Searching in Euclidean spaces with predictions*

Sergio Cabello[†1] and Panos Giannopoulos[‡2]

[1]University of Ljubljana and Institute of Mathematics, Physics and Mechanics, Slovenia
[2]Department of Computer Science, City, University of London, UK

The problem of searching for a target positioned at some unknown location in some region is a classic search game problem that has been well-studied in both fields of Computational Geometry and Operations Research [1, 3].

In this work, we consider the problem of finding a fixed but unknown *target* point $\boldsymbol{t}$ in Euclidean space $\mathbb{R}^d$ when additional information regarding the position of the target is available in the form of *predictions* [4]. Here, the predictions are the approximate distance to $\boldsymbol{t}$ for all points visited during the search. More precisely, we assume that there is a constant $c \geq 1$ and an unknown function

$$\lambda\colon \mathbb{R}^d \to \mathbb{R}_{\geq 0} \text{ such that } \forall p \in \mathbb{R}^d \ |p\boldsymbol{t}| \leq \lambda(p) \leq c{\cdot}|p\boldsymbol{t}|.$$

We refer to such a function λ as a *c-prediction* for the target $\boldsymbol{t}$. The constant c is the *prediction factor* of λ. (For $c = 1$, the function λ gives the exact distance to the target.)

Without loss of generality, we start the search at the origin, denoted by $\boldsymbol{o}$. We want to find a curve γ that starts at $\boldsymbol{o}$ and ends at $\boldsymbol{t}$. For each point p along the search path we have traversed so far, we obtain the value $\lambda(p)$, and the search strategy decides how to continue the search depending on that information. We know when we have reached the target because $\lambda(p) = 0$ holds only when $p = \boldsymbol{t}$. The *cost* of the search is the Euclidean length of the curve γ.

As it is common in search games, we are interested in the competitive ratio of the search strategy: how does the length of the search path compares to the straight-line distance from the origin to the target? A search strategy S is $\alpha = \alpha(S, c, d)$ *competitive* if for all $\boldsymbol{t} \in \mathbb{R}^d$ and all c-predictions λ for target $\boldsymbol{t}$, the length of the path defined by S to reach $\boldsymbol{t}$ from $\boldsymbol{o}$ is at most $\alpha|\boldsymbol{ot}|$. Note that we have two possible regimes, depending on whether the prediction factor c is known or unknown to the search strategy.

Our main contributions are the following. First, we introduce a natural, new search problem in $\mathbb{R}^d$ for $d \geq 1$, under a predictions model.

Second, we show that for each dimension d and each constant prediction factor $c \geq 1$ there is a search strategy with competitive ratio smaller than $(5c)^{d+1}$. To achieve this, we use ε-nets from metric spaces, and provide a path of finite length but an infinite number of pieces. This result holds assuming that we know the prediction factor c. For *unknown* prediction factor, a slightly different search strategy leads to a competitive ratio smaller than $(10c)^{d+1}$.

Third, we show that for $c \geq 4$, any deterministic search strategy in $\mathbb{R}^d$ with c-predictions will have a competitive ratio of at least $(c/4)^{d-1} \cdot \min\{\sqrt{\pi/d}, 1\}$. For this result, we construct an infinite family of c-predictions and use a volume argument to give a lower bound on the length of any path that can discern which c-prediction from the family is the actual one. The approach is motivated by the techniques used to obtain approximation algorithms for the Euclidean Traveling Salesperson with Neighbourhoods. However, in our setting we have an infinite number of neighbourhoods, one for each c-prediction. With a slightly worse constant, the lower bound holds also for randomized search strategies.

The full version is available at [2].

*We would like to thank Kostas Tsakalidis and Konrad Swanepoel for helpful discussions.
Research supported in part by the Slovenian Research and Innovation Agency (P1-0297, J1-2452, N1-0218, N1-0285), in part by the Erdős Center, and in part by the European Union (ERC grant KARST, 101071836; ERC grant 882971, GeoScape). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

†Email: sergio.cabello@fmf.uni-lj.si.
‡Email: Panos.Giannopoulos@city.ac.uk.

References

[1] Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*. Kluwer, 2003.

[2] Sergio Cabello and Panos Giannopoulos. Searching in Euclidean spaces with predictions. Presented at WAOA 2024. Full version available at `https://arxiv.org/abs/2408.04964`.

[3] Subir Kumar Ghosh and Rolf Klein. Online algorithms for searching and exploration in the plane. *Comput. Sci. Rev.*, 4(4):189–201, 2010.

[4] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020.