

The Core in Max-Loss Non-Centroid Clustering Can Be Empty

Robert Brederick¹, Eva Deltl¹, Leon Kellerhals¹, and Jannik Peters²

¹TU Clausthal, Germany

²National University of Singapore, Singapore

Abstract

We study core stability in non-centroid clustering under the max-loss objective, where each agent’s loss is the maximum distance to other members of their cluster. We prove that for all $k \geq 3$ there exist metric instances with $n \geq 9$ agents, with n divisible by k , for which no clustering lies in the α -core for any $\alpha < 2^{\frac{1}{k}} \sim 1.148$. The bound is tight for our construction. Using a computer-aided proof, we also identify a two-dimensional Euclidean point set whose associated lower bound is slightly smaller than that of our general construction. This is, to our knowledge, the first impossibility result showing that the core can be empty in non-centroid clustering under the max-loss objective.

1 Introduction

Clustering is a fundamental task in data analysis, optimization, and beyond: we are given n points (or agents) with the goal to partition these points into k groups, the so-called *clusters*. Most works in clustering focus on the *centroid model*. In this model, each cluster has a representative center. For a given point, the *loss* of this point is its distance to the nearest center or equivalently to the center assigned to its partition. Recent work in clustering has introduced and studied many variations of *proportional fairness* [Chen, Fain, Lyu, and Munagala, 2019, Micha and Shah, 2020, Aziz, Lee, Chu, and Vollen, 2024, Kellerhals and Peters, 2024], where a clustering is considered proportional if there is no coalition of at least $\frac{n}{k}$ agents that can agree on a new representative which strictly reduces the loss of each coalition member.

The above-mentioned work solely focuses on centroid clustering which always selects a representative for each of the clusters. However, there are also settings where clusters do not have natural representatives, for example in settings of team formation [Hajduková, 2006, Woeginger, 2013] or clustered federated learning [Sattler, Müller, and Samek, 2021]. In these *non-centroid* (representative-free) settings, an agent’s loss depends only on distances to other members of its cluster. Caragiannis, Micha, and Shah [2024] extended proportional fairness to this setting: A (non-centroid) clustering is in the α -core if no coalition of at least $\frac{n}{k}$ agents can form a new cluster in which each member’s loss decreases by a factor of more than α .

Among other possible definitions of loss, Caragiannis et al. [2024] consider the *max-loss objective*, where each agent’s loss in a cluster is the maximum distance to another agent in the cluster. For this objective, they show that there always exists a clustering which is in the 2-core. Moreover, they state the following open problem.

Question 1 (Caragiannis et al. [2024]). For the max-loss objective, does there always exist a clustering in the 1-core?

The same question was repeated by Cookson, Shah, and Yu [2024, page 8] (who studied how to combine both proportional centroid and non-centroid clustering).

As our main contribution, we answer the question in the negative.

Theorem 1. *For every $k \geq 3$ there exists a metric space with $n \geq 9$ agents, where n is divisible by k , for which no k -clustering is in the α -core under the max-loss objective, for any $\alpha < 2^{\frac{1}{5}} < 1.149$.*

The metric space we use to prove Theorem 1 is not Euclidean, whereas most metric spaces considered in the clustering setting are. Using a computer-aided proof, we were able to find a set of points in two-dimensional Euclidean space which yield a slightly smaller lower bound. These points were obtained by optimizing a system of inequalities and are based on the metric space used for proving Theorem 1.

Theorem 2. *For $k = 3$ there is a two-dimensional Euclidean space with $n = 9$ many agents for which no k -clustering is in the α -core under the max-loss objective, for any $\alpha < 1.138$.*

The points for the counterexample, along with a Python script for verifying the claim, can be found in Appendix A.

Notation. Let (N, d) be a (pseudo)-metric space with a distance function $d: N \times N \rightarrow \mathbb{R}_{\geq 0}$ satisfying $d(i, i) = 0$, $d(i, j) = d(j, i)$ and $d(i, j) + d(j, h) \geq d(i, h)$ for all $i, j, h \in N$. We call a metric space (N, d) t -dimensional Euclidean if $N \subseteq \mathbb{R}^t$ and if d is the Euclidean distance function, i.e., $d(i, j) = \sqrt{(i - j)^2}$.

A (non-centroid) k -clustering of a metric space (N, d) of n agents is a partition $\mathcal{C} = (C_1, \dots, C_k)$ of N into k (nonempty) clusters. Denote by $\mathcal{C}(i)$ the cluster including i . For $i \in N$ and a set $S \subseteq N$, the *individual max-loss* of i in S is

$$\text{loss}_i(S) := \max_{y \in S} d(i, y).$$

For a clustering $\mathcal{C}$, we write $\text{loss}_i(\mathcal{C}) := \text{loss}_i(\mathcal{C}(i))$.

We say that a coalition $S \subseteq N$ of agents α -blocks a clustering $\mathcal{C}$ if $\alpha \cdot \text{loss}_i(S) < \text{loss}_i(\mathcal{C})$ for all $i \in S$. A k -clustering $\mathcal{C}$ is in the α -core for $\alpha \geq 1$ if there is no α -blocking coalition of size at least $\frac{n}{k}$. We refer to the 1-core simply as the core.

2 Proof of Theorem 1

Fix $k = 3$ and let $n \geq 9$ divisible by 3. We create three sets of agents, G_1 , G_2 , and G_3 , which we will henceforth call *groups*, with

$$|G_1| = |G_3| = \frac{n}{k} - 1, \quad |G_2| = \frac{n}{k} - 2,$$

and $d(i, j) = 0$ for each $i, j \in G_t$ and each $t \in \{1, 2, 3\}$. The remaining four agents will be called a, b, c , and w . We have the following distances for all agent-group pairs (as illustrated in Figure 1).

	G_1	G_2	G_3
a	$2^{\frac{1}{5}}$	1	$(1 + 2^{\frac{4}{5}})$
b	$(1 + 2^{\frac{1}{5}})$	2	$2^{\frac{4}{5}}$
c	$2^{\frac{2}{5}}$	$(1 + 2^{\frac{3}{5}})$	$2^{\frac{3}{5}}$

Moreover, we have $d(a, b) = 1$, and agent w is very far away from all other agents, that is, $d(w, i) = M$ for all $i \neq w$ and a sufficiently large $M \gg 1$. The remaining undefined distances can be chosen as the lengths of shortest paths on the weighted graph induced by the specified edges.

In our proof, we focus on the five following coalitions of size $\frac{n}{k}$:

$$\begin{aligned} S_1 &:= G_1 \cup \{a\}, & S_2 &:= G_1 \cup \{c\}, & S_3 &:= G_2 \cup \{a, b\}, \\ S_4 &:= G_3 \cup \{b\}, & S_5 &:= G_3 \cup \{c\}. \end{aligned}$$




Figure 1: Counterexample for non-centroid max-loss: edge labels specify their length; the distance between any two points is the length of the shortest path.

For the sets S_1 , S_2 , S_4 and S_5 , and each agent i within that set, we have

$$\text{loss}_i(S_1) = 2^{\frac{1}{5}}, \text{loss}_i(S_2) = 2^{\frac{2}{5}}, \text{loss}_i(S_4) = 2^{\frac{4}{5}}, \text{ and } \text{loss}_i(S_5) = 2^{\frac{3}{5}}.$$

For S_3 and each $i \in S_3 \setminus \{a\}$ we have $\text{loss}_i(S_3) = 2$ while the loss of agent a is $\text{loss}_a(S_3) = 1$.

Let $\mathcal{C} = (C_1, C_2, C_3)$ be any 3-clustering of the gadget. One cluster must contain w , and any such cluster has a loss of M for all its members, unless w is placed alone as a singleton. Observe that the cluster containing w contains at most $\frac{n}{k}$ members, otherwise the cluster's agents without w form a blocking coalition for an α -core, with $\alpha \geq \frac{M}{2}$.

We distinguish between two cases, namely whether or not $\mathcal{C}$ contains a *mixed cluster*. Here, we call a cluster in $\mathcal{C}$ mixed if it contains agents from more than one of the three groups G_1, G_2, G_3 .

Case 1: There is no mixed cluster. We claim that, if $\mathcal{C}$ has a separate cluster for each of the groups G_1, G_2, G_3 , then one of $S_1, \dots, S_5$ is an α -blocking coalition.

If each group lies in its own cluster, then one of the three must also contain w . As mentioned above, we assume without loss of generality the cluster containing w has size at most $\frac{n}{k}$.

- G_1 with w . Then coalition $S_2 = G_1 \cup \{c\}$ α -blocks $\mathcal{C}$: Clearly, the loss of each agent $i \in G_1$ in S_2 is smaller than in $\mathcal{C}$ by a factor greater than α . Since $|\mathcal{C}(w)| \leq \frac{n}{k}$, agent c must be clustered with agents from G_2 or G_3 ; thus its loss is

$$\text{loss}_c(\mathcal{C}) \geq 2^{\frac{3}{5}} > \alpha 2^{\frac{2}{5}} = \alpha \text{loss}_c(S_2).$$

- G_3 with w . Suppose first that b is clustered with G_1 . Then due to our above assumption of having no mixed cluster, the third cluster contains G_2 . For c , this implies that it is clustered either with b or with G_2 , thus

$$\text{loss}_c(\mathcal{C}) \geq 2^{\frac{3}{5}} > \alpha 2^{\frac{2}{5}} = \alpha \text{loss}_c(S_2).$$

And for all $i \in G_1$

$$\text{loss}_i(\mathcal{C}) \geq d(b, i) > 2^{\frac{3}{5}} > \alpha 2^{\frac{2}{5}} = \alpha \text{loss}_i(S_2).$$

So $S_2 = G_1 \cup \{c\}$ α -blocks $\mathcal{C}$.

If b is clustered with G_2 , then we claim that $S_4 = G_3 \cup \{b\}$ α -blocks $\mathcal{C}$. Again this is obvious for the members of G_3 , and for b , we have

$$\text{loss}_b(\mathcal{C}) \geq 2 > \alpha 2^{\frac{4}{5}} = \alpha \text{loss}_b(S_4).$$

- G_2 with w .

- If b is clustered with G_2 and w , then $S_3 = G_2 \cup \{a, b\}$ α -blocks $\mathcal{C}$. The improvement is clear for all agents clustered with w . For agent a we have $\text{loss}_a(\mathcal{C}) \geq 2^{\frac{1}{5}} > \alpha = \alpha \text{loss}_a(S_3)$.
- If b is clustered with G_1 , leading to a loss of at least $(1 + 2^{\frac{1}{5}})$ for all its members, then $S_2 = G_1 \cup \{c\}$ α -blocks, as $\text{loss}_c(\mathcal{C}) \geq 2^{\frac{3}{5}} > \alpha 2^{\frac{2}{5}} = \alpha \text{loss}_c(S_2)$.
- Let b be clustered with G_3 . If a and c are clustered with G_1 , then $S_1 = G_1 \cup \{a\}$ α -blocks as $\text{loss}_i(\mathcal{C}) \geq 2^{\frac{2}{5}} > \alpha \cdot 2^{\frac{1}{5}} = \text{loss}_i(S_1)$ for all $i \in S_1$. Otherwise, one of a and c is not clustered with G_1 . If c is with $G_3 \cup \{b\}$ or with $G_2 \cup \{w\}$, then $S_5 = G_3 \cup \{c\}$ α -blocks $\mathcal{C}$ as $\text{loss}_i(\mathcal{C}) \geq 2^{\frac{4}{5}} > \alpha \cdot 2^{\frac{3}{5}} = \alpha \text{loss}_i(S_5)$ for all $i \in S_5$. If instead a is clustered with $G_3 \cup \{b\}$ or with $G_2 \cup \{w\}$, then $S_1 = G_1 \cup \{a\}$ α -blocks $\mathcal{C}$ as $\text{loss}_i(\mathcal{C}) \geq 2^{\frac{2}{5}} > \alpha \cdot 2^{\frac{1}{5}} = \alpha \text{loss}_i(S_1)$ for all $i \in S_1$.

In all cases, we find an α -blocking coalition for some $\alpha \geq 2^{\frac{1}{5}}$.

Case 2: There is a mixed cluster. We claim that one of $S_1, \dots, S_4$ α -blocks $\mathcal{C}$.

Suppose that there is only one mixed cluster and call it C_1 . As the other two clusters are not mixed, C_1 must contain all members of one of the groups. Assume the mixed cluster contains w . Since the cluster containing w cannot have more than $\frac{n}{k}$ agents, C_1 must contain the agents of G_2 , and one other agent x from G_1 or G_3 . Now consider the clustering $\mathcal{C}'$ where x is moved to the cluster containing the remaining members of its group. Then $\mathcal{C}'$ does not contain a mixed cluster and therefore, by Case 1, it is not in the α -core. Now observe that for every agent $i \in N$ we have $\text{loss}_i(\mathcal{C}') \leq \text{loss}_i(\mathcal{C})$; thus if a coalition S α -blocks $\mathcal{C}'$, then it also α -blocks $\mathcal{C}$. Consequently, $\mathcal{C}$ is also not in the α -core.

Hence, we may assume our clustering $\mathcal{C}$ contains a mixed cluster that does not contain w . In any such clustering $\mathcal{C}$, every agent in a mixed cluster has loss at least $1 + 2^{\frac{1}{5}}$, and every agent in the cluster containing w has loss at least $M \gg (1 + 2^{\frac{1}{5}})$. If there is an agent with a smaller loss, then it lies in the third cluster, which then is not mixed. It follows that there are at least two groups whose members all incur a loss of at least $1 + 2^{\frac{1}{5}}$ in $\mathcal{C}$: One group in the cluster containing w and at least one group in the mixed cluster. We call such groups *deviating*.

- Suppose that G_1 is deviating. If a and c are part of the same cluster, then $S_1 = G_1 \cup \{a\}$ is α -blocking: For each $i \in G_1$, we have

$$\alpha \cdot \text{loss}_i(S_1) = \alpha \cdot 2^{\frac{1}{5}} < 2^{\frac{2}{5}} < 1 + 2^{\frac{1}{5}} \leq \text{loss}_i(\mathcal{C}),$$

and for a we have

$$\alpha \cdot \text{loss}_a(S_1) = \alpha \cdot 2^{\frac{1}{5}} < 2^{\frac{2}{5}} < 2^{\frac{1}{5}} + 2^{\frac{2}{5}} = d(a, c) \leq \text{loss}_a(\mathcal{C}).$$

Similarly, if b and c are part of the same cluster, then $S_2 = G_1 \cup \{c\}$ is α -blocking as for each $i \in G_1$,

$$\alpha \cdot \text{loss}_i(S_2) = \alpha \cdot 2^{\frac{2}{5}} < 2^{\frac{1}{5}} \cdot 2^{\frac{2}{5}} \leq 1 + 2^{\frac{1}{5}} = \text{loss}_i(\mathcal{C}),$$

and for c we have

$$\alpha \cdot \text{loss}_c(S_2) = \alpha \cdot 2^{\frac{2}{5}} < 2^{\frac{3}{5}} + 2^{\frac{4}{5}} = d(b, c).$$

Finally, if a or c are clustered with w or with G_3 , then again S_1 or S_2 α -blocks, because both a and c face losses exceeding $2^{\frac{1}{5}}$ or $2^{\frac{2}{5}}$, respectively, i.e., by a factor strictly larger than $2^{\frac{1}{5}}$.

- If G_1 is not among the deviating groups, then b cannot be clustered with G_1 (otherwise G_1 would deviate, as $d(b, G_1) = 1 + 2^{\frac{1}{5}}$). Hence b is either clustered with w or in a mixed cluster, and in both cases its loss exceeds 2. Thus $S_4 = G_3 \cup \{b\}$ is α -blocking: For each $i \in S_4$, we have

$$\alpha \cdot \text{loss}_i(S_4) = \alpha \cdot 2^{\frac{4}{5}} < 2 \leq \text{loss}_i(\mathcal{C}).$$

Thus in every clustering where at least one cluster is mixed, one of the coalitions $S_1, \dots, S_4$ strictly blocks by a factor of at least $2^{\frac{1}{5}}$, contradicting α -core stability.

Finally, to extend to arbitrary $k \geq 3$ and n divisible by k , add $(k - 3)$ dummy clusters of size $\frac{n}{k}$ each, with intra-distance 0 and distance M to every outside node. In any feasible k -clustering, these dummy clusters must stand alone, so the effective choice reduces to clustering the gadget $\{a, b, c, w, G_1, G_2, G_3\}$.

3 Outlook

A couple of questions emerge from the two results. Can we improve upon the upper bound of $\alpha = 2$ for the core guarantee? Also, can we possibly prove a better core guarantee for Euclidean spaces than for metric spaces, as is the case in proportional (centroid) clustering setting [Chen et al., 2019, Micha and Shah, 2020]?

References

- H. Aziz, B. E. Lee, S. M. Chu, and J. Vollen. Proportionally representative clustering. In *Proceedings of the 20th International Conference on Web and Internet Economics (WINE)*, 2024. URL http://www.cse.unsw.edu.au/~haziz/faircluster_report.pdf. Forthcoming.
- I. Caragiannis, E. Micha, and N. Shah. Proportional fairness in non-centroid clustering. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*, volume 37, pages 19139–19166, 2024. URL https://papers.nips.cc/paper_files/paper/2024/hash/220cbc7435d6a56205c87d73d15d9eda-Abstract-Conference.html.
- X. Chen, B. Fain, L. Lyu, and K. Munagala. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 1032–1041, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/chen19d.html>.
- B. Cookson, N. Shah, and Z. Yu. Unifying proportional fairness in centroid and non-centroid clustering. In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS)*, 2024. URL <https://www.cs.toronto.edu/~nisarg/papers/fair-semicentroid-clustering.pdf>. Forthcoming.
- J. Hajduková. Coalition formation games: A survey. *International Game Theory Review*, 08(04): 613–641, 2006. URL <https://doi.org/10.1142/S0219198906001144>.
- L. Kellerhals and J. Peters. Proportional fairness in clustering: A social choice perspective. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*, pages 111299–11317, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/c981fd12b1d5703f19bd8289da9fc996-Abstract-Conference.html.
- E. Micha and N. Shah. Proportionally Fair Clustering Revisited. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168, pages 85:1–85:16, 2020. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2020.85>.
- F. Sattler, K. Müller, and W. Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Trans. Neural Networks Learn. Syst.*, 32(8):3710–3722, 2021. URL <https://doi.org/10.1109/TNNLS.2020.3015958>.
- G. J. Woeginger. Core stability in hedonic coalition formation. In *Proceedings of the 39th International Conference on Current Trends in Theory and Practice of Computer Science*, pages 33–50. Springer, 2013. URL https://doi.org/10.1007/978-3-642-35843-2_4.




Figure 2: The points for the counterexample in two-dimensional Euclidean space (rotated by 90 degrees). w was placed closer to the remaining points for an easier illustration. The black ellipses depict the clustering that is in the 1.139-core and the blue (dashed) ellipse is the 1.139-blocking coalition.

A Points for Theorem 2

The counterexample to prove Theorem 2 for $n = 9$ and $k = 3$ uses the same set of points as in Theorem 1, with each (group of) points having roughly the same “function”; see fig. 2 for an illustration. The points were obtained by optimizing a system of inequalities and are as follows.

$$\begin{aligned}
 a &: (0.0, 0.0), \\
 b &: (-0.52633831, 0.70311454), \\
 c &: (1.18678884, -1.12314325), \\
 G_1 &: (1.0, 0.0), \\
 G_2 &: (0.04545127, 0.87711815), \\
 G_3 &: (0.21399904, -1.97998602), \\
 w &: (10.0, 10.0).
 \end{aligned}$$

To find an α -core clustering with the smallest α for the set of points, we applied Algorithm 1 to each possible clustering $\mathcal{C}$.

The code for the computation of the point coordinates and verification is available at: <https://github.com/evamichelle30/CounterExampleMaxLossCore>.

Algorithm 1 CORECHECK($N, d, \mathcal{C}$)

- 1: $\alpha \leftarrow 1$; $n \leftarrow |N|$; $k \leftarrow |\mathcal{C}|$
 - 2: **for** each coalition $S \subseteq N$, $|S| \geq \frac{n}{k}$ **do**
 - 3: **for** each $i \in S$ **do**
 - 4: $\rho_i \leftarrow \text{loss}_i(\mathcal{C}) / \text{loss}_i(S)$
 - 5: **end for**
 - 6: $\alpha \leftarrow \max\{\alpha, \min_{i \in S} \rho_i\}$
 - 7: **end for**
 - 8: **return** α
-