



Competitive Search in the Line and the Star with Predictions

Spyros Angelopoulos

► To cite this version:

Spyros Angelopoulos. Competitive Search in the Line and the Star with Predictions. 48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023), Aug 2023, Bordeaux, France. pp.12:1–12:15, 10.4230/LIPIcs.MFCS.2023.12 . hal-04248709

HAL Id: hal-04248709

<https://hal.science/hal-04248709>

Submitted on 18 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Competitive Search in the Line and the Star with Predictions

Spyros Angelopoulos 

CNRS and LIP6, Sorbonne University, Paris, France

Abstract

We study the classic problem of searching for a hidden target in the line and the m -ray star, in a setting in which the searcher has some *prediction* on the hider's position. We first focus on the main metric for comparing search strategies under predictions; namely, we give positive and negative results on the *consistency-robustness* tradeoff, where the performance of the strategy is evaluated at extreme situations in which the prediction is either error-free, or adversarially generated, respectively. For the line, we show tight bounds concerning this tradeoff, under the *untrusted advice* model, in which the prediction is in the form of a k -bit string which encodes the responses to k binary queries. For the star, we give tight, and near-tight tradeoffs in the *positional* and the *directional* models, in which the prediction is related to the position of the target within the star, and to the ray on which the target hides, respectively. Last, for all three prediction models, we show how to generalize our study to a setting in which the performance of the strategy is evaluated as a function of the searcher's desired tolerance to prediction errors, both in terms of positive and inapproximability results.

2012 ACM Subject Classification Theory of computation → Online algorithms; Theory of computation → Theory and algorithms for application domains

Keywords and phrases Search problems, line and star search, competitive ratio, predictions, consistency and robustness

Digital Object Identifier 10.4230/LIPIcs.MFCS.2023.12

Funding This work was partially funded by the grant ANR-19-CE48-0016 from the French National Research Agency (ANR).

Acknowledgements I am thankful to Shahin Kamali for several helpful discussions.

1 Introduction

Searching for a hidden target is one of the original disciplines within the field of Operations Research, but also a topic of significant study in Computer Science, both from the point of view of theoretical analysis and applications. This class of problems typically involves a mobile *searcher* that must locate an immobile *target* (often called *hider*) which hides in some unknown point of the search *environment*. Search problems provide natural formulations of real-life applications such as search-and-rescue missions [41], de-mining operations [2], and robot-based exploration [27].

Among the most well-studied search problems is *searching on the line*, in which the environment is the unbounded line, and its generalization, the *m -ray search*, or *star search* problem. In the *m -ray search* problem, the environment consists of m unbounded and concurrent rays, with a common point O , which is called the *origin*. Starting from O , the searcher must locate the target by following a *strategy*, defined as an infinite sequence of the form $(x_i, u_i)_i$, where $x_i \in \mathbb{R}^+$ and $u_i \in \{0, \dots, m-1\}$, and with the following semantics: in iteration i , the searcher starts from O , traverses the ray u_i up to distance x_i from O , then returns back to O , before continuing with iteration $i+1$, until the target is eventually located. Note that for $m=2$, the star environment reduces to the infinite line.



© Spyros Angelopoulos;

licensed under Creative Commons License CC-BY 4.0

48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023).

Editors: Jérôme Leroux, Sylvain Lombardy, and David Peleg; Article No. 12; pp. 12:1–12:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Since the search environment is unbounded, the standard framework for evaluating the performance of a search strategy is *competitive analysis*, first introduced in [11]. Given a target t hiding at some unknown point of the star, define $d(t)$ as the distance of t from O , and $d(X, t)$ as the distance covered (or *cost* incurred) by a searcher that follows X , until t is located (i.e., the first time the searcher is reached, assuming a unit-speed searcher). The *competitive ratio* of X is formally defined as

$$\text{Cr}(X) = \sup_t \frac{d(X, t)}{d(t)}. \quad (1)$$

Searching on the line has a long history of study, going back to the work of Bellman [12] and Beck [10]. Beck and Newman [11] were the first to show that an optimal competitive ratio equal to 9 can be obtained by a simple doubling strategy, i.e., a strategy of the form $x_i = 2^i$. The m -ray search problem was first studied in the seminal works of Gal [24, 25] and independently by Baeza-Yates *et al.* [8]. Both problems have been extended in a variety of settings and generalizations related to TCS, AI and OR since the 1960s, due to their useful abstraction of resource allocation under uncertainty. For instance, linear and ray searching have connections to the design of *interruptible systems* in AI [4, 13], the design of *hybrid* algorithms [32], and pipelined *filter ordering* in databases [18]. They are also involved in the analysis of strategies for more complex search problems, such as *spiral* search on the plane [38]. There are numerous studies on variants of linear and star search; see, e.g., [31, 36, 29, 44, 33, 16, 20, 5, 19, 37, 15, 34, 46, 6, 43, 51, 36, 17] as well as the book [1] for a game-theoretic perspective of these problems. Note that the above are only some representative works, and that the problems have been studied under several other variants.

1.1 Searching with predictions

In this work, we study the power and limitations of search strategies with *predictions*, in which the searcher aims to improve the competitive ratio of its strategy by leveraging some inherently imperfect information on the target. This follows a very active line of research in online computation and algorithms with incomplete information, that was initiated with the works [45] and [49]. A very large number of problems have been studied under this model (see, e.g., the survey [47] and the online collection [42]).

In regards to the search problems we study, the nature of the prediction may vary according to the application at hand. We are interested in the following models, which were introduced in [3] in the context of linear search.

- (a) The prediction is a k -bit string. Here, the search is enhanced with a k -bit string that encodes some information on the target; alternatively, we may think of the prediction string as responses to binary *queries* given by k experts. For example, a single bit can provide a (potentially erroneous) response to queries such as “Is the target at distance at most d from O ”, or “Is the target on an even-indexed ray?”. This is a powerful model that generalizes the concept of *advice complexity* so as to allow for advice that may be erroneous. Note that search and exploration problems have been studied extensively under the standard advice complexity model (see, e.g., [21, 23, 28, 35, 48]), however all such studies rely crucially on advice that is error-free. Moreover, unlike works in which each query is *noisy* [14], i.e., the query responses are erroneous with some known probability, in our setting we do not rely on any probabilistic assumptions in regards to the quality of the advice.

- (b) The prediction is *directional*. Here, the prediction is an index in $\{0, \dots, m-1\}$ which describes the ray on which the target lies. This is a natural prediction can be useful, for example, in a search-and-rescue application, in which there is a hint about the direction a missing person may have taken when last seen.
- (c) The prediction is *positional*. Here, the prediction describes the position of the target within the environment, namely it is of the form (d, u) , where d corresponds to the predicted distance from the origin, and u corresponds to the predicted ray on which the target hides. This is, likewise, a very natural prediction (e.g., in a search-and-rescue mission, it provides a hint about the last reported whereabouts of the missing person).

We establish two objectives towards the evaluation of search strategies with predictions. The first objective is to find strategies of optimal, or near-optimal tradeoff between their *consistency* (namely, the competitive ratio assuming error-free prediction) and their *robustness* (namely, the competitive ratio assuming adversarially generated predictions). This is one of the standard methods of analyzing algorithms with predictions, since it establishes strong guarantees on worst-case (extreme) situations with respect to the quality of the prediction; see, e.g., [52, 40, 39, 3, 7] for applications to other online problems, and settings of incomplete information, more generally. Specifically, we are interested in showing both positive and negative results on the best-possible consistency that can be achieved by r -robust strategies, for any given r .

Our second objective goes beyond the consistency/robustness tradeoffs, and we evaluate the performance of the search strategy beyond the two extreme scenarios of error-free and adversarial error. Specifically, we study the novel setting in which the searcher defines an application-specific *tolerance* parameter H that determines its desired tolerance to errors or, equivalently, an anticipated upper bound on the prediction error (that may be known by historical data on previous searches). This parameter is defined appropriately for each of the three prediction models we study. Namely, in the untrusted advice model, H is related to the number of erroneous advice bits (or query responses); in the directional model H describes the distance of the predicted ray index to the one of the actual hiding target; and in the positional model, H is related to the distance between the predicted and the actual target position. The tolerance model is motivated by recent works in learning-enhanced online algorithms with *weak predictions*, in which the prediction is an upper bound of some pertinent parameter of the input (see e.g., online knapsack with frequency predictions [30], where the prediction is an upper bound on the size of items that appear online). Our objective is thus to quantify the tradeoff between the competitive ratio and the robustness as a function of the tolerance and other parameters of the problem (e.g., the number of queries, in the query-based model). Following [30], we will make use of the term *weak prediction* to refer to this setting.

The problems we study have applications in more general decision-making settings that go beyond the confines of search theory. This is since m -ray search, as discussed above, is an abstraction of resource allocation among m different tasks. To illustrate with an example, consider a researcher who has to allocate time among m different projects, without knowing ahead of time which project will be completed successfully. The researcher, however, may have some intuition about which of these tasks is the most likely to succeed. This problem fits the m -ray search abstraction with a directional prediction. In the weak prediction setting, H describes, more generally, the specific projects which the researcher believes are more likely to be successful.

Learning-augmented search has received attention in recent years. [3] studied consistency/robustness (Pareto) tradeoffs for linear search in the three prediction models described above. [9] studied a graph search setting where every node in the graph provides a prediction of its distance to the target vertex. [22] showed how to robustify graph exploration algorithms, where the prediction is related to the spanning tree of the explored graph.

1.2 Contribution

Our first results apply to the untrusted advice model (i.e, the k -query model). We prove tight upper and lower bounds on the best consistency that an r -robust strategy for linear search can achieve, for any $r \geq 9$, any size of advice k , and with no assumptions on the nature of the strategy. This improves upon both the upper and the lower bounds of [3], which gave a non-tight lower bound for $k = 1$ and $r = 9$, and a non-tight lower bound for $r > 9$ and $k = 1$ for a restricted class of strategies called *asymptotic*. Here, the challenge is on the lower bound side. Specifically, we reduce the problem to a *parallel search* problem that involves 2^k searchers, and we rely on a novel application of Gal’s functional theorem [26] to prove an information-theoretic tight lower bound. While this theorem has been previously applied in parallel search problems [44], its application in our setting is much more challenging, since we require that each of the 2^k searchers must be individually r -robust. Specifically, unlike previous works, the proof requires an explicit *labeling* scheme that maps the search lengths of each parallel searcher to lengths of a “global” sequence. We also extend our upper bound to weak predictions, by applying tools from the theory of games with a *lying responder* [50], in order to bound the effect of erroneous query responses to the performance.

Our second class of results is on the directional prediction model of m -ray search. We give the first upper and lower bounds on the consistency-robustness tradeoffs, which extend those of [3] to star search. Here, the main challenge is again on the lower bound side, and specifically in the weak predictions setting. We show how a generalization of a biased search approach, in which the searcher allocates more time towards the predicted ray, allows us to prove an asymptotically tight bound on the competitive ratio as a function of the tolerance and the number of rays.

Last, we show tight (Pareto-optimal) consistency-robustness tradeoffs for m -ray search in the positional model. As with the directional model, the only previous known results applied to linear search [3]. The proof uses tools that circumvent the exact study of linear recurrence relations inherent in m -ray search problems. To our knowledge, this is a new approach towards impossibility results on this type of search games. As with the other models, we also provide tight upper and lower bounds on the competitive ratio under weak predictions. We emphasize that beyond the tight and near-tight results, the generalization to star search and the accompanied analysis under weak predictions are conceptually novel aspects of this work and extend the performance guarantees beyond the consistency/robustness tradeoffs.

Due to space limitations, we omit or sketch technical details in some of the proofs.

2 Preliminaries

We review some notation and known results concerning m -ray searching. Without predictions, a strategy is described by a sequence of the form $X = (x_i, u_i)_{i \geq 0}$; we refer to i as the *iteration*¹ of the strategy, to x_i as the *length* of the i -th search *segment*, to u_i as the ray searched in iteration i , and to the point at which the searcher turns in the i -th iteration as the

¹ We will consider a numbering of iterations that starts either with 0, or with 1, depending on which simplifies the presentation.

corresponding *turn point*. Note that for linear search ($m = 2$), we may assume, without loss of generality, that $u_i = i \bmod m$, and that $x_{i+2} > x_i$. We make the standing assumption that the target lies within distance at least a fixed value, otherwise every strategy has unbounded competitive ratio. It is well-known that the worst-case hiding positions of the target, i.e., the positions that maximize $\text{Cr}(X)$, are infinitesimally beyond the turn points of a searcher that follows X , namely, at distances $x_i + \epsilon$ on rays u_i , for $\epsilon > 0$.

A strategy for m -ray search is called *cyclic*, if it explores the rays in a fixed permutation of $\{0, \dots, m-1\}$, e.g., if $u_i = i \bmod m$. The competitive ratio of a cyclic strategy of the form $X = (x_i, i \bmod m)_i$ is easily shown to be equal to

$$\text{Cr}(X) = 1 + \sup_i \frac{2 \sum_{j=0}^{i+m-1} x_j}{x_i}. \quad (2)$$

In particular, a cyclic strategy is called *geometric* if $x_i = b^i$, for some fixed $b > 1$ which is called the *base* of the strategy; we will denote such strategies by G_b . Geometric strategies are significant since they are often optimal for several variants of linear search. From (2), it follows that the competitive ratio of G_b is therefore equal to $1 + 2b^m/(b-1)$. This expression is minimized for $b = m/(m-1)$, and the resulting optimal competitive ratio, denoted by r_m^* , is equal to

$$1 + 2\rho_m^*, \quad \text{where } \rho_m^* = \frac{m^m}{(m-1)^{m-1}}.$$

Thus, given $r \geq r_m^*$, strategy G_b has competitive ratio at most r if $b^m/(b-1) \leq \rho_r$, where ρ_r is defined to be equal to $(r-1)/2$. We will denote by b_r the largest such b for which G_b is r -competitive, i.e., the largest real root of the function $f(x) = x^m/(m-1) - \rho_r$.

Under the prediction framework, the searcher is given some information h in regards to the target t , and determines a strategy X_h (we will often omit h when it is clear from context). Following [3], we define the *consistency* of a strategy as its competitive ratio assuming no prediction error, and its *robustness* as its competitive ratio assuming adversarial prediction error. Note, in particular, that the robustness of a strategy is equal to its competitive ratio without any prediction, and we will thus use these two terms interchangeably. We say that a strategy is r -robust if its robustness is at most r (similarly for the consistency), and that it is *Pareto-optimal* if its consistency and robustness are in an optimal tradeoff relation.

Let $Y = (y_i)_{i=0}^\infty$ denote a sequence in \mathbb{R}^+ . We define α_Y as $\alpha_Y = \limsup_{n \rightarrow \infty} y_n^{1/n}$. This parameter appears prominently in Gal's theorem [26] which, informally, gives a lower bound on the supremum of a set of functionals by the supremum of these functionals over geometrically increasing sequences. From it, it follows that any m -ray search strategy Y with search lengths $(y_i)_i$ has competitive ratio at least $1 + 2 \frac{\alpha_Y^m}{\alpha_Y - 1}$, hence if the strategy is r -competitive it must be that $\alpha_Y \leq b_r$.

3 Linear search with untrusted advice

In this section, we study linear search in a model in which the prediction is an untrusted advice string of size k . We first show optimal upper and lower bounds on the best consistency of r -robust strategies, then in Section 3.1 we study the extension to weak predictions.

Our results will show and exploit connections between a single-searcher strategy with k -bit advice, and a *multi-searcher* strategy with 2^k parallel searchers, but no advice. Hence, we first present some definitions and notation concerning the setting of $p > 1$ parallel searchers, labeled from the set $\{0, \dots, p-1\}$. In a p -searcher strategy, each searcher j defines its

own strategy of the form $X_j = (x_{j,i}, u_{j,i})_{i=0}^\infty$. We thus denote the p -searcher strategy as $\mathcal{X} = \{X_j\}_{j=0}^{p-1}$, or equivalently, we say that it is defined by the set $\{X_j\}_{j=0}^{p-1}$. The competitive ratio of a p -searcher strategy is the worst-case ratio of the first time one of the p searchers finds the target t (assuming unit-speed searchers) and the distance $d(t)$ of the target from the origin [44].

Observe that the optimal consistency of an r -robust strategy with k advice bits is equal to the competitive ratio of a parallel search strategy that is defined by 2^k searchers, each of which is individually r -robust. Namely, if the advice is error-free, it can be used to select the single-searcher strategy, among the 2^k ones, that reaches the target at optimal cost. Note that, by construction, the robustness of this strategy is at most r , since each individual searcher is r -robust. This observation applies to both positive and negative results on the consistency/robustness tradeoffs.

We first show an upper bound on the consistency of r -robust strategies:

► **Theorem 1.** *For any $r \geq 9$, there is an r -robust strategy for searching on the line with k -bit advice that has consistency at most $1 + 2^{\frac{b_r^{1/q}}{b_r - 1}}$, where $q = 2^{k-1}$.*

Proof sketch. Let \mathcal{S} denote the 2^k -parallel strategy as defined by the set S_0, \dots, S_{2^k-1} where

$$S_j = (b^{j+iq}, i \bmod 2), \text{ if } j \text{ is even and } S_j = (b^{j+iq}, (i+1) \bmod 2), \text{ if } j \text{ is odd,}$$


for some $b > 1$ that will be specified later. That is, each individual strategy is near-geometric, and half of the searchers explore ray 0 in their first iteration, whereas the other half explore ray 1. We require that each strategy in \mathcal{S} is r -robust which implies, from the discussion in Section 2, that b must satisfy $b^q \leq b_r$, hence $b \leq b_r^{1/q}$. ◀

Note that this upper bound is not only of theoretical significance, but can be obtained in practice via a query-based implementation. This is because the i -th advice bit can be interpreted, equivalently, as a response to a subset query that asks whether the target is hiding within a specific subset of the infinite line. Informally, the theorem shows which questions to ask to k different experts² about the whereabouts of the target so as to maximize the efficiency of search, while remaining robust to adversarial responses.

We now move to the lower bound. We first show a useful property of parallel search. To illustrate the property, consider Figure 1, which shows the first segments of a p -parallel strategy defined by S_1, \dots, S_p . In this example, the first segment of strategies S_1, \dots, S_i is to the left ray of the line, whereas the first segment of S_{i+1}, \dots, S_p is to the right ray of the line. Furthermore, the lengths of these segments are in increasing order, in the left and the right ray, respectively, as illustrated. We observe that, without loss of generality, a target that hides infinitesimally beyond the first turnpoint in S_j , with $j \in \{1, \dots, p-1\}$ is first discovered by S_{j+1} , and if $j = p$, it is first discovered by S_1 . This is because, if this was not the case, then one of the strategies would mark its second turn before it had explored any new parts of the line, which would mean that the corresponding second segment would be redundant and thus could be omitted.

We can argue, inductively, that the same property extends not only to targets hiding infinitesimally beyond the first turn points, but beyond *every* turn point. To formalize this concept, let \mathcal{S} be a p -parallel strategy defined by single-searcher strategies S_1, \dots, S_p . We say that S_j is *responsible for the i -th turn point of S_i* if S_j is the first strategy in \mathcal{S} to discover a

² Experts may be inherently erroneous; in Theorem 4 we extend the result to account for query errors.



■ **Figure 1** A snapshot of the first iteration in a p -parallel strategy.

target hiding infinitesimally beyond the i -th turnpoint of S_i . The following lemma shows that it suffices to consider p -parallel strategies in which a “snapshot” of the first iteration of the p individual strategies provides a global picture about the relative turnpoints of all individual strategies, for all subsequent iterations. This will help us setup the lower bound.

► **Lemma 2.** *For any p -searcher strategy \mathcal{S} , there is a p -searcher strategy $\mathcal{S}' = \{S'_1, \dots, S'_p\}$ such that there is a bijection $\pi : \{1, \dots, p\} \rightarrow \{1, \dots, p\}$ with the property that for any $j \in \{1, \dots, p\}$, $S'_{\pi(j)}$ is responsible for the i -th turn point of S'_j for all $i \in \mathbb{N}^+$, and \mathcal{S}' has competitive ratio no worse than \mathcal{S} .*

We now show how to prove the lower bound.

► **Theorem 3.** *For any $r \geq 9$, every r -robust strategy for searching on the line with untrusted advice of size k has consistency at most $1 + 2^{\frac{b_r^{1/q}}{b_r - 1}}$, where $q = 2^{k-1}$.*

Proof sketch. For convenience of notation, let $n = 2^k$, and let \mathcal{S} be an n -parallel strategy defined by S_1, \dots, S_n , and which satisfies the property of Lemma 2. Let i be such that strategies S_1, \dots, S_i start their first iteration to the left (in increasing order of this length), whereas S_{i+1}, \dots, S_n start their first iteration to the right (again in increasing length order). Thus, we have that S_j is responsible for S_{j-1} , for all $j \in \{2, \dots, 2^k\}$, whereas S_1 is responsible for S_n . The situation is depicted in Figure 1, where $p = n$.

Let $s_{j,m}$ denote the search length of the m -th iteration of S_j . For any fixed m , consider a target hiding infinitesimally beyond the m -th turn point of S_j , for each j . Since S_{j+1} is responsible for S_j , for all $j \in \{1, \dots, i-1, i+1, \dots, n-1\}$, we have

$$\text{Cr}(\mathcal{S}) \geq 1 + 2^{\frac{\sum_{l=1}^{m-1} s_{j+1,l}}{s_{j,m}}}, \quad \text{for all } j \in \{1, \dots, i-1, i+1, \dots, n-1\}. \quad (3)$$

In addition, since S_{i+1} and S_1 are responsible for S_i and S_n , respectively, we have that

$$\text{Cr}(\mathcal{S}) \geq 1 + 2^{\frac{\sum_{l=1}^{m-1} s_{i+1,l}}{s_{i,m-1}}} \quad \text{and} \quad \text{Cr}(\mathcal{S}) \geq 1 + 2^{\frac{\sum_{l=1}^{m-1} s_{1,l}}{s_{n,m-1}}}, \quad (4)$$

and note the subtle, but important differences in the indexing of the denominators between (3) and (4). This motivates our next step, in which we *label* the lengths of all search segments of the 2^k strategies in \mathcal{S} in a way that will allow us to use the above lower bounds. Let $\{x_l\}_{l=1}^\infty$ denote the set of all segment lengths in the parallel strategy \mathcal{S} . We map bijectively each such length to a segment length of one of the strategies S_j , according to the following function.

$$s_{j,m} = \begin{cases} x_{j+mn}, & \text{if } j \in [1, \dots, i-1] \\ x_{j-1+mn} & \text{if } j \in [i+1, \dots, n-1] \\ x_{n-1+mn}, & \text{if } j = i \\ x_{n+mn}, & \text{if } j = n. \end{cases} \quad (5)$$

Combining (3), (4) and the mapping (5), we can show that

$$\text{Cr}(\mathcal{S}) \geq 1 + 2 \sup_m \frac{\sum_{l=1}^{mn} x_l}{\sum_{l=0}^{n-1} x_{mn-1+l}}. \quad (6)$$

Define the functional F_m to be $F_m = \frac{\sum_{l=1}^{mn} x_l}{\sum_{l=0}^{n-1} x_{mn-1+l}}$. It is easy to see that this functional satisfies the conditions of Gal's Theorem [26]. Therefore,

$$\text{Cr}(\mathcal{S}) \geq 1 + 2 \sup_m \frac{\sum_{l=1}^{mn} \alpha^l}{\sum_{l=0}^{n-1} \alpha^{mn-1+l}}, \quad \text{where } \alpha = \limsup_{l \rightarrow \infty} x_l^{1/l}. \quad (7)$$

If $\alpha \leq 1$, the RHS of (7) is unbounded. If $\alpha > 1$, (7) gives

$$\text{Cr}(\mathcal{S}) \geq 1 + 2 \sup_m \frac{\alpha^{mn+1} - \alpha}{(\alpha - 1)\alpha^{mn-1} \frac{\alpha^n - 1}{\alpha - 1}} \geq 1 + 2 \frac{\alpha^2}{\alpha^n - 1}. \quad (8)$$

We will now use the fact that each strategy in \mathcal{S} is individually r -robust, in order to bound α from below, and thus $\text{Cr}(\mathcal{S})$ as well. From [31] we know that any r -competitive single searcher strategy of the form $Y = (y_j)_{j=1}^\infty$ satisfies $y_j = O(b_r^j)$. Given the labeling scheme (5), it follows that $x_j = O(b_r^{j/n})$, hence from the definition of α , $\alpha \leq b_r^{1/n}$. Therefore, (8) gives

$$\text{Cr}(\mathcal{S}) \geq 1 + 2 \frac{b_r^{2/n}}{b_r - 1} = 1 + 2 \frac{b_r^{1/q}}{b_r - 1},$$

which completes the proof. \blacktriangleleft

We give some intuition behind the significance of the mapping (5) in the proof. The labeling accomplishes two goals: First, it leads to (6), whose sums in the numerator and the denominator contain summands with “contiguous” indices: this is an essential requirement for the application of Gal's theorem. Second, it implies that each strategy S_j is of the form $(x_{\pi(j)+nl})_{l=0}^\infty$, where π is a bijection over $\{1, \dots, n\}$, which allows us to argue that $\alpha \leq b_r^{1/n}$.

3.1 Extension to weak predictions

We show how to extend the upper bound to incorporate weak predictions. In this setting, as discussed in Section 1, given advice of size k , and robustness requirement $r \geq 9$, the searcher specifies a tolerance parameter $H \leq k/2$. The objective is to obtain an r -robust strategy of minimum competitive ratio assuming that at most H advice bits are erroneous.

To address this problem, we will make use of a result by Rivest *et al.* [50], who studied games with a *lying responder*. In their setting, given $k \in \mathbb{N}^+$, $H \leq k/2$, and a domain $\mathcal{D} = \{1, \dots, m\}$, the objective is to find the index of an unknown $x \in \mathcal{D}$, using k queries, of which up to H may receive incorrect responses. A query can be a comparison query of the form “Is $x \leq M$?”, for some given $M \in [1, m]$, or more generally, a subset query of the form “Is x in S ?”, where S is a subset of the domain \mathcal{D} . Define the sum of binomial coefficients $\binom{N}{m} := \sum_{j=0}^m \binom{N}{j}$, for $m \leq N$. In [50] it was shown that as long as $m \leq 2^k / \binom{k}{H}$, k comparison queries suffice to find x in the above game, in the presence of at most $H \leq k/2$ errors. This leads to the following extension of Theorem 1.

► **Theorem 4.** *For any $r \geq 9$, and any $H \leq k/2$, there is an r -robust strategy for searching on the line with k -bit advice that has competitive ratio at most $1 + 2 \frac{b_r^{1/q}}{b_r - 1}$, where $q = 2^{k-1} / \binom{k}{H}$.*

We can further show that

$$\frac{1}{q} \leq \frac{1}{2^{k(1-\mathcal{H}(\frac{H}{k}))}-1},$$

where \mathcal{H} denotes the *binary entropy* function. This allows for a direct comparison to the Pareto-optimal upper bound of Theorem 1. In particular, we observe that as k increases, the effect of the advice error in the competitive ratio becomes marginal, even if H is as high as linear in k .

4 Ray search with directional prediction

In this section, we study m -ray search in the setting in which the prediction is the ray of the hiding target. Without loss of generality, we suppose that the prediction is the ray indexed 0.

For the upper bound, we consider the following strategy that generalizes the biased search approach of [20]. The searcher fixes some $b > 1$ and $\delta > 1$, to be specified later, and explores the rays in the cyclic order $0, 1, \dots, m-1$. If the ray visited in the i -th iteration is ray 0, it explores it to a length equal to δb^i , otherwise, i.e., if the visited ray is in $\{1, \dots, m-1\}$, it explores it to a length equal to b^i . Thus the search combines elements of geometric search with a bias towards the predicted ray, as expressed by the parameter δ .

► **Theorem 5.** *For every $b > 1$, $\delta > 1$, the strategy described above has consistency at most $1 + 2\frac{b^m}{b^m-1} + \frac{2}{\delta}\frac{b^m}{b^m-1}\frac{b^m-b}{b-1}$, and robustness at most $1 + 2\delta\frac{b^{m+1}}{b^m-1} + 2\frac{b^{m+1}}{b^m-1}\frac{b^m-b}{b-1} - 2b^m$.*

Observe that if $\delta = 1$, then both the consistency and the robustness of the above strategy are equal to $1 + 2\frac{b^m}{b^m-1}$, as expected (i.e., the competitive ratio of a geometric strategy with base b). For any fixed b , by increasing δ , the consistency of the resulting strategy improves, at the expense of its robustness. We would like thus to optimize the robustness by choosing δ as a function of b , m and the desired consistency, however it is not obvious that this is possible analytically. Instead, suppose that we choose $b = (m+1)/m$, namely the base of the geometric strategy that results in an optimal competitive ratio for m -ray search equal to $r_m^* = 1 + 2\rho_m^*$. Then $\frac{b^m}{b^m-1} \leq e/(e-1)$, which implies that the consistency of the strategy is at most

$$1 + 2\frac{e}{e-1} + \frac{2}{\delta}\frac{e}{e-1}(\rho_m^* - m) = \frac{2}{\delta}\frac{e}{e-1}(\rho_m^* - m) + O(1).$$

On the other hand, the robustness of the strategy is at most

$$1 + 2\delta\frac{81}{36} + 2\frac{81}{36}(\rho_m^* - m) - 4 = \frac{9}{2}(\delta + \rho_m^* - m) + O(1).$$

Therefore, if we would like the strategy to be c -consistent, where $c = O(1) + 2\tilde{c}$, for some \tilde{c} , we can choose δ to be equal to $\frac{e}{(e-1)\tilde{c}}(\rho_m^* - m)$, and the resulting robustness is then at most $\frac{9}{2}(\rho_m^* - m)(1 + \frac{e}{(e-1)\tilde{c}}) + O(1)$. We can also obtain more precise tradeoffs as $m \rightarrow \infty$, since in this case, it is known that $r_m^* = 1 + 2\rho_m^* = 1 + 2em$.

► **Corollary 6.** *For $m \rightarrow \infty$, the above strategy has consistency at most $1 + 2\frac{e}{e-1} + \frac{2}{\delta}em$, and robustness at most $2e(m + \frac{\delta}{e-1})$. In particular, given \tilde{c} , the strategy is $(O(1) + 2\tilde{c})$ -consistent, and $(O(1) + 2em(1 + \frac{e}{(e-1)\tilde{c}}))$ -robust.*

Next, we show a negative result on the tradeoff between the consistency and robustness that any strategy can achieve. The proof follows an approach that we generalize in the weak predictions setting (proof of the lower bound in Theorem 8).

► **Theorem 7.** *Any c -consistent strategy for searching with directional prediction, where $c = 1 + 2\tilde{c}$, has robustness at least $1 + 2\rho_{m-1}^*(1 + \frac{1}{\tilde{c}-1})$. In particular, for $m \rightarrow \infty$, its robustness is at least $1 + 2e(m-1)(1 + \frac{1}{\tilde{c}-1})$.*

4.1 Extension to weak predictions

We consider the model in which if the predicted ray is $h \in \{0, \dots, m-1\}$, then the searcher would like to minimize its competitive ratio, assuming that the target hides in one of the rays in the interval $[h - H \bmod m, h + H \bmod m]$, where $H \leq m/2$ is the tolerance parameter. This captures the case that the hiding ray is expected to be in a vicinity of the predicted ray, with respect to the tolerance of the searcher. We denote this set of rays by R_H , and its complement, by \bar{R}_H , and note that $|R_H| = 2H + 1$ and $|\bar{R}_H| = m - 2H - 1$. Without loss of generality, we may assume that $R_H = \{0, \dots, 2H\}$.

We prove an asymptotically tight bound on the tradeoff between the competitive ratio and the robustness, that generalizes the error-free setting. Note that every strategy has competitive ratio at least $1 + 2\rho_{2H+1}^*$, since the weak prediction may incur a search in a $(2H+1)$ -ray star. We also obtain an interesting corollary: as the competitive ratio approaches the optimal bound of $1 + 2\rho_{2H+1}^*$, the robustness increases dramatically.

► **Theorem 8.** *For every $\tilde{c} > \rho_{2H-1}^*$ there exists a strategy with directional hint that has competitive ratio $1 + 2\tilde{c}$, and robustness at most $O(\frac{\rho_{2H+1}^*}{\tilde{c} - \rho_{m-2H-1}^*}(m - 2H))$. Furthermore, this bound is tight, i.e., every strategy of competitive ratio $1 + 2\tilde{c}$ has robustness at least $\Omega(\frac{\rho_{2H+1}^*}{\tilde{c} - \rho_{m-2H-1}^*}(m - 2H))$.*

Proof. We first prove the upper bound, which generalizes the strategy we used in the context of consistency/robustness tradeoffs. Define $b = \frac{2H+1}{2H}$, i.e., the optimal base of a geometric strategy for searching in a $(2H+1)$ -ray star, and $\delta > 1$, to be specified later. Consider a cyclic strategy which visits rays $0, \dots, m-1$ in this order, and which works in rounds. Specifically, in round i , it explores ray $j \in R_H$ to length $\delta b^{(2H+1)i+j}$, and every ray in \bar{R}_H to length $b^{(2H+1)i+2H}$. The competitive ratio of this strategy, assuming error at most H , is maximized for targets hiding infinitesimally beyond the turn points on ray $2H$ in R_H . Simple calculations show that the competitive ratio is

$$1 + 2\rho_{2H+1}^* + \Theta\left(\frac{1}{\delta} \frac{b^{2H+1}}{b^{2H+1}-1}(m - 2H - 1)\right),$$

and note that $\frac{b^{2H+1}}{b^{2H+1}-1}$ is at most $e/(e-1)$, by the choice of b . Thus, for the competitive ratio to be at most $1 + 2\tilde{c}$, it must be that

$$\delta \in \Omega\left(\frac{m - 2H}{\tilde{c} - \rho_{2H+1}^*}\right). \quad (9)$$

The robustness of the strategy is evaluated for a target hiding at distance infinitesimally beyond the turn points of the searcher on ray $m-1$. After simple calculations we obtain that the robustness is at most $O(\delta \rho_{2H+1}^*(m - 2H + 1))$, which from (9) is at most $O(\frac{\rho_{2H+1}^*}{\tilde{c} - \rho_{2H+1}^*}(m - 2H))$, and which proves the upper bound.

We now proceed with the lower bound. Any strategy for the problem consists of *phases*, which alternate between searching a subset of R_H and a subset of \bar{R}_H . Namely, every strategy X is of the form $X = (x_i)_{i \geq 0}$, in which x_i , for even i , describes the aggregate explored length of X on rays that belong exclusively in R_H , and x_i for i odd, describes the aggregate explored length on rays that belong exclusively in \bar{R}_H . Thus, in each phase i , the searcher incurs a cost of $2x_i$, for all i (except for the phase at which the target is found).

Define $S_i = \sum_{j=0}^i x_{2j}$, and $S'_i = \sum_{j=0}^{i-1} x_{2j+1}$. Then from the competitiveness of the strategy, for all even i , and for any target t_H hiding in R_H that is discovered in phase $2(i+1)$,

$$1 + 2 \frac{S_i + S'_i}{d(t_H)} \leq c \Rightarrow \tilde{c} \geq \frac{S_i + S'_i}{d(t_H)}. \quad (10)$$

We know that searching in a $(2H+1)$ -ray has competitive ratio at least $1 + 2\rho_{2H+1}^*$. This means that for all i , there exists a target d_i that is discovered in phase $2(i+1)$ such that

$$1 + 2 \sup_i \frac{S_i}{d_i} \geq 1 + 2\rho_{2H+1}^*,$$

hence there exists some \bar{i} for which the above inequality gives $1 + 2 \frac{S_{\bar{i}}}{d_{\bar{i}}} \geq 1 + 2\rho_{2H+1}^* - \epsilon$, where $\epsilon \rightarrow 0$, as \bar{i} is allowed to be unbounded. To simplify the exposition, we can thus assume that $\epsilon = 0$, and obtain

$$d(t_{\bar{i}}) \leq \frac{S_{\bar{i}}}{\rho_{2H+1}^*}. \quad (11)$$

Moreover, for any i , there exists a hiding position for a target t'_i in \bar{R}_H that is first discovered in phase $2i+1$ it must be that

$$1 + 2 \frac{S'_i}{d(t'_i)} \geq m - 2H - 1 \Rightarrow d(t'_i) = O\left(\frac{S'_i}{m - 2H - 1}\right), \quad (12)$$

since $2S'_i$ describes the cost incurred by the searcher on rays in \bar{R}_H , right before phase $2i+1$ starts. Note also that this inequality holds *for all* i , unlike (11), that holds only for \bar{i} .

To bound the robustness, consider the phase $2\bar{i}+1$, with \bar{i} as defined above, and the target $t'_{\bar{i}}$, again as defined above. Then we have that

$$\text{Robustness} \geq 1 + 2 \frac{S_{\bar{i}} + S'_{\bar{i}}}{d(t'_{\bar{i}})} = \Omega\left(\left(1 + \frac{S_{\bar{i}}}{S'_{\bar{i}}}\right)(m - 2H)\right), \quad (13)$$

from (12). Moreover, from (10) and (11) we have that

$$\frac{S_{\bar{i}} + S'_{\bar{i}}}{S'_{\bar{i}}} \leq \frac{\tilde{c}}{\rho_{2H+1}^*} \Rightarrow \frac{S_{\bar{i}}}{S'_{\bar{i}}} \geq \frac{\rho_{2H+1}^*}{\tilde{c} - \rho_{2H+1}^*},$$

and substituting the above inequality to (10) yields the result. \blacktriangleleft

5 Ray search with positional prediction

In this section we study m -ray searching in the setting in which the prediction is the position of the target in the star environment. Namely, the prediction h is a pair (d_h, u_h) , where d_h is the predicted distance from O and u_h is the predicted ray. We first show the upper bound.

► **Theorem 9.** *For any $r \geq r_m^*$, there is an r -robust strategy of consistency at most $1 + 2 \frac{1}{b_r - 1}$.*

We will now show that the strategy of Theorem 9 is Pareto-optimal. The proof of the following theorem generalizes, but also simplifies the lower bound of [3] which applies only to linear search ($m = 2$). The crux in the proof is to exploit the properties of the parameter α_Y , where Y will be defined as the sequence of search lengths of a cyclic strategy defined by a linear recurrence relation. In particular, these properties allow us to bypass technical complications related to the study of such relations, by establishing appropriate lower bounds (as opposed to solving the recurrence relation).

► **Theorem 10.** *For any $r \geq r_m^*$, no r -robust strategy for searching with positional prediction has consistency better than $1 + 2\frac{1}{b_r - 1}$.*

Proof sketch. Using techniques rooted in previous studies of star search, we can show that there exists a *cyclic* strategy of the form

$$y_{j_h} = h, \text{ for some index } j_h \quad \text{and} \quad 1 + 2\frac{\sum_{i=0}^j y_i}{y_{j-m+1}} = r, \text{ for all } j \in [m-1, j_h],$$

which is r -robust and minimizes the consistency (among all r -robust strategies). It follows that for all $j \leq j_h$, the search length y_j is determined by the recurrence

$$y_j = r(y_{j-m+1} - y_{j-m}), \quad (14)$$

with some initial conditions y_0, \dots, y_{m-1} . We have that

$$d(Y, h) = \sum_{j=0}^{j_h} y_j = d_h + 2 \sum_{j=0}^{j_h-1} y_j = d_h + 2r(y_{j_h-m+1} - y_0) + \sum_{j=0}^{m-1} y_j, \quad (15)$$

where we used the fact that $\sum_j y_j$ is telescoping, as seen by (14). From (15) we have $\frac{d(Y, h)}{d_h} = 1 + 2\frac{ry_{j_h-m+1}}{d_h} + \frac{1}{d_h}(\sum_{j=0}^{m-1} y_j - y_0)$. Since Y is cyclic, and $d_h = y_{j_h}$, we obtain that

$$\sup_h \frac{d(Y, h)}{d_h} \geq 1 + 2 \sup_{j_h} \frac{ry_{j_h-m+1}}{y_{j_h}} = 1 + 2r \sup_{j_h} \frac{y_{j_h-m+1}}{y_{j_h}}. \quad (16)$$

Define the functional $F_j(Y) = \frac{y_{j-m+1}}{y_j}$. This functional satisfies the conditions of Gal's functional theorem [26], hence $\sup_j \frac{y_{j-m+1}}{y_j} \geq \frac{\alpha_Y^{j-m+1}}{\alpha_Y} = \alpha_Y^{-m}$. Since Y is r -robust, as discussed in Section 2 $r \geq \text{Cr}(X) \geq \frac{\alpha_Y^m}{\alpha_Y - 1}$, where it must be $\alpha_Y \leq b_r$, from the definition of b_r . Thus, (16) gives

$$\sup_h \frac{d(Y, h)}{d_h} \geq 1 + 2\frac{\alpha_Y^m}{\alpha_Y - 1} \alpha_Y^{-m} = 1 + 2\frac{1}{\alpha_Y - 1} \geq 1 + 2\frac{1}{b_r - 1}.$$

Hence, the consistency of Y is at least $1 + 2\frac{1}{b_r - 1}$, and thus so is the consistency of X . ◀

5.1 Extension to weak predictions

Given the prediction h related to a target t , we define the prediction *error* as the distance between t and h in the star, normalized by the distance $d(h)$ of the prediction's position from the origin. Namely, $\eta = \frac{|d(t) - d(h)|}{d(h)}$. We distinguish between different types of the error: If d and h are in the same ray, but $d(t) > d(h)$ we call the error *positive*, whereas if d and h are in the same ray, but $d(t) < d(h)$ we call the error *negative*. If the error is neither positive or negative, then h and t are in different rays.

Let X_h denote the Pareto-optimal strategy of Theorem 9. Let $H > 0$ denote the tolerance parameter that is specified by the searcher, and consider the strategy $X_{h(1+H)}$, i.e., the strategy that pretends that the prediction is at the same ray as h , but at distance $d(h)(1+H)$ from O . The following result is a corollary of Theorem 9.

► **Corollary 11.** *For any $H > 0$ and $r \geq r_m^*$, strategy $X_{h(1+H)}$ is r -robust and has competitive ratio at most $\min\{1 + 2\frac{1+H}{b_r - 1}, r\}$, if the error is either positive or negative, and at most H . Otherwise, its competitive ratio is at most r .*

Last, we can show that the above tradeoffs are tight.

► **Theorem 12.** *For any r -robust strategy with positional prediction, there exists $q > 0$ such that its competitive ratio is no better than $\min\{1 + 2\frac{1+q}{b_r - 1}, r\}$ for positive or negative error at most q . Otherwise, its competitive ratio is at least r .*

References

- 1 Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*, volume 55 of *International series in operations research and management science*. Kluwer, 2003.
- 2 Steve Alpern and Thomas Lidbetter. Mining coal or finding terrorists: The expanding search paradigm. *Operations Research*, 61(2):265–279, 2013.
- 3 Spyros Angelopoulos. Online search with a hint. In *Proceedings of the 12th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 51:1–51:16, 2021.
- 4 Spyros Angelopoulos. Further connections between contract-scheduling and ray-searching problems. *Journal of Scheduling*, 25(2):139–155, 2022.
- 5 Spyros Angelopoulos, Christoph Dürr, and Shendan Jin. Best-of-two-worlds analysis of online search. In *Proceedings of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 7:1–7:17, 2019.
- 6 Spyros Angelopoulos, Christoph Dürr, and Thomas Lidbetter. The expanding search ratio of a graph. In *Proceedings of the 33rd Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 9:1–9:14, 2016.
- 7 Spyros Angelopoulos and Shahin Kamali. Contract scheduling with predictions. *Journal of Artificial Intelligence Research*, 77:395–426, 2023.
- 8 Ricardo A. Baeza-Yates, Joseph C. Culberson, and Gregory G.E. Rawlins. Searching in the plane. *Information and Computation*, 106:234–244, 1993.
- 9 Siddhartha Banerjee, Vincent Cohen-Addad, Anupam Gupta, and Zhouzi Li. Graph searching with predictions. In *Proceedings of the 14th Conference on Innovations in Theoretical Computer Science (ITCS)*, volume 251 of *LIPICs*, pages 12:1–12:24, 2023.
- 10 Anatole Beck. On the linear search problem. *Naval Research Logistics*, 2:221–228, 1964.
- 11 Anatole Beck and Donald J. Newman. Yet more on the linear search problem. *Israel Journal of Mathematics*, 8:419–429, 1970.
- 12 Richard Bellman. An optimal search problem. *SIAM Review*, 5:274, 1963.
- 13 Daniel S. Bernstein, Lev Finkelstein, and Shlomo Zilberstein. Contract algorithms and robots on rays: Unifying two scheduling problems. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1211–1217, 2003.
- 14 Lucas Boczkowski, Amos Korman, and Yoav Rodeh. Searching a tree with permanently noisy advice. In *ESA 2018-26th Annual European Symposium on Algorithms*, pages 1–32, 2018.
- 15 Anthony Bonato, Konstantinos Georgiou, Calum MacRury, and Pawel Pralat. Probabilistically faulty searching on a half-line – (Extended abstract). In *LATIN 2020: Theoretical Informatics – 14th Latin American Symposium, São Paulo, Brazils*, volume 12118 of *Lecture Notes in Computer Science*, pages 168–180. Springer, 2020.
- 16 Prosenjit Bose, Jean-Lou De Carufel, and Stephane Durocher. Searching on a line: A complete characterization of the optimal solution. *Theoretical Computer Science*, 569:24–42, 2015.
- 17 Marek Chrobak, Leszek Gasieniec, Thomas Gorry, and Russell Martin. Group search on the line. In *International Conference on Current Trends in Theory and Practice of Informatics*, pages 164–176. Springer, 2015.
- 18 Anne Condon, Amol Deshpande, Lisa Hellerstein, and Ning Wu. Algorithms for distributional and adversarial pipelined filter ordering problems. *ACM Transaction on Algorithms*, 5(2):24:1–24:34, 2009.
- 19 Jurek Czyzowicz, Evangelos Kranakis, Danny Krizanc, Lata Narayanan, and Jaroslav Opatrny. Search on a line with faulty robots. *Distributed Comput.*, 32(6):493–504, 2019.
- 20 Erik D Demaine, Sándor P Fekete, and Shmuel Gal. Online searching with turn cost. *Theoretical Computer Science*, 361:342–355, 2006.
- 21 Stefan Dobrev, Rastislav Kráľovič, and Euripides Markou. Online graph exploration with advice. In *Structural Information and Communication Complexity: 19th International Colloquium, SIROCCO 2012*, pages 267–278. Springer, 2012.

- 22 Franziska Eberle, Alexander Lindermayr, Nicole Megow, Lukas Nölke, and Jens Schlöter. Robustification of online graph exploration methods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36(9), pages 9732–9740, 2022.
- 23 Pierre Fraigniaud, David Ilcinkas, and Andrzej Pelc. Tree exploration with advice. *Information and Computation*, 206(11):1276–1287, 2008.
- 24 Shmuel Gal. A general search game. *Israel Journal of Mathematics*, 12:32–45, 1972.
- 25 Shmuel Gal. Minimax solutions for linear search problems. *SIAM Journal on Applied Mathematics*, 27:17–30, 1974.
- 26 Shmuel Gal. *Search Games*. Academic Press, 1980.
- 27 Subir Kumar Ghosh and Rolf Klein. Online algorithms for searching and exploration in the plane. *Comput. Sci. Rev.*, 4(4):189–201, 2010.
- 28 Barun Gorain and Andrzej Pelc. Deterministic graph exploration with advice. *ACM Transactions on Algorithms*, 15(1):1–17, 2018.
- 29 C. Hipke, C. Icking, R. Klein, and E. Langetepe. How to find a point in the line within a fixed distance. *Discrete Applied Mathematics*, 93:67–73, 1999.
- 30 Sungjin Im, Ravi Kumar, Mahshid Montazer Qaem, and Manish Purohit. Online knapsack with frequency predictions. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 2733–2743, 2021.
- 31 Parick Jaillet and Matthew Stafford. Online searching. *Operations Research*, 49:234–244, 1993.
- 32 Ming-Yang Kao, Yuan Ma, Michael Sipser, and Yiqun Yin. Optimal constructions of hybrid algorithms. *Journal of Algorithms*, 29(1):142–164, 1998.
- 33 Ming-Yang Kao, John H Reif, and Stephen R Tate. Searching in an unknown environment: an optimal randomized algorithm for the cow-path problem. *Information and Computation*, 131(1):63–80, 1996.
- 34 David G. Kirkpatrick. Hyperbolic dovetailing. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, pages 616–627, 2009.
- 35 Dennis Komm, Rastislav Kráľovič, Richard Kráľovič, and Jasmin Smula. Treasure hunt with advice. In *Structural Information and Communication Complexity: 22nd International Colloquium, SIROCCO 2015*, pages 328–341. Springer, 2015.
- 36 E. Koutsoupias, C.H. Papadimitriou, and M. Yannakakis. Searching a fixed graph. In *Proc. of the 23rd Int. Colloq. on Automata, Languages and Programming (ICALP)*, pages 280–289, 1996.
- 37 Andrey Kupavskii and Emo Welzl. Lower bounds for searching robots, some faulty. In *Proceedings of the 37th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 447–453, 2018.
- 38 Elmar Langetepe. On the optimality of spiral search. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1–12. SIAM, 2010.
- 39 Russell Lee, Jessica Maghakian, Mohammad H. Hajiesmaili, Jian Li, Ramesh K. Sitaraman, and Zhenhua Liu. Online peak-aware energy scheduling with untrusted advice. In *Proceedings of the 12th ACM International Conference on Future Energy Systems (eEnergy)*, pages 107–123. ACM, 2021.
- 40 Tongxin Li, Ruixiao Yang, Guannan Qu, Guanya Shi, Chenkai Yu, Adam Wierman, and Steven H. Low. Robustness and consistency in linear quadratic control with untrusted predictions. *Proc. ACM Meas. Anal. Comput. Syst.*, 6(1):18:1–18:35, 2022.
- 41 Thomas Lidbetter. Search and rescue in the face of uncertain threats. *Eur. J. Oper. Res.*, 285(3):1153–1160, 2020.
- 42 Alexander Lindermayr and Nicole Megow. Repository of works on algorithms with predictions. <https://algorithms-with-predictions.github.io/about>, 2023. Accessed: 2023-04-01.
- 43 Alejandro López-Ortiz and Svem Schuierer. The ultimate strategy to search on m rays? *Theoretical Computer Science*, 261(2):267–295, 2001.

- 44 Alejandro López-Ortiz and Sven Schuierer. On-line parallel heuristics, processor scheduling and robot searching under the competitive framework. *Theoretical Computer Science*, 310(1–3):527–537, 2004.
- 45 Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. *J. ACM*, 68(4):24:1–24:25, 2021.
- 46 Andrew McGregor, Krzysztof Onak, and Rina Panigrahy. The oil searching problem. In *Proceedings of the 17th European Symposium on Algorithms (ESA)*, pages 504–515, 2009.
- 47 Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. In *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020.
- 48 Andrzej Pelc and Ram Narayan Yadav. Advice complexity of treasure hunt in geometric terrains. *Information and Computation*, 281:104705, 2021.
- 49 Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ML predictions. In *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*, pages 9661–9670, 2018.
- 50 Ronald L. Rivest, Albert R. Meyer, Daniel J. Kleitman, Karl Winklmann, and Joel Spencer. Coping with errors in binary search procedures. *J. Comput. Syst. Sci.*, 20(3):396–404, 1980.
- 51 Sven Schuierer. Lower bounds in online geometric searching. *Computational Geometry: Theory and Applications*, 18(1):37–53, 2001.
- 52 Alexander Wei and Fred Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*, 2020.