

AI-Generated Compromises for Coalition Formation

Eyal Briman, Ehud Shapiro, and Nimrod Talmon

Abstract

The challenge of finding compromises between agent proposals is fundamental to AI sub-fields such as *argumentation* [26], *mediation* [21], and *negotiation* [20]. Building on this tradition, Elkind et al. [9] introduced a process for coalition formation that seeks majority-supported proposals preferable to the status quo, using a metric space where each agent has an ideal point. The crucial step in this iterative process involves identifying *compromise proposals* around which agent coalitions can unite. How to effectively find such compromise proposals, however, remains an open question. We address this gap by formalizing a holistic model that encompasses agent bounded rationality and uncertainty and developing AI models to generate such compromise proposals.

We focus on the domain of collaboratively writing text documents – e.g., to enable the democratic creation of a community constitution. We apply NLP (Natural Language Processing [7]) techniques and utilize LLMs (Large Language Models [35]) to create a semantic metric space for text and develop algorithms to suggest suitable compromise points. To evaluate the effectiveness of our algorithms, we simulate various coalition formation processes and demonstrate the potential of AI to facilitate large-scale democratic text editing, such as collaboratively drafting a constitution—an area where traditional tools are limited.

1 Introduction

We propose a framework for iterative *compromise-based coalition formation* that enables a set of agents to collaboratively develop a single text document. Each agent starts with an *ideal document*, modeled as a point in a (potentially high-dimensional) metric space. At each step, certain agents may collectively switch to a newly proposed *compromise document*—also represented as a point in the metric space.

Our work generalizes the model of Elkind et al. [9], who examine an iterative coalition formation process wherein agents only move to a new coalition (i) if its compromise document is closer to their ideal points than the status quo, and (ii) if the new coalition is at least as large as their current one. These two conditions ensure certain theoretical guarantees (e.g., on the convergence of the process; see footnote 5 for a formal description of the generalization). However, in many realistic scenarios, agents may not behave strictly according to these criteria. E.g., an agent might rationally move to a smaller coalition if it yields a document that more strongly aligns with its preferences, or it might stochastically deviate from strict rationality due to partial information, uncertainty, or other behavioral considerations.

Generalizing Agent Behavior. In contrast to Elkind et al. [9], we allow for more flexible coalition formation. Agents may:

- Move to a new coalition even if that coalition is smaller than their current one,
- Take actions probabilistically, representing bounded rationality or incomplete knowledge,

These relaxed conditions capture a broader range of real-world behaviors. Thus, our modeling goal is to develop a framework robust enough to accommodate both purely rational and partially rational agents, while still facilitating majority-supported text generation.

Collaborative Text Editing and the Mediator Concept. Although Elkind et al. [9] discuss the theoretical dynamics of forming coalitions via compromise proposals, they do not specify *how* such proposals are generated. We address this gap by introducing the notion of a *mediator* that systematically produces compromise documents. Specifically, we embed potential texts in a semantic metric space and employ modern natural language processing (NLP) techniques to measure distances between documents, thereby identifying compromise points that better align with multiple agents’ preferences. This approach is particularly relevant for large-scale, democratic text editing tasks—such as drafting a constitution for a decentralized autonomous organization (DAO) [13]—which existing collaborative platforms (e.g., Google Docs, Notion, Wikipedia) do not address in a structured, consensus-driven manner. By modeling documents as points in a high-dimensional embedding space, the mediator can propose new drafts that balance diverse viewpoints, thus paving the way for a more democratic editing process.

Main Contributions. Our contributions can be summarized as follows:

1. **Generalizing the Coalition-Formation Model.** We extend the work of Elkind et al. [9] by permitting less restrictive movement rules, thus supporting bounded rationality and agents who may move to smaller coalitions.
2. **AI-Mediated Proposal Generation.** We introduce algorithms that employ large language models (LLMs) and other NLP tools to embed and manipulate text documents in a semantic metric space, enabling the discovery of meaningful compromise drafts.
3. **Empirical Evaluation in Euclidean and Textual Domains.** We present simulations in both a simplified 2D Euclidean space and a more realistic text-editing environment. Our findings show that—even under relaxed decision rules—agents converge to a majority-supported document that improves upon the status quo.

For space considerations, some text is deferred to the appendix: a more detailed exposition of certain related work; application of the model to the Euclidean setting; more examples; and some details regarding the simulation results.

1.1 Model State

The model is defined by the following fixed components¹:

- A metric space X .
- A distance function $\mathbf{d} : X \times X \rightarrow \mathbb{R}_{\geq 0}$ defining the metric on X .
- A point $r \in X$, representing the fixed status quo.
- A set $V = \{v_1, \dots, v_n\}$ of n agents. Each agent $v \in V$ is associated with an ideal point $p^v \in X$ and has *Euclidean preferences* [4], meaning that preferences are determined by distance from the ideal point.

The *state* of the process is given by a coalition structure:

¹This is a centralized description for ease of presentation. Conceptually, we envision a decentralized setting, where the mediator operates as a non-centralized tool available to individual coalitions. That is, coalitions may grow in a bottom-up manner, each using an instance of the mediator independently.

- A set $D = \{d_1, \dots, d_z\}$, where each $d_i = (C_i, p_i) \in D$ for $i \in [z] := \{1, \dots, z\}$, is a *coalition*. Here, $C_i \subseteq V$ denotes the set of agents in the coalition and $p_i \in X$ the compromise point around which the coalition is formed. The coalition structure D is a partition of the agents: for all $i \neq j \in [z]$, we have $C_i \cap C_j = \emptyset$, and $\bigcup_{i \in [z]} C_i = V$.

We assume $z \in [n] := \{1, \dots, n\}$; that is, the number of coalitions does not exceed the number of agents. The notation $[z]$ refers to the index set of the current coalition structure.

1.2 Initialization, Iterative Process, and Halting Conditions

Next, we describe a specific modeling and configuration. This approach allows us to present the capabilities of the model in a clear, specific, and traceable manner, making it easier to understand. By focusing on a concrete example, we aim to illustrate the potential applications and advantages of the model, while leaving room to discuss its broader possibilities in the outlook section.

Initialization Initially, the process starts with each agent forming its own singleton coalition: namely, $D = \{d_1, \dots, d_n\} = \{(C_1, p_1), \dots, (C_n, p_n)\}$ with $C_i = \{v_i\}$ and $p_i \in X$.

Process The model contains an entity—the **mediator**—which is the workhorse of the process.

Definition 1 (mediator). *An mediator M is a function that gets as input a coalition structure D and returns as output a tuple (d_i, d_j, p) with $d_i, d_j \in D$ and p a point in the metric space.*

Intuitively, the mediator suggests that two coalitions, d_i and d_j , merge around a compromise point p . Given the current coalition structure D , the mediator returns a triple (d_i, d_j, p) , where p is proposed as a new joint position.

Each coalition responds to this suggestion according to a predefined *constitution*, which governs how agents decide whether to join the new coalition. Specifically, agents in d_i and d_j vote on whether they prefer the proposed compromise p over remaining in their current coalition. Based on these votes and the constitution, some agents may transition to the new coalition while others remain.

We first define the voting behavior of an agent before specifying the constitutions that aggregate these votes.

Definition 2 (Agent, vote). *An agent v corresponds to some ideal point p^v ; and, furthermore, a vote of agent v regarding some point p is $\text{vote}(v, p) \in \{0, 1\}$ (where $\text{vote}(v, p) = 1$ means that v accepts the suggestion to move to a coalition to be formed around p).*

Now, a constitution *const* is defined as follows.

Definition 3 (Constitution). *A constitution const gets as input a tuple (d_i, d_j, p) by the mediator and, when applied on d_i – and based on the votes of the agents in d_i , as described by $\{\text{vote}(v, p) : v \in d_i\}$ – returns an assignment to a coalition for each $v \in d_i$, namely $\text{const}(v) \in \{d_i, d^p\}$, where d^p describes the coalition to be possibly-formed around the suggested compromise point p .*

Following a suggestion of (d_i, d_j, p) and an application of the constitution *const* on d_i and d_j (which internally depends on the votes of the agents in both coalitions), the resulting Markov state contains a new coalition structure D' that is defined as follows:² $D' := D \setminus \{d_i, d_j\} \cup \{d'_i, d'_j, d^p\}$, where $d'_i := (\{v \in d_i : \text{const}(v) = d_i\}, p^i)$; $d'_j := (\{v \in d_j : \text{const}(v) = d_j\}, p^j)$; $d^p := (\{v \in d_i \cup d_j : \text{const}(v) = d^p\}, p)$.

²A coalition with no members can safely be removed from a coalition structure (such that d'_i , d'_j , and d^p may be empty).

That is, agents from d_i whom the constitution assigns to d_i remain in it; agents from d_j whom the constitution assigns to d_j remains in it; and agents from $d_i \cup d_j$ whom the constitution assigns to the new coalition around p are being moved there.

A halting condition The process halts whenever a coalition that contains an agent majority is being formed; i.e., whenever some $d \in D$, $d = (C, p)$, exists with $|C|/|V| \geq Q$, where the fraction $Q \in [0, 1]$ can be set to be majority, super majority, or consensus (in our simulation we implement a simple majority).

2 Concrete Model Realizations

We provide concrete realizations of the following ingredients: agent models (in Section 2.1), coalition constitutions (in Section 2.2), and mediators (in Section 2.3). These concrete realizations are used later, for the 2D Euclidean setting presented in the appendix and the setting that involves text documents. Furthermore, some of the details next are needed for the computer-based simulations that follow.³

2.1 One Concrete Agent Model

As abstractly stated above, an agent v corresponds to an ideal point p^v and shall have the ability to vote on a proposal p by returning a binary answer in the form of $vote(v, p) \in \{0, 1\}$ – if $vote(v, p) = 1$, then we say that v *approves* p . Naturally, various realizations of agent models are possible. Below we describe the agent model we use in our theoretical realization (later, in Section 4 we use a different, LLM-based agent model). Let us first define a simple, deterministic agent model.

Definition 4 (A deterministic agent model). *Under the deterministic agent model, an agent v within previous coalition d_i with ideal point p^v votes as follows: $vote(v, p) = 1$ if $\mathbf{d}(p^v, p) < \mathbf{d}(p^v, r)$.*

That is, an agent approves a proposal p if p is closer to its ideal point than the status quo r , and it disapproves of a proposal p otherwise. Next, as we are interested in modeling agent altruism and flexibility in the process in a naive and intuitive manner (influenced by [17]) we use a probabilistic generalization of the simple model, as described next.

In particular, given the status quo r , a proposal p , and an agent v with ideal point p^v , we define a function $F(r, p, p^v)$ that returns the probability of the agent approving p . Specifically, $F(r, p, p^v) \in [0, 1]$. (It may be helpful to note that the deterministic agent model corresponds to the probabilistic model if $F(r, p, p^v) = 1$ whenever $\mathbf{d}(p^v, p) < \mathbf{d}(p^v, r)$.)

Specifically, to model different types of non-deterministic agents, we introduce a parameter $\sigma \geq 0$, where larger values of σ results in a more altruistic agent behavior as compared to the simplest agent model described above. Mathematically, we use a half (positive) Gaussian distribution to “enlarge” a bit the region for which the agent approves the proposal (i.e., so that an agent will approve a proposal even if it is farther away from its ideal point, compared to the status quo; but with diminishing probability of doing so); formally, we have the following definition of F (note that the *else* case is 0 in case of $\sigma = 0$):

$$F(r, p, p^v) = \begin{cases} 1, & \text{if } \mathbf{d}(p^v, r) \geq \mathbf{d}(p^v, p) \\ \frac{2}{\sigma_v \sqrt{2\pi}} e^{-\frac{(\mathbf{d}(p^v, p))^2}{2\sigma_v^2}} & \text{else} \end{cases}$$

Definition 5 (A probabilistic agent model). *Under the probabilistic agent model, an agent v with ideal point p^v votes as follows: $vote(v, p) = 1$ with probability $F(r, p, p^v)$.⁴*

³We consider realizations of the model in which all agents share the same agent model; all coalitions share the same constitution; and there is one mediator throughout the process. We discuss other options in Section 5.

⁴Indeed, for $\sigma = 0$, the probabilistic agent model and the deterministic agent mode coincide.

Remark 6. *The current agent model assumes that voting behavior depends only on the distance between the proposed point p , the status quo r , and the agent’s ideal point p^v . A natural extension is to allow votes to depend on the anticipated composition of the new coalition. For instance, agents may approve p only if sufficiently many others are expected to join.*

2.2 Two Concrete Constitutions

As abstractly stated above, given a proposal for a coalition $d_i = (C_i, p_i)$ to move to a new coalition around a compromise point p , a constitution takes the votes of the agents and determines whether any of the coalition members shall move to the new coalition, and, if so, who. We explore two options for such constitutions.

- **Coalition Discipline:** A new coalition is formed only if at least $Q \in [0, |C_i|]$ members of C_i approve the proposal. Formally:

$$\text{if } |\{v \in C_i : \text{vote}(v, p) = 1\}| \geq Q, \text{ then for each } v \in C_i : \quad \text{const}(v) := \begin{cases} d_p & \text{if } \text{vote}(v, p) = 1 \\ d_i & \text{otherwise} \end{cases}$$

otherwise, $\text{const}(v) := d_i$ for all $v \in C_i$.⁵

- **No Coalition Discipline** is a special case of the above with $Q = 0$, where each agent independently decides whether to join the new coalition:

$$\text{const}(v) := \begin{cases} d_p & \text{if } \text{vote}(v, p) = 1 \\ d_i & \text{otherwise} \end{cases} \quad \text{for each } v \in C_i.$$

Remark 7. *We assume the coalition size $|C_i|$ is known to its members at the time of voting.*

Remark 8. *Agent preferences depend only on distance to their ideal point. Coalition discipline imposes coordination constraints but does not affect individual utility.*

2.3 Several Concrete Mediators

Recall that a mediator takes as input a coalition structure D and returns two coalitions, d_i and d_j , and a compromise point p . It is convenient to break the description of our realizations into the two main tasks of mediators, namely: (1) choosing the coalitions d_i and d_j to suggest p to; and (2) choosing the compromise point p to suggest to d_i and d_j . Note that the role of the AI in our design is in the implementation of such mediators.

Choosing the coalitions d_i, d_j Our mediators proceed by first computing the centroid of the coalitions’ ideal points, weighted by the coalition sizes. Formally: $\text{centroid}(D) = \arg \min_{x \in X} \frac{1}{n} \cdot \sum_{i \in |D|} |C_i| \cdot \mathbf{d}(x, p_i)$. Using the centroid, the mediators consider the distance of each coalition from the centroid, denoted by $\mathbf{d}(p_i, \text{centroid})$. The selection process is guided by a parameter $\alpha \in [-1, 1]$, intuitively ranging from whether the closest coalitions to the centroid are preferred ($\alpha = -1$), the furthest ones are preferred ($\alpha = 1$), or there is no significance ($\alpha = 0$) to the distance from the centroid.

⁵Our model builds upon and generalizes aspects of the framework proposed by Elkind et al. In the case of coalition discipline with deterministic agents and unanimous approval ($Q = |C_i|$), we recover their convergence results under the constraint that agents only transition to strictly preferred larger coalitions.

Each coalition i is assigned a score S_i based on its distance from the centroid using the following scoring function $S_i = e^{\alpha \cdot d'(p_i, \text{centroid})}$, where $d'(p_i, \text{centroid}) \in [0, 1]$ is the normalized distance; formally:

$$d'(p_i, \text{centroid}) = \frac{\mathbf{d}(p_i, \text{centroid})}{\arg \max_{j \in |D|} \mathbf{d}(p_j, \text{centroid})}.$$

Subsequently, the mediator assigns a probability $\text{prob}(d_i)$ to each coalition d_i , proportionate to its scores: $\text{prob}(d_i) = \frac{S_i}{\sum_{i=1}^{|D|} S_i}$. In practice, the mediator probabilistically chooses one coalition based on these probabilities and then selects the closest coalition to the initially chosen one.

Choosing the Compromise Point p Given coalitions $d_i = (C_i, p_i)$ and $d_j = (C_j, p_j)$, the mediator selects a compromise point $p \in X$ that minimizes the weighted sum of distances to p_i and p_j , with weights proportional to coalition sizes:

$$p = \arg \min_{x \in X} \left(\frac{|C_i|}{|C_i| + |C_j|} \cdot \mathbf{d}(p_i, x) + \frac{|C_j|}{|C_i| + |C_j|} \cdot \mathbf{d}(p_j, x) \right).$$

In Euclidean space, this reduces to the standard weighted average.

Remark 9. *The mediator is assumed to know all agents’ ideal points. This may result from voluntary disclosure or be treated as a modeling assumption. While this enables computation of globally optimal coalitions under a defined objective, the current mediator applies local, myopic merges. Designing optimal, forward-looking mediators is left for future work.*

3 Related Work

Coalition formation in a metric space has been studied from a multiagent system context [5, 36, 29]. We build upon the theoretical framework of Elkind et al. [9], which introduced a model for deliberative coalition formation in metric spaces. Their work presents a transition system to capture the dynamics of the coalition formation process. While Elkind et al. describe the formation process in detail, they assume that compromise points are provided by an external source (an oracle), without specifying how these points should be determined. Our contribution addresses this gap by introducing *mediators* that algorithmically suggest compromise points, allowing coalitions to unite around majority-supported proposals. We implement and optimize these mediators to make the coalition formation process both practical and efficient. We also demonstrate that, under a specific configuration of our model, it aligns with Elkind’s model, thereby inheriting their theoretical results for that configuration thus our model generalizes Elkind’s model. For the general case we show simulations that show convergence rates are very good even for large instances. In developing mediators, we utilize NLP techniques and LLMs; this relates to NLP-based recommendation systems [2, 32], where models suggest content based on user preferences, and to recent work in *Generative Social Choice* [12], which explores the use of LLMs to generate representative statements for social choice tasks. However, our work differs by focusing on identifying compromise points in the coalition formation process, where the goal is to find a majority-supported text or proposal. Another relevant line of work is by Bakker et al. [1], who study how machines can assist in finding agreements among individuals with diverse preferences. Their approach fine-tunes LLMs to generate statements that maximize the expected approval of a group, which is conceptually similar to our use of AI for proposing compromise points. However, our model incorporates an iterative coalition formation process, making it distinct in its operational dynamics. We also mention Yang et al. [34] that investigate how GPT-4 and LLaMA-2 behave in voting scenarios compared to human voters. They show that voting methods, presentation order, and temperature settings can significantly influence LLM choices, often reducing preference diversity and risking bias.

In the context of *Dynamic Coalition Formation*, there is significant prior work on how agents with diverse preferences form and adapt coalitions to achieve consensus [30, 19]. This is relevant, as coalition

formation plays a key role in decision-making processes, especially when agents aim to form majority-supported agreements [25]. Our approach to coalition formation in metric spaces also draws on existing research in spatial voting models [10, 15]. We are also motivated by psychological research on the ability of agents to objectively evaluate proposals. Mikhaylovskaya et al. [22] provide evidence that AI-based mediators can mitigate human biases, making AI a promising tool for generating neutral, data-driven compromise points. This motivates our use of AI-mediated coalition formation, where agents can evaluate AI-suggested compromise points to find collective agreements. Our work also draws inspiration from negotiation-based approaches to coalition formation [24, 14, 27, 9, 16, 33, 3, 11], which offer valuable insights into how agents with divergent preferences negotiate and form coalitions. These approaches further reinforce the relevance of mediation in improving the efficiency and effectiveness of coalition formation in multi-agent systems.

4 mediators in a Textual Space

As our ultimate goal relates to text aggregation – i.e., to enable an agent community to converge towards a majority-supported textual document. So, we wish to utilize the mediators framework (demonstrated also in a Euclidean space in the appendix) to a setting in which the metric space contains textual documents and coalitions form around different texts, until a majority-supported textual document is identified. We describe our specific model; and then report on computer-based simulations.

4.1 Modeling mediators in a Textual Space

Our general solution works as follows. We use word embedding (a standard NLP technique) to translate texts into numerical-valued vectors; this is crucial as, after applying such a word embedding, we are then able to compute distances between the embedded coalition ideal points and use the mediators of the Euclidean space. In this work, we use Google’s Universal Sentence Encoder [6]: this is a pre-trained model that converts sentences into fixed-size vectors, capturing their semantic meanings. (The Universal Sentence Encoder is designed to generate 512-dimensional embedding vectors, providing a semantic representation of sentences.) Thus our embedded metric space contains as elements all those 512-length vectors that can be the output of the Universal Sentence Encoder. As a distance measure in this space, we use the square root of the cosine-based dissimilarity [28], a commonly used pseudo metric in NLP.⁶

The mediator guides the coalition formation by proposing sentences to two coalitions within the given word limit (in our simulations, 15 words). In each iteration, the mediator’s goal is to find two coalitions to suggest a sentence that minimizes the *squared cosine similarity* pseudo metric. between the embedding vector of the chosen sentence and the weighted average of the embedding vectors of the two coalition points.

Our approach to coalition formation in the domain of text relies on the integration of OpenAI’s GPT-3.5-turbo-1106 model (<https://openai.com/blog/chatgpt>) with a temperature parameter (responsible for the randomness of results) of 0.75 as recommended in some of the documentation to provide a good trade-off for applications like ours where the output should be coherent but still allow for some diversity and creativity. This LLM takes 3 key roles within our framework:

⁶Formally, we define

$$d_{\cos}(A, B) := \sqrt{2 - 2 \cdot \frac{A \cdot B}{\|A\| \cdot \|B\|}} \in [0, 2],$$

which corresponds to $\sqrt{2} \cdot \sin(\theta/2)$, where θ is the angle between vectors A and B . This is a metric when restricted to the unit sphere (or more generally, to rays through the origin), but only a pseudo metric over the entire embedding space: colinear vectors have distance zero, even if they differ in magnitude. Since our mediator constructs compromise vectors via (non-normalized) weighted averages, we do not assume that vectors lie on a sphere.

1. It generates sentences that act as agent ideal points.
2. It constructs initial singleton coalitions mirroring these ideal sentences if the coalition formation process introduces noise (i.e., for simulations runs with $I = \text{true}$).
3. It proposes diverse options for aggregating two sentences, presenting methods to combine opinions from different coalitions represented as text. The process then determines the most suitable sentence by evaluating which has the embedding that is the closest to the weighted average of the two coalition embedded sentences.

Remark 10. *We assume that Euclidean distance in the embedding space reflects agent preferences. That is, texts closer to an agent’s embedded ideal point are considered more preferable. This assumption connects the embedding to the distance-based agent model, though it may not hold uniformly across domains.*

4.2 Simulations-Based Analysis

We conducted simulations to assess the robustness and resilience of the model; done as follows:

- **Parameter Tests and Scale:** The simulations were conducted with different numbers of agents, specifically $n \in \{10, 20, 30, 40, 50, 100, 1000\}$. We varied the parameters $\sigma \in \{0, 1, 1.5\}$ and $\alpha \in \{-1, 0, 1\}$, while also setting the boolean variable C —enforcing coalition discipline or not. Each parameter combination was tested across 50 repetitions, with all ideal sentences generated sharing a predetermined topic of ways to address global warming.
- **Coalition Formation Process:** We employed an iterative pursuit in an embedding space using *squared cosine similarity* pseudo metric; the prompt given to GPT was: “Give me T different sentences that are well structured about how to deal with Y with at most of 15 words” (T being the number of agents, and Y being any topic – global warming in our case); to initialize the singleton coalition with introducing noise ($I = \text{True}$), the LLM was requested to provide a sentence resembling the ideal sentence of each agent, rather than introducing additional noise through a normal distribution as conducted in the euclidean case presented in the appendix. These function as the singleton coalition sentences to be embedded into the Euclidean space. The prompt given to the GPT was: “Give me a well-structured sentence with a maximum of 15 words, resembling this sentence: Z” (where Z represents an ideal sentence of an agent).
- **Sentence Selection Process:** For each proposed sentence to the two coalitions, the LLM was tasked with generating 10 sentences that effectively combined both coalition sentences. We followed best practices for structured prompt design and multi-step reasoning [18], including these concepts:
 - *Structured Prompt Design:* Prompts should provide clear and concise instructions, ensuring that the LLM produced well-structured sentences.
 - *Encouraging Multi-Step Reasoning:* Prompts should be designed to guide the LLM through step-by-step reasoning, leveraging Zero-shot Chain of Thought (CoT) techniques to handle the task effectively.
 - *No Task-Specific Examples Needed:* Prompts should avoid the need for specific examples or task-specific training, enabling the model to generalize across different tasks.

We used the following prompts and messages given to GPT 3.5 (5 options in total):

- Mediator 1: **Prompt:** Generate 10 possible different well-structured sentences that aggregate the following two sentences. Make sure each sentence has at most 15 words. Number your answers (i.e., 1), 2), 3), 4), 5), and so on) for each sentence you propose. **Message:** You

Option	Mean Number of Iterations
Option 1	4.8000
Option 2	5.0750
Option 3	5.5250
Option 4	7.5000
Option 5	41.8125

Table 1: A comparison of different mediators-each corresponding to different prompts and LLM-usage strategy. The mean number of iterations until convergence is shown, validated with 95% confidence using ANOVA and Tukey HSD.

are a mediator trying to find agreed wording for how to deal with global warming based on existing sentences. Give a straightforward answer with no introduction to help people reach an agreed wording of a coherent sentence. (The proposed sentence was selected based on the minimal *squared cosine similarity* between its embedding and the weighted average embedding vector, considering the two embedding vectors of the coalitions and their sizes.)

- Mediator 2: **Prompt:** Generate 10 concise and clear sentences that blend the following two sentences into one coherent idea: Ensure each sentence is no longer than 15 words. Number your answers (i.e., 1), 2), 3), 4), 5), and so on) for each sentence you propose. **Message:** As a mediator, you need to find a consensus on global warming solutions. Provide straightforward and numbered suggestions to help reach a clear and agreed-upon sentence.
- Mediator 3: **Prompt:** Create 10 unique, well-structured sentences that combine these two sentences into one unified thought: Each sentence should be a maximum of 15 words. Number your answers (i.e., 1), 2), 3), 4), 5), and so on) for each sentence you propose. **Message:** You are acting as a mediator to achieve a common statement on global warming. Give direct and numbered suggestions to assist in forming a unified and coherent sentence.
- Mediator 4: This baseline mediator involved soliciting several possibilities for sentence aggregation from the GPT and then selecting the sentence that minimized the distance from the average embedding vector of the two coalitions. Instead, we simply requested GPT to provide a single sentence. The prompt and message given to GPT were the same as those given for Option 1, but instead of 10 sentences, it was asked for 1 sentence only.
- Mediator 5: This second baseline mediator denoted by **Option 5**, was to ask GPT for a completely random sentence.

We tested the number of iterations needed for coalitions to converge on a compromise, and the average distance between the compromise document and the ideal document of each agent within the coalition that halted the process.

5 Outlook and Discussion

Our findings closely align with the Euclidean case presented in the appendix:

- 1) Processes with coalition discipline and deterministic agents always exhibit some cases of non-convergence (defined as exceeding 10,000 iterations), whereas all other combinations result in convergence;
- 2) A higher number of agents leads to more iterations;
- 3) Increasing α enlarges the mean average distance between the compromise sentence of the largest coalition and the ideal sentences of its members;
- 4) Coalition discipline reduces this mean average distance.

We also analyze the performance of different mediators, summarized in Table 1. An ANOVA test confirms statistically significant differences in the number of iterations across mediation approaches. A

post-hoc Tukey HSD test (see appendix) further identifies significant pairwise differences, revealing that Option 1 achieves the fewest iterations on average.

Remark 11. *We conducted additional experiments using GPT-3 Davinci and GPT-4o Mini, both with a temperature of 0.75. As their results followed the same patterns and led to identical conclusions, we omit them here for brevity.*

5.1 Interpretation

Our simulations demonstrate the effectiveness of *mediated coalition formation*, particularly when leveraging *Large Language Models (LLMs)*. AI-mediation significantly reduces the number of iterations required for coalitions to reach a compromise while minimizing the average distance between the final compromise and individual agents’ ideal documents. Notably, LLMs combined with distance-based optimization consistently accelerate convergence compared to simpler AI based mediator approaches.

The statistical tests reinforce that meaningful differences exist among mediation strategies, underscoring the adaptability of mediation to different scenarios. The superior performance of option 2 in minimizing iterations further suggests that careful tuning of the AI based mediator’s behavior can yield substantial efficiency gains.

5.2 Future Work

Several directions for future research are outlined below:

- **Theoretical Guarantees:** Analyze convergence under relaxed rationality and probabilistic behavior; study stability and fairness properties.
- **Strategic Behavior:** Extend the model to a game-theoretic setting that accounts for strategic agents who may misreport preferences and anticipate coalition dynamics.
- **Scalability:** Develop efficient mediator selection, distributed implementations, and hierarchical coalition structures for large-scale settings.
- **Bias and Interpretability:** Address AI-induced bias, enforce fairness constraints, and improve mediator transparency.
- **Application Domains:** Apply the model to other contexts, such as participatory budgeting, resource allocation, and collaborative drafting.
- **Empirical Evaluation:** Test the framework in real-world environments (e.g., DAOs, Wikipedia); assess adoption via human studies.
- **Adaptive Mediators:** Use reinforcement learning or game-theoretic tools to adapt mediator strategies over time.
- **Decentralized Use:** Support coalition-local mediator usage in decentralized systems with autonomous agent groups.
- **Proportionality:** Mitigate majority dominance using methods like Phragmén’s rule [23] to ensure proportional influence in aggregation.
- **Forward-Looking Planning:** Extend mediators to evaluate merge sequences that optimize objectives (e.g., minimum distance or maximal support), under different assumptions about agent information.

- **Deliberation and Communication:** Extend the model to allow agent-to-agent communication, enabling persuasion or belief updates during the process.
- **Context-Sensitive Voting Behavior:** Our model compares proposals to a fixed status quo, following Elkind et al. [9]. A natural extension is to compare proposals to the current coalition point p_c , i.e., accept p if $d(p^v, p) < d(p^v, p_c)$. A more refined model would also consider coalition size and composition—agents may prefer large coalitions for influence, or avoid those with ideologically distant members.

References

- [1] Michiel Bakker, Martin Chadwick, Hannah Sheahan, Michael Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt Botvinick, et al. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems*, 35:38176–38189, 2022. doi: <https://doi.org/10.48550/arXiv.2211.15006>.
- [2] Illia Balush, Victoria Vysotska, and Solomiia Albota. Recommendation system development based on intelligent search, nlp and machine learning methods., 2021.
- [3] Martin Beer, Mark d’Inverno, Michael Luck, Nick Jennings, Chris Preist, and Michael Schroeder. Negotiation in multi-agent systems. *The Knowledge Engineering Review*, 14(3):285–289, 1999.
- [4] Anna Bogomolnaia and Jean-François Laslier. Euclidean preferences. *Journal of Mathematical Economics*, 43(2):87–98, 2007.
- [5] Laurent Bulteau, Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. Aggregation over metric spaces: Proposing and voting in elections, budgeting, and legislation. *Journal of Artificial Intelligence Research*, 70:1413–1439, 2021.
- [6] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*, N/A(N/A):N/A, 2018. doi: 10.18653/v1/D18-2029.
- [7] K. R. Chowdhary. Natural language processing for word sense disambiguation and information extraction, 2020. URL <https://arxiv.org/abs/2004.02256>.
- [8] Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, page N/A, Palo Alto, CA, 2017. AAAI Press. doi: <https://doi.org/10.1609/aaai.v31i1.10612>.
- [9] Edith Elkind, Davide Grossi, Ehud Shapiro, and Nimrod Talmon. United for change: deliberative coalition formation to change the status quo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5339–5346, Palo Alto, CA, 2021. AAAI Press. doi: <https://doi.org/10.1609/aaai.v35i6.16673>.
- [10] James M Enelow and Melvin J Hinich. *The spatial theory of voting: An introduction*. Cambridge University Press, Cambridge, UK, 1984.
- [11] Bernard Espinasse, Guy Picolet, and Eugene Chouraqui. Negotiation support systems: A multi-criteria and multi-agent approach. *European Journal of Operational Research*, 103(2):389–409, 1997.
- [12] Sara Fish, Paul Gözl, David C Parkes, Ariel D Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. Generative social choice. *arXiv preprint arXiv:2309.01291*, N/A(N/A):N/A, 2023. doi: <https://doi.org/10.48550/arXiv.2309.01291>.
- [13] Samer Hassan and Primavera De Filippi. Decentralized autonomous organization, 2021.
- [14] Teresa W Haynes, Jason T Hedetniemi, Stephen T Hedetniemi, Alice A McRae, and Raghuveer Mohan. Introduction to coalitions in graphs. *AKCE International Journal of Graphs and Combinatorics*, 17(2):653–659, 2020.

- [15] Jobst Heitzig, Forest W Simmons, and Sara M Constantino. Fair group decisions via non-deterministic proportional consensus. *Social Choice and Welfare*, N/A(N/A):1–27, 2024. doi: <https://doi.org/10.1007/s00355-024-01524-3>.
- [16] Pavel Janovsky and Scott A DeLoach. Multi-agent simulation framework for large-scale coalition formation. In *2016 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, pages 343–350, N/A, 2016. IEEE.
- [17] Anna Maria Kerkmann. Stability, fairness, and altruism in coalition formation. In Dorothea Baumeister and Jörg Rothe, editors, *Multi-Agent Systems*, pages 427–430, Cham, 2022. Springer International Publishing. ISBN 978-3-031-20614-6.
- [18] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in Neural Information Processing Systems*, 35: 22199–22213, 2022. doi: <https://doi.org/10.48550/arXiv.2205.11916>.
- [19] Hideo Konishi and Debraj Ray. Coalition formation as a dynamic process. *Journal of Economic Theory*, 110(1):1–41, 2003. ISSN 0022-0531. doi: [https://doi.org/10.1016/S0022-0531\(03\)00004-8](https://doi.org/10.1016/S0022-0531(03)00004-8). URL <https://www.sciencedirect.com/science/article/pii/S0022053103000048>.
- [20] Sarit Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94 (1–2):79–97, 1997.
- [21] Francisco P Maturana and Douglas H Norrie. Multi-agent mediator architecture for distributed manufacturing. *Journal of Intelligent Manufacturing*, 7:257–270, 1996.
- [22] Anna Mikhaylovskaya and Élise Rouméas. Building trust with digital democratic innovations. *Ethics and Information Technology*, 26(1):1, 2024.
- [23] Dominik Peters. Proportional representation for artificial intelligence. In *ECAI 2024*, pages 27–31. IOS Press, 2024.
- [24] Talal Rahwan. *Algorithms for coalition formation in multi-agent systems*. PhD thesis, University of Southampton, N/A, 2007.
- [25] Talal Rahwan, Tomasz P Michalak, Michael Wooldridge, and Nicholas R Jennings. Coalition structure generation: A survey. *Artificial Intelligence*, 229:139–174, 2015.
- [26] Ariel Rosenfeld and Sarit Kraus. Strategical argumentative agent for human persuasion. In *ECAI 2016*, pages 320–328. IOS Press, 2016. doi: 10.3233/978-1-61499-672-9-320.
- [27] Samriddhi Sarkar, Mariana Curado Malta, and Animesh Dutta. A survey on applications of coalition formation in multi-agent systems. *Concurrency and Computation: Practice and Experience*, 34(11): e6876, 2022.
- [28] Erich Schubert and Arthur Zimek. Elki: A large open-source library for data analysis-elki release 0.7.5 "heidelberg". *arXiv preprint arXiv:1902.03616*, N/A(N/A):N/A, 2019. doi: <https://doi.org/10.48550/arXiv.1902.03616>.
- [29] Ehud Shapiro and Nimrod Talmon. Foundations for grassroots democratic metaverse. *arXiv preprint arXiv:2203.04090*, N/A(N/A):N/A, 2022. doi: <https://doi.org/10.48550/arXiv.2203.04090>.
- [30] Onn M Shehory, Katia Sycara, and Somesh Jha. Multi-agent coordination through coalition formation. In *Intelligent Agents IV Agent Theories, Architectures, and Languages: 4th International Workshop, ATAL’97 Providence, Rhode Island, USA, July 24–26, 1997 Proceedings 4*, pages 143–154, N/A, 1998. Springer.
- [31] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne, 2008.
- [32] Shuohang Wang and Jing Jiang. A compare-aggregate model for matching text sequences. *arXiv preprint arXiv:1611.01747*, N/A(N/A):N/A, 2016. doi: <https://doi.org/10.48550/arXiv.1611.01747>.
- [33] Tom Wanyama and Behrouz H Far. Negotiation coalitions in group-choice multi-agent systems. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 408–410, N/A, 2006. N/A.
- [34] Joshua C Yang, Marcin Korecki, Damian Dailisan, Carina I Hausladen, and Dirk Helbing. Llm voting: Human choices and ai collective decision making. *arXiv preprint arXiv:2402.01766*, 2024. doi: <https://doi.org/10.1609/aies.v7i1.31758>.

- [35] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. doi: <https://doi.org/10.48550/arXiv.2303.18223>.
- [36] Gil Ben Zvi, Eyal Leizerovich, and Nimrod Talmon. Iterative deliberation via metric aggregation. In *International Conference on Algorithmic Decision Theory*, pages 162–176, N/A, 2021. Springer.

A Missing Text

A.1 A Concrete Example

Consider the following, toy example.

Example 12. Consider a metric space X with a set of elements P and a given distance d . We have a status quo $r \in P$ and three agents A , B , and C , each with its ideal point, p^A , p^B , p^C . Furthermore:

- each agent is non-altruistic ($\sigma = 0$);
- there is no coalition discipline;
- the mediator’s α is set to be 0;
- $\{p^A, p^B, p^C\}$ serve as the initial singleton coalition points.

The distances between the different ideal points of the agents and the status quo within the metric space are as follows (note that it is indeed a metric): $d(p^A, p^A) = d(p^B, p^B) = d(p^C, p^C) = d(p^D, p^D) = 0$; $d(p^A, p^B) = 3$; $d(p^A, p^C) = 5$; $d(p^A, r) = 9$; $d(p^B, p^C) = 2$; $d(p^C, r) = 6$; $d(p^C, r) = 8$. Consider another element of the metric space, d^{BC} , with $d(p^B, p^{BC}) = d(p^C, p^{BC}) = \frac{2}{2} = 1$.

1. **Initialization:** Each agent starts with its own coalition.

$$D = \{(C_A, p^A), (C_B, p^B), (C_C, p^C)\}$$

$$C_A = \{A\}, \quad C_B = \{B\}, \quad C_C = \{C\}.$$

2. **Iteration 1:** The mediator suggests the compromise point p^{BC} to the coalitions (C_B, p^B) and (C_C, p^C) . Both agents approve since $1 < 6$ and $1 < 8$. We arrive to the following coalition structure D' :

$$D' = \{(C_A, p^A), (C_{BC}, p^{BC})\},$$

$$C_{BC} = \{B, C\}.$$

3. **Halting condition:** A coalition with an agent majority has been formed (as $|C_{BC}|/|D| > 0.5$), thus the process halts.

A.2 Mediators in a 2D Euclidean Space

In this section, we consider a rather simple setting where the metric space X contains points in a 2D Euclidean space and the distance is ℓ_2 . This serves to illustrate the fundamental properties of our model and showcases the operation of our algorithms. As a usecase, consider a scenario in which an agent community collaborates to mutually select a location for a social event (e.g., a picnic).

A.2.1 Simulation-Based Analysis

We describe the design of the computer-based simulations we have conducted; and report and discuss the results.

We have generated instances of our model for the realization described above for a 2-dimensional Euclidean space. Next are details of the specific configuration used:

- **Status Quo:** Generated uniformly at random between $(0, 200) \times (0, 200)$.
- **Ideal Points:** Drawing inspiration from the literature [8], each agent was assigned an ideal point (x, y) with both coordinates sampled from either the uniform distribution between $(0, 200)$ or from a 2-dimensional Gaussian Mixture Model (GMM). The GMM represents the overall probability distribution as a weighted sum of several Gaussian components with multiple peaks. In our simulations, we considered GMMs with g combined Gaussian distributions, for $g \in \{0, 1, 2, 3, 4\}$ with the mean of each Gaussian being distributed uniformly between 0 and 200 in each dimension, its deviation distributed uniformly between 0 and 50, and the weights signifying the importance of each Gaussian are distributed from the $Dirichlet(\alpha^g \in \mathbb{R}_0^{+g})$ distribution with α set to 1 (resulting in g numbers that sum to 1). This sampling of ideal points is demonstrated in the supplementary material. Note that we treat GMM with $g = 0$ (i.e., 0 peaks) as the uniform distribution.
- The different number n of agents used in the simulations was $n \in \{10, 20, 30, 40, 50, 100, 250, 1000\}$.
- **Coalition Discipline:** we evaluated and compared instances with coalition discipline and without (as described in the realization Subsection in the main text).
- For the mediator, we have used $\alpha = \{-1, 0, 1\}$ (as described in realization Subsection in the main text).
- We have used $\sigma_v \in \{0, 10, 20, 30\}$ as the degree of altruism, representing the smoothing of agents' approval functions (as described in the realization Subsection in the main text).
- For the initialization of the singleton coalitions we set a parameter $I \in \{True, False\}$: for $I = False$ the initial singleton coalition points were set to be the ideal points of each agent; while for $I = True$, the initial singleton coalition points were generated using a 2-dimensional Gaussian distribution with a mean being the ideal point p^v and with a covariance matrix with $\sigma_x \sim U(0, 10)$ and $\sigma_y \sim U(0, 10)$ on the diagonal (and zeros off-diagonal).

We conducted 100 independent repetitions for each configuration. Next we present our two evaluation metrics (the first measures the process speed, while the second measures the process quality):

- **Speed of convergence:** average number of iterations until the halting condition is met.
- **Quality of converged state:** average distance between the proposal of the coalition containing an agent majority to the ideal points of the agents within that coalition; formally, for the single coalition $d = (C, p)$ in the halting state, with $\frac{|C|}{n} \geq 0.5$ we compute $\frac{1}{|C|} \sum_{v=1}^{|C|} d(p, p^v)$.

For efficiency, we halt our simulations whenever the number of iterations exceeds a threshold of 10,000 (i.e., we treat an instance for which no convergence is reached within 10,000 iterations as an instance that does not converge at all).




Figure 1: Coalition formation result- "Dealing with Global Warming", for $n = 10$, $C = \text{False}$, $\alpha = 0$, $\sigma = 0$, $I = \text{True}$.

A.2.2 Results and Discussion

Next we discuss the main conclusions, drawn at a 5% significance level:

1. Processes with coalition discipline and non-altruist agents always result in some non-convergences (i.e., the number of iterations is greater than 10,000) while all other combinations result in convergence.
2. More agents result in more iterations (linearly), shorter mean distances, and higher log-odds of a converging process before 10,000 iterations.
3. Higher α leads to a larger mean average distance.
4. Coalition discipline shortens the mean average distance.
5. High interaction between n and α , n and coalition discipline, and σ and coalition discipline results in more iterations until the halting condition.
6. High interaction between n and number of peaks (GMM), and coalition discipline and number of peaks (GMM), leads to fewer iterations until the halting condition.

B Illustrating the Simulation Framework

Example 13. To better illustrate the process, we present one of the simulations conducted with fixed parameters outlined in Figure 1. The simulation involves 10 ideal sentences of agents regarding dealing with global warming (of maximum 15 words) projected onto a 2D Euclidean space, showcasing the coalitions each agent belongs to by the time the process concludes (i.e., the halting condition is satisfied). The visualization method employed for multi-dimensional data is adapted from [31].

C Illustrations

We include several figures – tables and illustrations:

- Figure 2 contains an illustration of the abstract model.
- Figure 3 and Figure 4 contains graphs regarding the location based simulations. In particular, Figure 3 illustrates the average number of iterations required for different parameter combinations to reach a majority coalition. This is depicted as a function of the total number of agents, assuming convergence within 10,000 iterations. To gain a deeper understanding of how different parameter combinations influence the convergence, we conducted a detailed analysis as well. In Figure 4, we present the mean of the average distances between each ideal point of an agent and the coalition points, specifically those that have reached a majority.

Note that,

- The linear regression model for predicting the average number of iterations resulted in an adjusted R^2 of 0.683. Coefficients provide insights into the linear connection. However, for the linear regression model on mean distances, the adjusted R^2 is 0.177, reflecting a noisier and less explainable relationship, as observed in Figure 5, 7, and 8 (in the supplementary material).
- The logistic regression model demonstrates $F1$ score 0.89, reflecting a strong ability of prediction.
- Figure 5 contains a graph regarding the text based simulation.
- We illustrate in Table 2 the example given in the main text.
- We illustrate another example to show the capabilities of our model with fixed parameters outlined in Figure 7 and Table 2. This example also involves 10 ideal sentences of agents but this time deals with ways to fight for rights of minorities (of maximum 15 words).
- Figures 6, 8, 9, and 10 shows raw statistical analysis.
- Figure 11 shows the different agent distributions used in the experiments.




Figure 2: Illustration of the Abstract Model.




Figure 3: Location Use case- Speed of convergence (i.e., average number of iterations) as a function of the number n of agents; each of the line corresponds to different combination of the parameters α , g , σ , coalition discipline, I .




Figure 4: Location Use case- A noisy outcome: number of agents VS average distance chosen combinations of α , g , σ , coalition discipline, I .




Figure 5: Text use case: Number of agents as a function of the average number of iterations and chosen combinations of α , σ , C , I .

R-squared:	0.685		
Model:	OLS	Adj. R-squared:	0.683
Method:	Least Squares	F-statistic:	345.0
Prob (F-statistic):	6.76e-235		
Time:	13:26:49	Log-Likelihood:	-3965.6
No. Observations:	960	AIC:	7945.
Df Residuals:	953	BIC:	7979.
Df Model:	6		

	coef	std err	t	P> t	[0.025	0.975]
intercept	-6.6379	1.364	-4.866	0.000	-9.315	-3.961
n	1.5814	0.066	24.030	0.000	1.452	1.711
n·α	0.1633	0.022	7.485	0.000	0.120	0.206
n·peaks	-0.0339	0.016	-2.146	0.032	-0.065	-0.003
n·C	0.5116	0.057	9.051	0.000	0.401	0.622
peaks·C	-1.6322	0.572	-2.856	0.004	-2.754	-0.510
sigma·C	0.4137	0.058	7.085	0.000	0.299	0.528

Figure 6: Linear Regression for Average Iterations Results- Location use case.

Iter. 1	Iter. 2	Iter. 3	Iter. 4	Iter. 5	Iter. 6	Iter. 7
0: "Educate yourself on the issues facing minority communities"		1,0: Amplify the voices of marginalized groups				
1: "Use your privilege to amplify marginalized voices"		1,0: Amplify the voices of marginalized groups				
2: Advocate for policies that promote equality						
3: "Join or support organizations advancing minority rights"					3,7,8,4: "Engage in advocacy and support minority rights"	6,3,7,8,4: "Participate in protests to show solidarity"
4: "Engage friends and family in conversations"	8,4: "Engage friends and family while standing in solidarity"		7,8,4: "Support businesses for marginalized empowerment"		3,7,8,4: "Support advocacy and minority rights"	6,3,7,8,4: "Participate in protests to show solidarity"
5: "Challenge discriminatory behavior"				5,9: "Confront discrimination and learn intersectional issues"		
6: "Vote for political candidates who prioritize minority rights"						6,3,7,8,4: "Participate in protests to show solidarity"
7: "Support minority-owned businesses"			7,8,4: "Support businesses for marginalized empowerment"		3,7,8,4: "Support advocacy and minority rights"	6,3,7,8,4: "Participate in protests to show solidarity"
8: "Stand in solidarity during protests"	8,4: "Engage friends and stand in solidarity"		7,8,4: "Support businesses for marginalized empowerment"		3,7,8,4: "Support advocacy and minority rights"	6,3,7,8,4: "Participate in protests to show solidarity"
9: "Learn about intersectional issues"				5,9: "Confront discrimination and learn intersectional issues"		

Table 2: Matrix representation of a coalition formation process fighting for the rights of minorities. Each column represents an iteration, and each row corresponds to an agent. Initially, each agent is in its singleton coalition. The table shows the compromise sentences accepted by pairs of coalitions as the process evolves.




Figure 7: Coalition formation result- "Fighting for Rights of Minorities", for $n = 20$, $C = \text{False}$, $\alpha = 0$, $\sigma = 0$, $I = \text{True}$.

Model:	Logit	Df Residuals:	762			
F1 Score:	0.89	Pseudo R-squ.:	0.5432			
Log-Likelihood:	-212.55					
converged:	True	LL-Null:	-465.36			
LLR p-value:	4.928e-107					
	coef	std err	z	P> z	[0.025	0.975]
n	-0.0425	0.009	-4.785	0.000	-0.060	-0.025
α	0.3311	0.153	2.163	0.031	0.031	0.631
peaks	-0.2525	0.080	-3.140	0.002	-0.410	-0.095
sigma	-0.1620	0.015	-11.152	0.000	-0.190	-0.134
C	5.3229	0.404	13.170	0.000	4.531	6.115
I	-1.2725	0.253	-5.037	0.000	-1.768	-0.777

Figure 8: Logistic Regression for the probability of encountering non-convergent results-Location use case.

R-squared:	0.180		
Model:	OLS	Adj. R-squared:	0.177
Method:	Least Squares	F-statistic:	69.75
Prob (F-statistic):	8.53e-41		
Log-Likelihood:	-3986.9		
No. Observations:	960	AIC:	7982.
Df Residuals:	956	BIC:	8001.
Df Model:	3		

	coef	std err	t	P> t	[0.025	0.975]
const	53.4245	1.317	40.552	0.000	50.839	56.010
n	0.4051	0.045	9.096	0.000	0.318	0.493
α	2.8802	0.610	4.723	0.000	1.683	4.077
C	-10.1666	0.996	-10.209	0.000	-12.121	-8.212

Figure 9: Linear Regression for Average Distance.

R-squared:	0.999		
Model:	OLS	Adj. R-squared:	0.999
Method:	Least Squares	F-statistic:	2.404e+04
Prob (F-statistic):	4.35e-89		
Time:	11:41:12	Log-Likelihood:	-109.19
No. Observations:	64	AIC:	224.4
Df Residuals:	61	BIC:	230.9
Df Model:	2		

	coef	std err	t	P> t	[0.025	0.975]
const	-4.0914	0.288	-14.188	0.000	-4.668	-3.515
n	1.0078	0.005	218.869	0.000	0.999	1.017
alpha·C	0.9930	0.308	3.224	0.002	0.377	1.609

Figure 10: Linear Regression for Average Iterations Results- Text use case.




Figure 11: 1000 agents sampled from a GMM with a mixture of 0,1,2,3,4 Gaussians over a Euclidean 2D grid over (0,200) square, using weights distributed by Dirichlet distribution. The color of the contour lines represents different density levels in the mixture model; for 0 peaks, the density is constant across the grid.

Eyal Briman
Ben-Gurion University of the Negev
Beer Sheva, Israel
Email: briman@post.bgu.ac.il

Ehud Shapiro
Weizmann Institute of Science
Rehovot, Israel
Email: udi.shapiro@gmail.com

Nimrod Talmon
Ben-Gurion University of the Negev
Beer Sheva, Israel
Email: nimrodtalmon77@gmail.com