# Optimal Metric Search Is Equivalent to the Minimum Dominating Set Problem 

Magnus Lie Hetland<br>Norwegian University of Science and Technology, mlh@ntnu.no


#### Abstract

In metric search, worst-case analysis is of little value, as the search invariably degenerates to a linear scan for ill-behaved data. Consequently, much effort has been expended on more nuanced descriptions of what performance might in fact be attainable, including heuristic baselines like the AESA family, as well as statistical proxies such as intrinsic dimensionality. This paper gets to the heart of the matter with an exact characterization of the best performance actually achievable for any given data set and query. Specifically, linear-time objective-preserving reductions are established in both directions between optimal metric search and the minimum dominating set problem, whose greedy approximation becomes the equivalent of an oracle-based AESA, repeatedly selecting the pivot that eliminates the most of the remaining points. As an illustration, the AESA heuristic is adapted to downplay the role of previously eliminated points, yielding some modest performance improvements over the original, as well as its younger relative iAESA2.


Keywords: Metric indexing • Baselines • Hardness • Dominating set

## 1 Introduction

Mapping out the complexity of a computational problem is generally a twopronged affair. On the one hand, there will be algorithms solving the problem, whose performance is evaluated theoretically or empirically, providing evertightening pessimistic bounds on what is possible. On the other hand, there may be lower bounds, based on reasonable complexity-theoretical assumptions, as in the case of edit distance, for example [1], or on reasoning about the fundamentals of the computational model, as in the case of sorting [11]. The endgame is when these bounds meet, showing some algorithm to be optimal.

Such bounds generally apply to the worst case, as the best-case performance tends to be trivial. For metric search, however, both the best case and the worst are quite uninformative. For a range query, one could always construct an input where examining a single object is enough-or one where there is no escaping a full linear scan. The main thrust of research attempting to describe what performance is possible has thus been directed toward empirical baselines like the AESA family $[25,10]$ and statistical hardness measures such as intrinsic dimensionality [5], ${ }^{1}$ or in some cases restricting the type of structure studied, to permit a more nuanced analysis [20].

[^0]It is, however, possible to describe exactly what performance is attainable for a given data set and query, as I show in what follows. The main equivalence result, between metric search and dominating sets, provides just such a description, i.e., the lowest number of distance computations that can resolve the query. This performance will not, in general, be attainable without some lucky guesses, but it is attainable. In addition, it is possible to give a bound on how close to this performance a polytime algorithm may come in the worst case, under reasonable complexity assumptions. The bound is tight for a sufficiently precise pivot selection heuristic, i.e., one that is able to predict which point will eliminate the most of the remainder, if used as a pivot.

In the AESA method, the index is a distance matrix, and search alternates between heuristically selecting points close to the query and eliminating remaining objects that are shown to be irrelevant. The results in this paper are based on an idea developed by Ole Edsberg, ${ }^{2}$ which involves computing an elimination matrix for a given query, with which one may implement an "oracle AESA," selecting pivots greedily based on elimination power, rather than on similarity to the query object. I build on this idea, establishing equivalence to the minimum dominating set problem. ${ }^{3}$ The main results and contributions of the paper are summarized in the following.

Reduction to domination. Sections 2 and 3 establish a linear-time objectivepreserving reduction from the problem of resolving metric range queries (and certain $k \mathrm{NN}$ queries) with as few distance computations as possible to that of finding minimum dominating sets in directed graphs. This reduction applies to an offline variant of metric search, where all query-object distances are already known. It does, however, make it possible to compute the exact optimum attainable for the online version as well. Some experimental results are provided as an illustration.

Reduction from domination. Section 4 describes a reduction in the other direction, from the dominating set problem in undirected graphs to minimizing distance computations, establishing the hardness of metric search. While it may in many cases still be feasible to determine the optimum using efficient solvers of various kinds, this does mean that under reasonable complexity-theoretical assumptions, no search method can, in general, guarantee attaining this optimum.

The reduction preserves the objective value, and for range search, the number of data objects equals the number of vertices, which means that inapproximability results for the dominating set problem carry over to metric search, with approximation bounds for the former applying to the performance of the latter, i.e., the number of distance computations. Thus, for range search, one cannot even expect to get closer than within a log-factor of the optimum.

AESA and greedy approximation. Because the objective is preserved also in reducing to domination, and the number of objects equals the number of

[^1]vertices, approximability results also translate, meaning that in principle the standard greedy selection strategy would yield the best feasible metric range search algorithm (or very close to it), in terms of distance computations in the worst case. ${ }^{4}$ As discussed in section 5, the greedy approach corresponds to the AESA family of algorithms, given the right selection heuristic, i.e., one that accurately estimates the elimination power of a potential pivot, among the remaining objects. An exact estimate here is, of course, not possible without knowing the query-pivot distances, but this correspondence does demonstrate that, in the limit, AESA is, indeed, as good as it gets. As an illustration, inspired by the greedy approximation, greedy AESA (gAESA) is proposed, taking into account which points remain to be eliminated.

## 2 Pivoting Is, of Course, Optimal

A range search using a metric $\delta$ over a set X means finding all points $x \in \mathrm{X}$ within some search radius $r$ of a given query point $q$, i.e., all points $x$ for which $\delta(q, x) \leqslant r$. Given the distances between a query $q$ and a set P of pivots, the distance $\delta(q, x)$ for any point $x$ is bounded as follows:

$$
\begin{equation*}
\max _{p \in \mathrm{P}}|\delta(q, p)-\delta(p, x)| \leqslant \delta(q, x) \leqslant \min _{p \in \mathrm{P}} \delta(q, p)+\delta(p, x) \tag{1}
\end{equation*}
$$

Leaving $q$ and P implicit, we may refer to the lower and upper bounds as $\ell(x)$ and $u(x)$, respectively. If our search radius falls outside this range, there is no need to compute $\delta(q, x)$; either the radius is small enough that we simply eliminate $x$ $(r<\ell(x))$, or it is great enough that $x$ is "eliminated" by adding it to the search result, sight unseen $(r \geqslant u(x))$.

This very direct approach of using exact, stored distances $\delta(p, x)$, pivoting, is the gold standard for minimizing the number of distance computations needed. Other approaches, which all involve coarsening the stored information in some way, may reduce the computational resources needed to eliminate candidate objects, but it should be obvious that they cannot require fewer distance computations. As the following lemma shows, the lower and upper bounds are necessarily valid values for $\delta(q, x)$, so if $\ell(x) \leqslant r \leqslant u(x), x$ cannot safely be eliminated.

Lemma 1. Let $(\mathrm{X}, \delta)$ be a metric space, with $\mathrm{X}=\left\{p_{1}, \ldots, p_{m}, q, z\right\}$, and let the distances $\delta_{1}, \delta_{2}: \mathrm{X} \times \mathrm{X} \rightarrow \mathbb{R}_{\geqslant 0}$ be defined as follows:

$$
\begin{aligned}
& \delta_{1}(x, y)= \begin{cases}\max _{i}\left|\delta\left(q, p_{i}\right)-\delta\left(p_{i}, z\right)\right| & \text { if }\{x, y\}=\{q, z\} \\
\delta(x, y) & \text { otherwise } .\end{cases} \\
& \delta_{2}(x, y)= \begin{cases}\min _{i} \delta\left(q, p_{i}\right)+\delta\left(p_{i}, z\right) & \text { if }\{x, y\}=\{q, z\} \\
\delta(x, y) & \text { otherwise } .\end{cases}
\end{aligned}
$$

Then $\delta_{1}$ is a pseudometric and $\delta_{2}$ is a metric. If $\delta\left(q, p_{i}\right) \neq \delta\left(p_{i}, z\right)$ for some $i$, or if $q=z$, then $\delta_{1}$ is a metric.

[^2]Proof. We have $\delta_{j}(x, y)=\delta_{j}(y, x)$ and $\delta_{j}(x, x)=0$, for $x, y \in \mathrm{X}, j \in\{1,2\}$. We also have $\delta_{2}(x, y)=0 \Longrightarrow x=y$, and if $\delta\left(q, p_{i}\right) \neq \delta\left(p_{i}, z\right)$ for some $i$, or if $q=0$, then $\delta_{1}(x, y)=0 \Longrightarrow x=y$. We have $\delta_{1}(q, z) \leqslant \delta(q, z)$, so triangularity can only be broken for $\delta_{1}$ in the cases $\delta_{1}\left(q, p_{k}\right) \leqslant \delta_{1}(q, z)+\delta_{1}\left(z, p_{k}\right)$ or $\delta_{1}\left(z, p_{k}\right) \leqslant$ $\delta_{1}(z, q)+\delta_{1}\left(q, p_{k}\right)$, for some $k$. Consider the first of these. We maximize over $i$, so we need only show the following for some choice of $i$ :

$$
\begin{equation*}
\delta\left(q, p_{k}\right) \leqslant\left|\delta\left(q, p_{i}\right)-\delta\left(p_{i}, z\right)\right|+\delta\left(z, p_{k}\right) \tag{2}
\end{equation*}
$$

This is satisfied for $k=i$. The other case is handled symmetrically. For $\delta_{2}$, we have $\delta(q, z) \leqslant \delta_{2}(q, z)$, so triangularity can only be broken in $\delta_{2}(q, z) \leqslant$ $\delta_{2}\left(q, p_{k}\right)+\delta_{2}\left(p_{k}, z\right)$. We minimize over $i$, so this need only hold for some choice of $i$, and again we may choose $i=k$, producing an equation.

Corollary 1. No search method can resolve a metric range query with fewer distance computations than pivoting.

Proof. From lemma 1, we know that after a set of distance computations making pivots $p_{1}, \ldots, p_{m}$ available, the pivoting bounds are tight; if pivoting cannot eliminate an object, no method can safely do so. (Note that an adversary would be free to let $q=z$ in the case where $\delta(q, z)=\ell(z)=0$, ensuring that we are indeed dealing with a metric space.) And given that no method can eliminate more objects than pivoting for any distance count, no method can eliminate all the objects with a lower distance count than pivoting.

In other words, any method using fewer distance computations than pivoting could be made to fail by an adversary in charge of the data set. This argument covers range queries, and it is not hard to translate it to the $k$ NN case, where the $k$ nearest neighbors of $q$ are sought, as long as the result set is uniquely determined. A radius must then exist, separating the $k$ nearest neighbors from the others, and the tightest possible upper bound on this radius is the maximum of the $k$ lowest pivoting bounds we have. Pivoting must then be able to eliminate all points outside this radius, or our adversary might strike again. The following corollary covers the more general case.

Corollary 2. No search method can resolve a metric kNN query with fewer distance computations than pivoting.

Proof. We need to establish $\delta(q, x) \leqslant \delta(q, y)$ for every $x$ in the result and every $y$ outside it. Assume that, given some pivot set P , there is one such inequality that cannot be established by pivoting, i.e., $u(x)>\ell(y)$. An adversary could then ensure $\delta(q, x)>\delta(q, y)$, as follows. First, let $\delta(q, x)=u(x)$. The only effect on the valid range for $\delta(q, y)$ is found in the lower bound $\delta(q, y) \geqslant|u(x)-\delta(x, y)|$. If $\delta(x, y) \leqslant u(x)$, then the relevant lower bound is $u(x)-\delta(x, y)$, which is strictly less than $\delta(q, x)=u(x)$ (because $\delta(x, y)>0$, as $x \neq y$ ), and so it is still possible to have $\delta(q, y)<\delta(q, x)$.

If, however, $\delta(x, y)>u(x)$, the relevant lower bound is $\delta(x, y)-u(x)$. Let $p$ be the pivot that produced the pivoting bound $u(x)$. We then have:

$$
\begin{aligned}
\delta(x, y)-u(x) & =\delta(x, y)-(\delta(q, p)+\delta(p, x)) \\
& =(\delta(x, y)-\delta(p, x))-\delta(q, p) \leqslant \delta(p, y)-\delta(q, p) \leqslant \ell(y)
\end{aligned}
$$

In other words, $\delta(q, y)=\ell(y)$ is still a valid choice for the adversary, yielding the desired $\delta(q, y)<\delta(q, x)$.

The upshot is that the optimal distance count (for range and $k \mathrm{NN}$ queries) can be found by considering only elimination using individual pivots.

The range and $k \mathrm{NN}$ search modes are closely related, and yet there are cases where they behave quite differently, as shown in fig. 1.

(a) Range wins $(k=2)$

(b) $k \mathrm{NN}$ wins $(k=1)$

Fig. 1. Differences between range search and $k N N$ in the presence of ties for the $k$ th position, using $\left(\mathbb{R}^{2}, \mathrm{~L}_{1}\right)$. In both configurations, we have $r=8$. In (a), a range search need only compute $\delta(q, p)$, while $k \mathrm{NN}$ much also compute $\delta\left(q, x_{i}\right)$ for all but one of the $x_{i}$. In (b), the $k \mathrm{NN}$ search need only compute $\delta(q, p)$ and $\delta\left(q, x_{i}\right)$ for one of the $x_{i}$, while a range search must compute all distances $\delta(q,-)$

It is, however, possible to establish some correspondence between the two, when the $k \mathrm{NN}$ result set is uniquely determined.

Lemma 2. If the $k N N$ result is uniquely determined, the optimum number of distance computations for $k N N$ is no worse than for a range search with the smallest possible $k N N$ radius, even if the radius is unknown initially. Furthermore, there is a radius for which range queries and $k N N$ will produce the same search result using the same number of distance computations.

Proof. When the $k \mathrm{NN}$ result is unique, there is a radius $r$ corresponding to the $k$ resulting objects. Resolving a range query with radius $r$ must necessarily yield upper bounds of at most $r$ for returned objects and lower bounds greater than $r$ for the remainder. These same bounds can also be used to separate the $k$ nearest objects from the remainder, without a specified radius, so $k \mathrm{NN}$ cannot require more distance computations.

Conversely, consider a $k$ NN query. By corollary 2 , no method requires fewer distance computations than pivoting, so in the optimal case we will have actual distance bounds available, strictly separating the $k$ nearest from the remainder. Any range query with a search radius falling between the upper and lower bounds can then be resolved with the same number of distance computations.

It is possible to increase the radius such that a range query would require additional distance computations, while still just returning $k$ objects (cf. fig. 2).

## 3 Elimination as Domination

Given a (directed) graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$, a vertex $u$ is said to dominate another vertex $v$ if the graph has an edge from $u$ to $v$. The (directed) minimum dominating set problem involves finding a set $\mathrm{D} \subseteq \mathrm{V}$ of minimum cardinality, such that every vertex $v \in \mathrm{~V} \backslash \mathrm{D}$ is dominated by some vertex $u \in \mathrm{D}$. We call $\gamma(\mathrm{G})=|\mathrm{D}|$ the (directed) domination number of G.

For a given range query, computing the distance to a point may eliminate one or more other points. There are no interactions between such eliminations (see section 2), so an exhaustive listing of the potential eliminations gives us all the relevant information needed to determine which points to examine and which to eliminate. This corresponds to a directed graph - the elimination graph-whose minimum dominating set is the smallest pivot set, and thus the minimum number of distance computations, needed to resolve the query (cf. fig. 3).

Proposition 1. There is a linear-time reduction from the metric range search problem to the directed minimum dominating set problem, which preserves the objective values of the solutions exactly.

If the result of a $k \mathrm{NN}$ query is uniquely determined, and we ignore elimination based on upper bounds (usually done in practice), the number of distance computations correspond to a range query with the smallest $k N N$ radius. ${ }^{5}$ Of course, finding a minimum dominating set is NP-hard, ${ }^{6}$ and given the rather unusual clash between large-scale information retrieval and combinatorial optimization, we may quickly end up with overwhelming instance sizes. Still, with a suitable mixed-integer programming solver, for example, the optimization may very well be feasible in many practical cases. As an example, fig. 4 shows some computations made using the Gurobi solver [12]. Many of these optima were found rather quickly, as presumably the structure of the elimination graph was amenable to the solution methods of the solver. Others, such as those for the DNA data set, took several days to compute. And even for some of the easier cases, there were outliers. For example, for the 2NN radius in 15-dimensional Euclidean space, all of the 10 randomly selected queries led to computations lasting $10-200$ seconds, except for one, which took almost twenty hours. As with many

[^3]

Fig. 2. The nearest neighbor can be determined by examining $x_{1}$ and $p$, as we then have $u\left(x_{1}\right)<\ell\left(x_{2}\right)$. Range search with $r_{1}$ can be resolved similarly, but using $r_{2}$ requires three distance computations, while still returning the single nearest neighbor


Fig. 3. The directed elimination graph G resulting from a specific range query, with the domination number $\gamma(\mathrm{G})=5$ corresponding to the minimum number of distance computations needed to separate relevant objects from irrelevant ones
such cases, however, being satisfied with a solution that is a couple of percentage points shy of perfect could drastically cut down on the computation time (i.e., by setting the absolute or relative MIP gap), as illustrated in fig. 5 .

Figure 4 also includes results for several other methods, beyond the optimum. These are all versions of the AESA approach [25], as discussed in more depth in section 5. At the opposite end of the spectrum of the optimum, there's the incremental random selection of pivots. Separating the feasible from the infeasible, is an oracle AESA, which has access to the elimination power of each potential pivot, i.e., how many of the remaining objects will be eliminated if a given pivot is selected. In the feasible region we find AESA, iAESA2 [10], and the new gAESA, which is explained in section 5 .

It is worth noting that $\gamma(\mathrm{G})$ is a more precise lower bound than an ordinary best-case analysis, which only takes input size into account, and which is therefore always 1. Rather, this is the lowest possible number of distance computations needed for a given dataset and query. In order to guarantee using at most $\gamma(\mathrm{G})$ distance computations, you would need to somehow determine G, which is quite unrealistic. And, as the next section shows, it is also far from enough.

## 4 Metric Search Is Hard, Even If You're Omniscient

Obviously, a major challenge in choosing the right pivots is that you don't know what the elimination graph looks like-you can only make heuristic guesses. But what if you did know? As it turns out, that wouldn't be the end of your worries.

Section 3 showed that it is possible to find the optimum by framing the problem as that of looking for a minimum directed dominating set. Of course, this is an NP-hard problem, so there's no real surprise in that we can reduce to it. But what about reducing in the other direction? That is, unless $P=N P$, is there any hope of finding some feasible way of determining the optimum? Alas, no: reducing from the general minimum undirected dominating set problem to


Fig. 4. Number of distance computations as a function of $k$, the number of nearest neighbors covered by the chosen radius used for a range query. The first four datasets are uniformly random vectors, while the last two are taken from the SISAP dataset collection [9], with queries withheld. The listeria string lengths vary from 39 to 6579. The results are the average over 10 randomly selected queries. The oracle AESA uses elimination power among remaining points as its heuristic

Fig. 5. Bound on relative error (MIP gap) as a function of time, when computing the optimal number of distance computations in a particularly difficult instance with $k=2$ over uniformly random vectors in 15 -dimensional Euclidean space. Finding the optimum took over nineteen hours. After 41 s , the gap was $39.1 \%$, but already at 44 s , it was down to $4.34 \%$. Getting to $1 \%$ took 2.21 h

finding the optimum for metric search is quite straightforward, and the reduction preserves the both the objective value and the problem size exactly, ${ }^{7}$ meaning that approximation hardness results apply as well.

Theorem 1. There is a linear-time reduction from the undirected minimum dominating set problem on $n$ vertices to the metric range search problem on $n$ objects, which preserves the objective values of the solutions exactly.

Proof. We first consider range search. To encode any instance $G=(\mathrm{V}, \mathrm{E})$ of the minimum dominating set problem, we construct a metric space ( $\mathrm{X}, \delta$ ), where $\mathrm{X}=\mathrm{V} \cup\{q\}$, with $q \notin \mathrm{~V}$, and design the metric so that the elimination graph corresponds to G. We define the metric as follows:

$$
\delta(x, y)= \begin{cases}0 & \text { if } x=y \\ 1 & \text { if }\{x, y\} \in \mathrm{E} \\ 2 & \text { otherwise }\end{cases}
$$

In particular, $\delta(q, x)=2$ for all $x \in \mathrm{~V}$. This definition of $\delta$ satisfies all the metric properties. Specifically, note that triangularity holds, because for any objects $x, y, z \in \mathrm{X}$, we have $\delta(x, z) \leqslant 2 \leqslant \delta(x, y)+\delta(y, z)$ (assuming $x \neq y \neq z$; otherwise triangularity is trivial).

It should be clear that the elimination graph for $q$ with $r<1$ corresponds exactly to the original graph $\mathrm{G} .{ }^{8}$ The closed neighborhood $\mathrm{N}[x]$ of $x$ (that is, $x$ and the set of objects dominated or eliminated by $x)$ is $\{y: \delta(q, x)-\delta(x, y) \geqslant 1\}$, which corresponds exactly to the cases where $\delta(x, y)=0$ (that is, $x=y$ ) and where $\delta(x, y)=1$ (that is, $\{x, y\} \in \mathrm{E}$ ). In other words, any set of pivots that eliminate the remaining objects corresponds to a dominating set in G, and vice versa. If we find such a pivot set of minimum cardinality, we will have solved the undirected minimum dominating set problem. In other words, we have a valid reduction from the undirected dominating set problem to metric range search. It should also be obvious that the reduction can be performed in linear time, and that the size of the optimal solutions are identical. ${ }^{9}$

The previous reduction can be extended to a polytime reduction to $k \mathrm{NN}$ search quite easily, showing NP-hardness (though not necessarily preserving approximation results). We simply set $k=1$ and add another object $\bar{x}$ so that $\delta(q, \bar{x})=$ $r<1$ and $\delta(\bar{x}, y)=2$ for any other object $y$. Now the minimum $k \mathrm{NN}$ radius will automatically be $r$, which gives us the same reduction as before.

The reduction in the proof of theorem 1 constructs a metric range search problem on $n$ objects from an undirected dominating set problem on $n$ nodes, so that if (and only if) we can solve the search problem (that is, find a minimum pivot set), we have also solved the minimum dominating set problem.

[^4]Approximation bounds thus carry over from the dominating set problem, so for any $\epsilon>0$, finding solutions that are within a factor of $(1-\epsilon) \ln n$ is unfeasible, unless NP $\subseteq \operatorname{DTIME}\left(n^{\mathrm{O}(\lg \lg n)}\right)[6]$.

Corollary 3. For instances of the metric range search problem over $n$ objects where the optimal number of distance computations is $\gamma$, the worst-case running time of any algorithm is $\Omega(\gamma \log n)$, unless $\mathrm{NP} \subseteq \operatorname{DTIME}\left(n^{\mathrm{O}(\lg \lg n)}\right)$.

Proof. An algorithm with a (polynomial) running time of $o(\gamma \log n)$ would necessarily use $o(\gamma \log n)$ distance computations, yielding an approximation algorithm for the dominating set problem with an approximation ratio $o(\log n)$.

Note that the worst-case running time in general is still $\Omega(n)$, as we may very well have $\gamma=n$, in degenerate workloads.

## 5 Omniscience Is Overrated

In the discussion so far, what has been described is a scenario where all potential eliminations are known. Even then, as we have seen, it is only realistically feasible to get to within a log-factor of the optimum. And as it turns out, achieving this log-factor is possible, even without knowing all the potential eliminations. What is assumed instead is a more limited oracle that can tell us which of the remaining points has the highest elimination power, that is, the highest out-degree among the remaining vertices.

The thing is, a minimum dominating set may be approximated to within a log-factor using a simple greedy strategy - a strategy that most likely cannot be significantly improved upon; it gets within a factor of $\ln n+1$, and as discussed in the previous section, we have a lower bound of $(1-\epsilon) \ln n$ for any $\epsilon>0 .{ }^{10}$

What is more, this is exactly the approach taken by the AESA family of indexing methods: they greedily pick one point at a time, based on estimated elimination power, eliminating others as they go (cf. fig. 6). In other words, full omniscience wrt. the elimination graph is not needed; if we can formulate a heuristic returning the most useful next pivot at each step, the algorithm is already as good as it realistically can be, or at least very nearly so.

Proposition 2. Greedily selecting pivots based on high elimination power is an asymptotically optimal polytime strategy for minimizing distance computations in metric range search, unless $\mathrm{NP} \subseteq \operatorname{DTIME}\left(n^{\mathrm{O}(\lg \lg n)}\right)$.

To say that AESA picks pivots based on elimination power may be overstating it, however. Rather, Vidal Ruiz talks about "successive approximation to nearest points" [25], while Figueroa et al. state that their goal is "to define an order such that the first element is very close to the query," because " $[t]$ he closer

[^5]$\operatorname{AESA}(q, r ; \mathrm{V}, \delta)$
Greedy-Dom-Set(V, E)
$\mathrm{U} \leftarrow \mathrm{V}$
$\mathrm{U} \leftarrow \mathrm{V}$
$P=\emptyset ; R=\emptyset$
while $U \neq \emptyset$
$\mathrm{D} \leftarrow \emptyset$
while $U \neq \emptyset$
while $U \neq \emptyset$
$p \leftarrow \arg \min _{x \in \mathrm{U}} h_{\mathrm{P}}(x)$
$p \leftarrow \arg \max _{x \in \mathrm{U}}\left|\mathrm{N}^{+}(x) \cap \mathrm{U}\right|$
$\mathrm{D} \leftarrow \mathrm{D} \cup\{p\}$
if $\delta(q, p) \leqslant r: \mathrm{R} \leftarrow \mathrm{R} \cup\{p\}$
$\mathrm{P} \leftarrow \mathrm{P} \cup\{p\}$
$\mathrm{U} \leftarrow \mathrm{U} \backslash \mathrm{N}^{+}[p]$
$\mathrm{U} \leftarrow \mathrm{U} \backslash\left(\{p\} \cup\left\{x: \ell_{\mathrm{P}}(x)>r\right\}\right)$
return D

Fig. 6. Side-by-side comparison of the AESA metric search algorithm and the greedy approximation for the directed minimum dominating set problem
the pivot to the query $q$, the more effective the pruning is" [10]. Of course, all manner of regression and learning methods might be used with the specific goal of estimating which points are close to the query $[8,16]$, or which are likely to be part of the search result [17].

There has been work on pivot selection focusing directly on elimination power [4], but this does not seem to have been central in AESA-like methods, using a full distance matrix. One selection method, which maximizes the lower bound used for elimination, and skips over pivots that don't contribute, has been explored in the fixed, initial pivot list of PiAESA [22], but the second phase, where pivots are selected dynamically, still follows the heuristic of selecting those that seem close to the query.

Following the analogy with the greedy approximation for the directed dominating set problem, there are two modifications one might make. The first is to look for high elimination power in the data set overall, rather than closeness to the query. For example, it is quite possible that a pivot that is far away might be able to eliminate an entire nearby cluster. The second modification, which I will briefly explore, is to modify the selection based on redundancy, i.e., how much of a point's elimination power actually applies to remaining points. If one selects pivots that are as similar to the query as possible, they are bound to be similar to each other as well; and even if a pivot is able to eliminate many other points, that is of little use if those points have already been discarded.

A simple version of this second modification is the following: rather than merely minimizing the sum of lower bounds, as in AESA, we divide this by the sum of distances to remaining points. This will not only prefer pivots that seem to be close to $q$, but those that seem close to $q$ relative to how far they are from the remaining points, meaning they ought to be able to eliminate more of them. Some preliminary results on the performance of this greedy AESA (gAESA) are shown in fig. 4. As can be seen, it does seem to perform on par with AESA and iAESA2, at times outperforming both. Given the rather arbitrary nature of the heuristic, better variants might very well exist.

## 6 Concluding Remarks and Future Work

The previous sections have established an equivalence between the minimal number of distance computations needed to resolve an exact metric range query, on the one hand, and the size of a minimum dominating set in a directed graph on the other. ${ }^{11}$ The result also applies to uniquely determined $k$ NN queries, if upper bounds are ignored. One might object that the scenario is too limited-that in practice, one would be contented with an approximate or probabilistic search. In fact, the results do also apply for certain approximations, such as those that merely modify the query, resulting in a new, simpler exact search [18]. But beyond this, the main uses of these results are precisely in establishing the limits of exact search for given workloads; if one can show that any exact algorithm must examine an excessively large portion of the data set, that is a forceful argument in favor of approximation or randomization. What is presented here only scratches the surface, however. What follows is a sketch of possible directions for future research based on the established equivalence.

Heuristic development. The gAESA heuristic is somewhat arbitrary. While it picks pivots that seem close to the query, relative to the remaining points, the goal is to pick the pivot with the highest elimination power. There may be many ways of estimating this more directly, either using hand-crafted heuristics (e.g., including pivots that are far away from the query compared to remaining points) or machine learning (which has so far been focused on distance or relevance).

Algorithm development. In the interest of constructing better baselines, one might take the development further. Rather than going with the AESA approach, one might attempt to solve the dominating set problem without actually knowing the graph. This would be different from the more common forms of online dominating set problems [3], where vertices are provided in some arbitrary order. Rather, this would presumably involve link prediction [15], at each step selecting a pivot deemed likely to be included in the optimal solution or to provide good support for future predictions.

Problem variants. The dominating set problem provides a new perspective on the problem of metric search, and variants of the former might find analogies for the latter. For example, the weighted dominating set problem can also be approximated greedily, and the analogous metric search method would be a weighted AESA, where selection is based on the ratio of weight to elimination power. The weight could, for example, represent the actual cost of computing the query-pivot distance, which is the effort that is being minimized, after all. For many distances, this cost is identical for all points, but for, e.g., the signature quadratic form distance [2], it may vary wildly.

One might also look for analogies in the other direction. For example, probabilistic methods (such as probabilistic iAESA [10]) do not aim to eliminate all vertices; in these cases, one could instead consider partial domination [7].

[^6]Probabilistic analysis. There is a substantial literature on the topic of random graphs. For example, it is known that for random digraphs whose edges are independent Bernoulli variables with probability $p,{ }^{12}$ the domination number is logarithmic, with base $1 /(1-p)$ [14]. In fact, it is not hard to modify the results of Telelis and Zissimopoulos [23] to show that in this scenario, even AESA selecting pivots arbitrarily would yield a logarithmic number of pivots, staying within a doubly logarithmic additive term of the optimum, results that match those of Navarro [19].

Workload descriptions. Beyond finding $\gamma$, the dominating set perspective may inspire other hardness measures and workload descriptions. For example, the greedy approximation is, more precisely, logarithmic in the maximum degree $\Delta(\mathrm{G})$, a value that could be used as an indicator of how hard it is to get close to the optimum. And although the independence assumption on elimination may be too strong, one could still use the elimination probability $p$, perhaps estimated by averaging over several queries, as an indication of general workload hardness.

Acknowledgements. The author would like to thank Ole Edsberg, both for discussions providing the initial idea for this paper, and for substantial later input. He would also like to thank Jon Marius Venstad and Bilegsaikhan Naidan for reading early drafts of the paper and providing feedback.

## References

1. Backurs, A., Indyk, P.: Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In: Proceedings of the forty-seventh annual ACM symposium on Theory of computing (2015). doi:10.1145/2746539.2746612
2. Beecks, C., Uysal, M.S., Seidl, T.: Signature quadratic form distance. In: Proceedings of the ACM International Conference on Image and Video Retrieval. ACM, New York, NY, USA (2010). doi:10.1145/1816041.1816105
3. Boyar, J., Eidenbenz, S.J., Favrholdt, L.M., Kotrbčík, M., Larsen, K.S.: Online dominating set. Algorithmica 81(5) (2019). doi:10.1007/s00453-018-0519-1
4. Bustos, B., Navarro, G., Chávez, E.: Pivot selection techniques for proximity searching in metric spaces. Pattern Recognition Letters 24(14) (2003). doi: 10.1016/S0167-8655(03)00065-5
5. Chávez, E., Navarro, G., Baeza-Yates, R., Marroquín, J.L.: Searching in metric spaces. ACM Computing Surveys 33(3) (2001). doi:10.1145/502807.502808
6. Chlebík, M., Chlebíkova, J.: Approximation hardness of dominating set problems. In: European Symposium on Algorithms. Springer (2004). doi:10.1007/978-3-540-30140-0_19
7. Das, A.: Partial domination in graphs. Iranian Journal of Science and Technology, Transactions A: Science 43(4) (2019). doi:10.1007/s40995-018-0618-5
8. Edsberg, O., Hetland, M.L.: Indexing inexact proximity search with distance regression in pivot space. In: Proceedings of the Third International Conference on Similarity Search and Applications (2010). doi:10.1145/1862344.1862353

[^7]9. Figueroa, K., Navarro, G., Chávez, E.: Metric spaces library (2007), available at http://www.sisap.org/Metric_Space_Library.html
10. Figueroa, K., Chávez, E., Navarro, G., Paredes, R.: Speeding up spatial approximation search in metric spaces. Journal of Experimental Algorithmics 14 (2010). doi:10.1145/1498698.1564506
11. Ford Jr, L.R., Johnson, S.M.: A tournament problem. The American Mathematical Monthly 66(5) (1959). doi:10.1080/00029890.1959.11989306
12. Gurobi Optimization, LLC: Gurobi optimizer reference manual (2020), available at http://gurobi.com
13. Kann, V.: On the Approximability of NP-complete Optimization Problems. Ph.D. thesis, Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm (1992)
14. Lee, C.: Domination in digraphs. Journal of the Korean Mathematical Society 35(4) (1998)
15. Lü, L., Zhou, T.: Link prediction in complex networks: A survey. Physica A: statistical mechanics and its applications $\mathbf{3 9 0}(6), 1150-1170$ (2011). doi:10.1016/j. physa.2010.11.027
16. Mao, R., Liu, X., Tang, H., Luo, Q., Chen, J., Wu, W.: Multivariate regression for pivot selection: A preliminary study. In: 2011 3rd Symposium on Web Society. IEEE (2011). doi:10.1109/SWS.2011.6101281
17. Murakami, T., Takahashi, K., Serita, S., Fujii, Y.: Probabilistic enhancement of approximate indexing in metric spaces. Information Systems 38(7) (2013). doi: 10.1016/j.is.2012.05.012
18. Naidan, B., Hetland, M.L.: Shrinking data balls in metric indexes. DBKDA (2013)
19. Navarro, G.: Analyzing metric space indexes: what for? In: Proceedings of the 2009 Second International Workshop on Similarity Search and Applications. SISAP '09, IEEE Computer Society (2009). doi:10.1109/SISAP.2009.17
20. Pestov, V.: Lower bounds on performance of metric tree indexing schemes for exact similarity search in high dimensions. Algorithmica (2013). doi:10.1007/s00453-012-9638-2
21. Skopal, T.: Unified framework for exact and approximate search in dissimilarity spaces. ACM Transactions on Database Systems, TODS 32(4) (2007). doi:10.1145/ 1292609.1292619
22. Socorro, R., Micó, L., Oncina, J.: A fast pivot-based indexing algorithm for metric spaces. Pattern Recognition Letters 32(11) (2011). doi:10.1016/j.patrec.2011.04. 016
23. Telelis, O.A., Zissimopoulos, V.: Absolute $o(\log m)$ error in approximating random set covering: an average case analysis. Information Processing Letters 94(4) (2005). doi:10.1016/j.ipl.2005.02.009
24. Traina Jr., C.: Distance exponent: A new concept for selectivity estimation in metric trees. In: Proceedings of the 16th International Conference on Data Engineering (2000). doi:10.1109/ICDE.2000.839409
25. Vidal Ruiz, E.: An algorithm for finding nearest neighbours in (approximately) constant average time. Pattern Recognition Letters 4(3) (1986). doi:10.1016/0167-8655(86)90013-9
26. Williamson, D.P., Shmoys, D.B.: The Design of Approximation Algorithms. Cambridge university press (2011)


[^0]:    ${ }^{1}$ Other measures include the distance exponent [24] and the ball-overlap factor [21].

[^1]:    ${ }^{2}$ Personal communication, July 2012
    ${ }^{3}$ Note that the reductions are to and from two different versions of the dominating set problem (the directed and undirected version, respectively). At the price of slightly looser bounds, one could stick with just one of these.

[^2]:    ${ }^{4}$ This is the worst case given that the optimal number of distance computations is some value $\gamma$, not the more general, non-informative worst-case of $\Omega(n)$.

[^3]:    ${ }^{5}$ Optimal $k \mathrm{NN}$ with upper bounds does not map as cleanly to dominating sets.
    ${ }^{6}$ The undirected version is most commonly discussed, with a reduction, e.g., from set covering [13, Th. A.1]. A similar reduction to the directed version is straightforward.

[^4]:    ${ }^{7}$ In terms of vertices, not edges.
    ${ }^{8}$ Note that only the lower bound is relevant, as the upper bound is always greater than the search radius.
    ${ }^{9}$ If the new distance is allowed to use the original graph as part of its definition, the reduction can be performed in constant time-it is merely a reinterpretation.

[^5]:    ${ }^{10}$ The upper bound is easily shown by reinterpreting the minimum dominating set problem for a directed graph $\mathrm{G}=(\mathrm{V}, \mathrm{E})$ as the problem of covering V with the closed out-neighborhoods of G, translating the standard set covering approximation [26].

[^6]:    ${ }^{11}$ That is, for any range search instance, there is a directed graph with the objects as its nodes for which the equivalence holds. Reducing in the other direction preserves the objective value, but not necessarily the number of nodes/objects.

[^7]:    ${ }^{12}$ Chávez et al. say that such independence is a "reasonable approximation" [5].

