where (P) is a combinatorial or linear programming problem [8]. Fortunately, a well-known heuristic exists to obtain a 2-approximation algorithm for (MR): it uses an optimal solution of (P) under the so-called midpoint scenario [5,8]. The goal here is to adapt this heuristic for our considered uncertainty representation, for which we follow the approach in [5]. We denote

$$\overline{u}_i := \sum_{r=1}^{K} m(F_r) u_i^r \text{ and } \overline{l}_i := \sum_{r=1}^{K} m(F_r) l_i^r. \tag{3}$$

Proposition 2. *Let $\overline{c}$ be a vector in $\mathbb{R}^n$ such that $\overline{c}_i = \frac{\overline{u}_i + \overline{l}_i}{2}$ and y be an optimal solution of (P) under $\overline{c}$, i.e., $\overline{c}^T y = \min_{x \in \mathcal{X}} \overline{c}^T x$. Let x^* be any optimal solution of (GMR). Then $\overline{R}(y) \leq 2\overline{R}(x^*)$.*

To prove Proposition 2, we need some preliminary observations. First, notice that for any F_r, $R^{F_r}(x^*) = \max_{x \in \mathcal{X}} \max_{c \in F_r} c^T (x^* - x)$, and thus

$$R^{F_r}(x^*) \geq \max_{c \in F_r} c^T (x^* - y). \tag{4}$$

Note also that $\max_{c\in F_r} c^T(x^*-y) = \sum_{i:x_i^*>y_i} u_i^r(x_i^*-y_i) - \sum_{i:x_i^*<y_i} l_i^r(y_i-x_i^*)$. Therefore, $R^{F_r}(x^*) \geq \sum_{i:x_i^*>y_i} u_i^r(x_i^*-y_i) - \sum_{i:x_i^*<y_i} l_i^r(y_i-x_i^*)$. Using (1) and (3),

$$\overline{R}(x^*) \geq \sum_{i:x_i^*>y_i} \overline{u}_i(x_i^*-y_i) - \sum_{i:x_i^*<y_i} \overline{l}_i(y_i-x_i^*) \tag{5}$$

In the subsequent, we use the notation $\delta(y-x^*, F_r) := \max_{c\in F_r} c^T(y-x^*)$. Referring to [5, Property 2.2], we have that $R^{F_r}(y) \leq R^{F_r}(x^*) + \delta(y-x^*, F_r)\ \forall r$. From (1),

$$\overline{R}(y) \leq \overline{R}(x^*) + \sum_{r=1}^{K} m(F_r)\delta(y-x^*, F_r). \tag{6}$$

We are ready to prove Proposition 2.

Proof (Proof of Proposition 2). By the optimality of y, $\sum_{i=1}^n (\overline{u}_i + \overline{l}_i)x_i^* \geq \sum_{i=1}^n (\overline{u}_i + \overline{l}_i)y_i$. Equivalently, $\sum_{i=1}^n \overline{u}_i(x_i^*-y_i) \geq \sum_{i=1}^n \overline{l}_i(y_i-x_i^*)$. It follows that

$$\sum_{i:x_i^*>y_i} \overline{u}_i(x_i^*-y_i) - \sum_{i:x_i^*<y_i} \overline{u}_i(y_i-x_i^*) \geq \sum_{i:x_i^*<y_i} \overline{l}_i(y_i-x_i^*) - \sum_{i:x_i^*>y_i} \overline{l}_i(x_i^*-y_i) \tag{7}$$

$$\sum_{i:x_i^*>y_i} \overline{u}_i(x_i^*-y_i) - \sum_{i:x_i^*<y_i} \overline{l}_i(y_i-x_i^*) \geq \sum_{i:x_i^*<y_i} \overline{u}_i(y_i-x_i^*) - \sum_{i:x_i^*>y_i} \overline{l}_i(x_i^*-y_i). \tag{8}$$

It can be easily checked that the right hand side of (8) equals $\sum_{r=1}^K m(F_r)\delta(y-x^*, F_r)$. Hence, it follows from (5) that $\overline{R}(x^*) \geq \sum_{r=1}^K m(F_r)\delta(y-x^*, F_r)$. Finally, Proposition 2 is true because of (6).

□

6 Beyond belief functions

We still consider the case of finite $\Omega = \{c^1, \ldots, c^l\}$ as in Section 4. However, in this context, the partial knowledge about the true coefficient vectors is modeled by non-additive measures, namely capacities which are more general than belief functions. We quickly summarize some basics elements adapted from [6].

A capacity on Ω is a set function $\mu : 2^\Omega \to [0,1]$ such that $\mu(\Omega) = 1,\ \mu(\emptyset) = 0$ and if $A \subseteq B,\ \mu(A) \leq \mu(B)$. Note that μ is a probability measure if it is additive, i.e., $\mu(A\cup B) = \mu(A) + \mu(B)\ \forall A, B \in 2^\Omega$ with $A\cap B = \emptyset$. Furthermore, μ is a 2-monotone capacity if $\mu(A\cup B) + \mu(A\cap B) \geq \mu(A) + \mu(B)\ \forall A, B \subseteq \Omega$. A belief function, also known as a *complete monotonicity* capacity, is a special 2-monotone capacity [6,10].

Remark 2. In combinatorial optimizations [7], a 2-monotone capacity μ is called a supermodular set function while its dual $\bar{\mu}$, defined as $\bar{\mu}(A) = 1-\mu(\Omega\backslash A)\ \forall A \subseteq \Omega$, is called submodular.

In imprecise probability [4], μ is usually called a lower probability where values of μ are interpreted as lower bounds of values of the true (yet unknown) probability measure P^* on Ω. Under this view, the so-called *credal set* of μ consisting of all compatible probability measures with μ on Ω, is defined as $\mathcal{M}(\mu) := \{P : P(A) \geq \mu(A) \ \forall A \subseteq \Omega\}$. We will henceforth view any element in $\mathcal{M}(\mu)$ as a vector $p \in [0,1]^l$, and thus $\mathcal{M}(\mu)$ is a polytope:

$$\mathcal{M}(\mu) = \{p \in [0,1]^l : \sum_{i \in A} p_i \geq \mu(\{c^i : i \in A\}) \ \forall A \subseteq \{1, \ldots, l\}, \ \sum_{j=1}^{l} p_i = 1\}. \quad (9)$$

Because explicitly listing all 2^l values of $\mu(A)$ is intractable, we use a typical assumption from optimizations [7, Chapter 10].

Assumption 1 *We have access to an evaluation oracle that returns $\mu(A)$ for each query $A \subseteq \Omega$.*

We proceed to extend the generalized minimax regret criterion, discussed in Section 3.2, to incorporate the notion of capacity as follows. The expected regret of a solution $x \in \mathcal{X}$ with respect to a probability measure $p \in \mathcal{M}(\mu)$ is $\sum_{i=1}^{l} p_i R(x, c^i)$. Since the only available information is that the true probability measure lies in $\mathcal{M}(\mu)$, a reasonable approach is to seek a solution that minimizes the worst-case of expected regret among all compatible probabilities. In other words, we need to solve:

$$\min_{x \in \mathcal{X}} \max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i R(x, c^i) = \min_{x \in \mathcal{X}} \max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i) \right). \quad \text{(CGMR)}$$

If μ is a belief function on Ω and m is its associated mass function (see Section 2), a well-known result [6] states that

$$\max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i) \right) = \sum_{r=1}^{K} m(F_r) \max_{c \in F_r} (c^T x - val^*(c)). \quad (10)$$

Thus, in this case, CGMR reverts to (GMR).

Remark 3. If μ is 2-monotone, it is well-known that a p^* that maximizes the left hand side of (10) can be efficiently computed by using only l accesses to the oracle, as follows [6,7]. Reindex elements of Ω such that $(c^1)^T x - val^*(c^1) \geq \ldots \geq (c^l)^T x - val^*(c^l)$ and let $A_j = \left\{c^j, \ldots, c^l\right\} \ \forall j \in \{1, \ldots l\}$ and $A_{l+1} = \emptyset$. Finally, take $p_j^* = \mu(A_j) - \mu(A_{j+1}) \ \forall j \in \{1, \ldots l\}$. Moreover, such p^* is also an extreme point of $\mathcal{M}(\mu)$.

We now show that under the computational model described in Assumption 1, CGMR is tractable if (P) is a linear programming problem. Let $Ext(\mu)$ be the set of extreme points of $\mathcal{M}(\mu)$. Note that $Ext(\mu)$ is finite but can be very large, i.e., $|Ext(\mu)|$ is exponential in l. We first observe that

$$\max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i) \right) = \max_{p \in Ext(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i) \right). \quad (11)$$

Proposition 3. *If (P) is a linear programming problem and assume that μ is 2-monotone. Then (CGMR) can be solved in polynomial time.*

Proof. Using (11), we reformulate (CGMR) as:

$$\min t \tag{12}$$

$$t \geq \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i)\right) \;\; \forall p \in Ext(\mu) \tag{13}$$

$$Mx \leq b, \;\; x \in \mathbb{R}^n_{\geq 0}. \tag{14}$$

Problem (12-14) is a linear programming problem but it has a vast number of constraints due to (13). Because (P) is a linear programming problem, $val^*(c^i)$ is computed in polynomial time. To demonstrate the polynomial solvability of (12-14), we employ the celebrated ellipsoid method [7]. According to this method, we need to show that the separation problem associated with (12-14) can be solved in polynomial time: either confirms if given a point $(x^0, t^0) \in \mathbb{R}^n$ satisfies all the constraints (13-14) or return a constraint that it violates. Checking if (x^0, t^0) satisfies (14) can be easily done in polynomial time. Furthermore, checking if (x^0, t^0) satisfies (13) amounts to testing if $t^0 \geq \max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x^0 - val^*(c^i)\right)$, which can be done in polynomial time because of Remark 3. We conclude that the separation problem, and thus (12-14) is polynomial solvable. □

Remark 4. Because of the popularity of the minimax regret criterion, similar forms to (CGMR) have already appeared in the literature of optimization under distributional uncertainty, to cite only a few [2,3]. However, to the best of our knowledge, Proposition 3 is new.

While the ellipsoid algorithm is theoretically polynomial, it is known to be slow in practice [7]. Consequently, alternative approaches are necessary. Note that the function $f(x) := \max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i)\right)$ is convex in x (it is a pointwise maximum of affine functions). Therefore, (CGMR) is a convex optimization problem. A standard approach to solving it is using subgradient methods [9], where a subgradient of f is required at each iteration. Recall that a subgradient of f at x is a vector η such that $f(y) \geq f(x) + \eta^T (y - x) \; \forall y$. The next result follows from standard calculations in convex analysis. For the completeness, we include a proof.

Proposition 4. *For any x, let $p^* \in \text{argmax}_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i)\right)$. Then $\eta := \sum_{i=1}^{l} p_i^* c^i$ is a subgradient of f.*

Proof. For any y, $f(y) = \max_{p \in \mathcal{M}(\mu)} \sum_{i=1}^{l} p_i \left((c^i)^T x - val^*(c^i) + (c^i)^T (y - x)\right)$. By the optimality of p^*,

$$f(y) \geq \sum_{i=1}^{l} p_i^* \left((c^i)^T x - val^*(c^i)\right) + \sum_{i=1}^{l} (p_i^* c^i)^T (y - x) = f(x) + \eta^T (y - x).$$

□

Thanks to Remark 3 and Proposition 4, in case of 2-monotone capacities, a subgradient of f can be computed efficiently.

7 Conclusion

In this paper, we have used the generalized minimax regret criteria for optimization problems with uncertain objectives, where the uncertainty is modeled by belief functions and, more generally, capacities. We have identified some tractable cases and developed a 2-approximation method when focal sets of the considered belief functions are Catersian products of intervals. Future work includes applying subgradient methods to problem (CGMR) for linear programming problems or investigating problems (GMR) and (CGMR) for some practical combinatorial problems.

References

1. Adam, L., Destercke, S.: Possibilistic preference elicitation by minimax regret. In: Uncertainty in artificial intelligence. pp. 718–727. PMLR (2021)
2. Agarwal, A., Zhang, T.: Minimax regret optimization for robust machine learning under distribution shift. In: Conference on Learning Theory. pp. 2704–2729. PMLR (2022)
3. Al Taha, F., Yan, S., Bitar, E.: A distributionally robust approach to regret optimal control using the wasserstein distance. In: 2023 62nd IEEE Conference on Decision and Control (CDC). pp. 2768–2775. IEEE (2023)
4. Augustin, T., Coolen, F.P., De Cooman, G., Troffaes, M.C.: Introduction to imprecise probabilities. John Wiley & Sons (2014)
5. Conde, E.: A 2-approximation for minmax regret problems via a mid-point scenario optimal solution. Oper. Res. Lett. **38**(4), 326–327 (2010)
6. Grabisch, M.: Set functions, games and capacities in decision making, vol. 46. Springer (2016)
7. Grötschel, M., Lovász, L., Schrijver, A.: Geometric algorithms and combinatorial optimization. Springer (1993)
8. Kasperski, A., Zieliński, P.: Robust discrete optimization under discrete and interval uncertainty: A survey. Robustness analysis in decision aiding, optimization, and analytics pp. 113–143 (2016)
9. Nesterov, Y.: Lectures on convex optimization, vol. 137. Springer (2018)
10. Shafer, G.: A mathematical theory of evidence. Princeton university press (1976)
11. Vu, T.A., Afifi, S., Lefèvre, E., Pichon, F.: Optimization problems with evidential linear objective. Int. J. Approx. Reason. **161**, 108987 (2023)
12. Wasserman, L.A.: Belief functions and statistical inference. Can. J. Stat. **18**(3), 183–196 (1990)
13. Yager, R.R.: Decision making using minimization of regret. Int. J. Approx. Reason. **36**(2), 109–128 (2004)