Information distortion through aggregation

Jordan Ackerman, Sabina Sloman, and Pamela Bilo Thomas

June 20, 2019

1 Introduction

As news consumers, as scientists, and in daily life, we are surrounded by a diversity of opinions and perspectives, and faced with the task of reconciling often contradictory accounts. Is a news story true or fake? Which of several competing scientific hypotheses should we subscribe to?

In most cases, we do not have direct experience with the event or phenomenon we are attempting to reach a conclusion about. We rely on some condensed version of the relevant information—a *perceived consensus*. For example, we may read an article that describes a news event. We may rely on a meta-analysis or literature review to summarize the evidence in favor of one scientific theory or another. In a system of multiple agents, each trying to reach their own conclusion about the state of the world, this introduces a positive feedback loop: Our judgments influence the perceived consensus, which in turn influences our judgments.

Both of the examples above are aggregations of the conclusions arrived at by multiple individuals. However, the aggregation procedure introduces a degree of distortion and information loss: Some information is communicated at the expense of other information. For example, a news article may report on a conspiracy theory, but neglect to mention that it's an extremely fringe belief. Moreover, the nature of this distortion is typically unknown to us.

Our formalization of the perceived consensus incorporates modular group structure and hierarchical aggregation (see section 2 below). This is both a way to introduce systematic aggregation, and is reflective of the generative process of many aggregated forms of information. For example, a journalist may take care to gather the perspectives of many diverse groups, rather than reporting the results of an opinion poll. Similarly, a scientific literature review may delve into competing hypotheses, without reporting how many people in the field subscribe to either one, or the scope of nuance of beliefs.

With the ability to learn about events through various channels like social media, cable news, or the newspaper, it is possible for individuals to hold widely differing opinions about an event that has occurred. We model this scenario using aggregation techniques: where we identify perceived consensus about an event, and show the impact of that consensus on individuals in a group. Additionally, we show that this perceived consensus is an aggregation of the opinions of various groups of people who hold different ideas.

2 Model

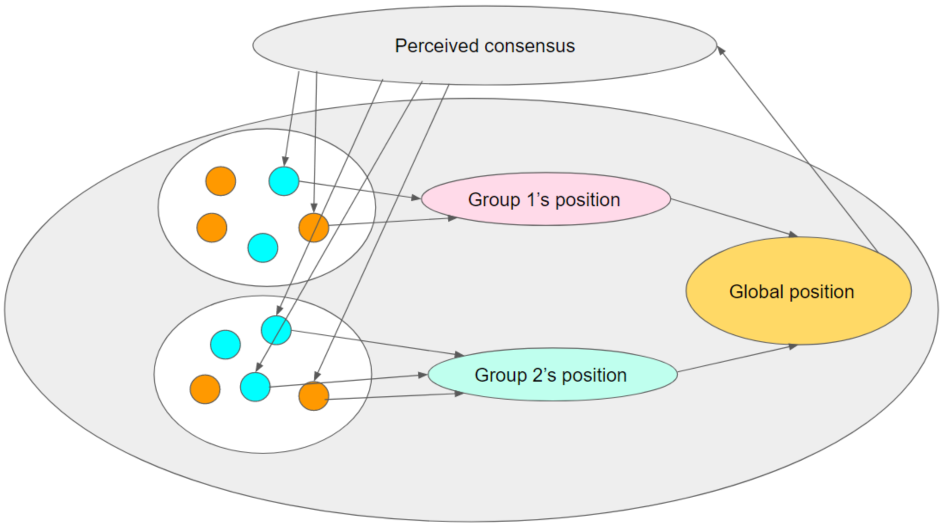

Figure 1: Our model depends on the aggregation of the opinions of groups. We see that for each 'round', we get the group position, which is aggregated into the global position. From there, we get the perceived consensus, which is used to inform agents in the next round.

We define the perceived consensus as the summarized recent history of global positions. The state of the world is given a value of 0, and in the case where we have two competing hypotheses, they are given the values of -1 and 1. Individuals will be assigned beliefs between -1 and 1, where values close to 0 represent more nuanced beliefs, and values further away represent a commitment to the respective hypothesis. We propose a model that has m groups that consist of n agents, which will influence the perceived consensus. In this way, the broadcasting patterns of the groups influence the views of the next generations via a feedback loop.

2.1 How do we arrive at a global position?

In the most simplistic scenario, the global position is equal to the sum of the position of the groups divided by the number of groups. In this way, each group has equal impact on the global position. We are aware that in a real-life scenario, some groups may have a larger effect on the global position than others. However, we use this simple model and assume that each group has equal representation at the global opinion. Further work can explore the effects of groups which have larger influence over the global position.

2.2 How do we arrive at a group position?

For every time step, group membership gets updated. We assign a 5 percent probability of an individual leaving the group and being replaced by another individual. Additionally, we identify which agents will be broadcasters of information, which will influence the position that the group will take. Each individual has a $\frac{1}{\sqrt{n}}$ chance of being a broadcaster. These broadcasters will use their own individual weights which will incorporate the impact of the global position, and their own opinion, on their final opinion. After these positions have been updated, the group position is the average of the opinions of the broadcasters.

2.3 How does each individual arrive at their position?

Each individual is assigned a random opinion between -1 and 1. These opinions are averaged with the perceived consensus to get to the informed opinion based on individual weights. In this way, we model the effect of consensus information being presented to the individual.

2.4 How is the perceived consensus computed?

The perceived consensus is computed by averaging the global position over a certain number of chosen timesteps, divided by that number of timesteps. In that way, the perceived consensus will be influenced by past events. The perceived consensus will be similar to a Bayesian model where the function is continuously updated based on new information, but information will also be forgotten as the model will only look back a certain number of timestamps PC. Depending on how we set PC, the perceived consensus will be more or less volatile. A longer view of history will present a model which is less resistant to change.

2.5 Components

Netlogo was utilized to build and run grid search (Behavior Space) over the model. Python was used for data analysis. Details of the model are presented below.

Parts

1. World: a discrete square grid, boundaries (periodic or not) and world size are unimportant.

2. Groups: Groups are invisible stationary agents and are placed relative to world size.

3. Agents: Agents are also stationary, and are assigned membership to a group.

Initialization

1. World is created with size (L x L)

2. M group agents are created (rows * columns) and placed in square grid relative to size of the world. Their main function is to position groups in space for easy viewing.

3. N agents are added in a circle around each invisible group agent. Agents are assigned a random float value (my-position) from [-1,1], from a uniform distribution.

4. C (Consensus) is initially initialized as an empty list - it will become a vector of global positions observed at each time step.

5. PC (Percieved Concensus) is the window size of most recent global position observations in C that agents are allowed to see. This value is used to re-initialize (bootstrap) C with PC random float value observations drawn from [-1,1], from a uniform distribution.

6. SW (Self Weight) is a value from [0,1] that represents the weight agents give to their own position when broadcasting their position. (1 - SW) is the weight given to the global position.

7. PL (Probability to live) is the probability that each agent will continue to live on each tick.

8. PB (Probability to Broadcast) is the probability that each agent will broadcast their position on the current round. This is set to 1/sqrt(N)

9. mandatory-broadcasting: If ON, all agents broadcast each round. If OFF, agents broadcast to their group with PB

Dynamics (Each timestep)

1. Draw broadcasters: Decide which agents will broadcast this round using PB
2. Each agents calculates their position to broadcast by combining their position with the last PC values of C as follows: (my-position * SW) + (C[-PC:] * (1-SW)) ——— [-PC:] is python notation...
3. Each group calculates their position by taking the average of all broadcast agent positions in their group.
4. Global position is calculated by taking the average of all group positions. This value is appended to C
5. Average private opinion is calculated by taking the average of all broadcasting agents in the system.

Outcomes

1. Distance between the average private position of all agents and the aggregated global position (consensus) at each timestep.

Parameters to vary for batch runs

1. SW
2. PC
3. M
4. N
5. mandatory-voting

3 Results

We are interested how information becomes *distorted* as a function of the various layers of aggregation. If judgments are independently sampled from some true distribution, the average of all the judgments will converge on the expected value of that distribution. Our measure of *information distortion* is *the degree to which the global position deviates from the average of all judgments* (we term this last quantity the *average private opinion*).

Our model introduces distortion at many layers: Global positions are averages taken over not individual observations, but group positions; group positions are calculated from unbalanced subsets of their members; and each individual observation is skewed away from its IID component and towards the perceived consensus. In this section, we unpack the effect of each of these layers of aggregation on information distortion.

3.1 Number of observations

Figures 2a, 2b and 4a show how information distortion varies as a function of the number of groups, and the number of agents in each group. In general, increasing the number of members of a group has no effect, while increasing group size reduces information distortion. Since the global position is an average of group positions, we would expect that when the group positions are good approximations of the truth, by the Law of Large Numbers, more observations will lead to a global position that more accurately reflects the state of the world. Note, however, that the reduction in distortion is small when $PCL = 1$. When $PCL > 1$, the perceived consensus has an opportunity to "stabilize" (converge to the true expected value of 0) before being incorporated into agents' judgments. However, when $PCL = 1$, initial asymmetries in the system are not cancelled out, and instead propagate through the system. In these cases, group positions are not such good approximations of the truth.

On the other hand, adding more agents to a group has a negligible, if any, effect on the degree of distortion. In the case they are not systematically driving the group position towards the perceived consensus, additional agents lead to more accurate *group* positions. However, the *accuracy* of the group decisions does not matter; it matters only that the group positions are representative samples of the distribution being estimated. While each additional agent may help their *group* be more accurate, they have no effect on the global position.

The above plot the distance between global position and the average of all private positions as a function of self-weight and consensus length. When self-weight is 0 (all information is coming from the global position) and agents are only allowed to see the single most recent global position (C-len 1), distortion is significantly higher than is any other case, hovering around 0.5.

As self-weight and C-len (perceived consensus) increase the overall distortion between all private positions dnd the global position decreases. Variation in this measure tends to fall away when self-weight is above 0.75, with C-len of 1 holding distortion slightly above other cases, especially when mandatory broadcasting is off. In addition when mandatory broadcasting is ON, the distortion levels seems to mildly increase, particularly where C-len is 2, 5, or 10.

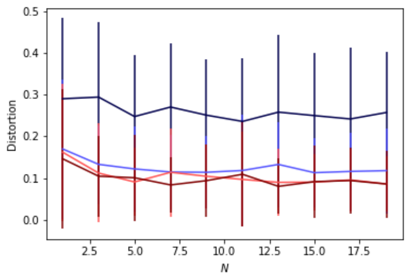

(a) Distortion as a function of the number of agents per group ($SW = .5$).

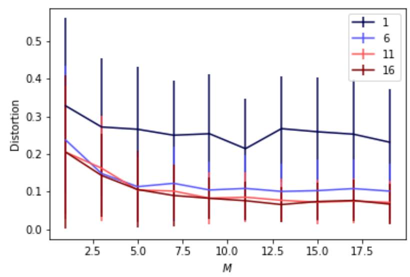

(b) Distortion as a function of the number of groups ($SW = .5$).

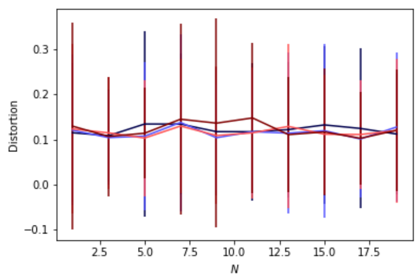

(c) Distortion as a function of the number of agents per group ($SW = 1$).

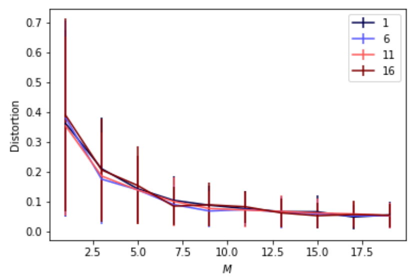

(d) Distortion as a function of the number of groups ($SW = 1$.

Figure 2: Distortion as a function of the number of members in each group (N), the number of groups (M) and the self-weight parameter (SW). “Distortion” is the absolute value of the difference between the global position and the average private opinion on the 100^{th} time step. The color of the lines indicates the PC value. Broadcasting is not mandatory. Lines and error bars indicate the means and standard deviations, respectively, of 10 simulations per parameter setting.

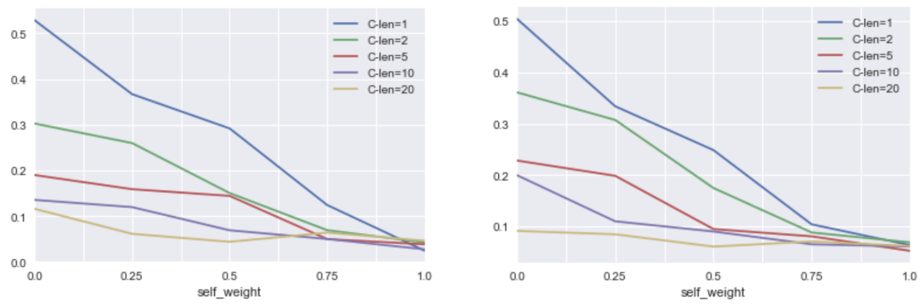

(a) Distortion as a function of the number of self-weight and perceived consensus length, mandatory broadcasting ON.

(b) Distortion as a function of the number of self-weight and perceived consensus length, mandatory broadcasting ON.

Figure 3: Here we demonstrate the impact of changing the ratio of information agents use from their private position versus the global position when calculating the position they will broadcast to their group.

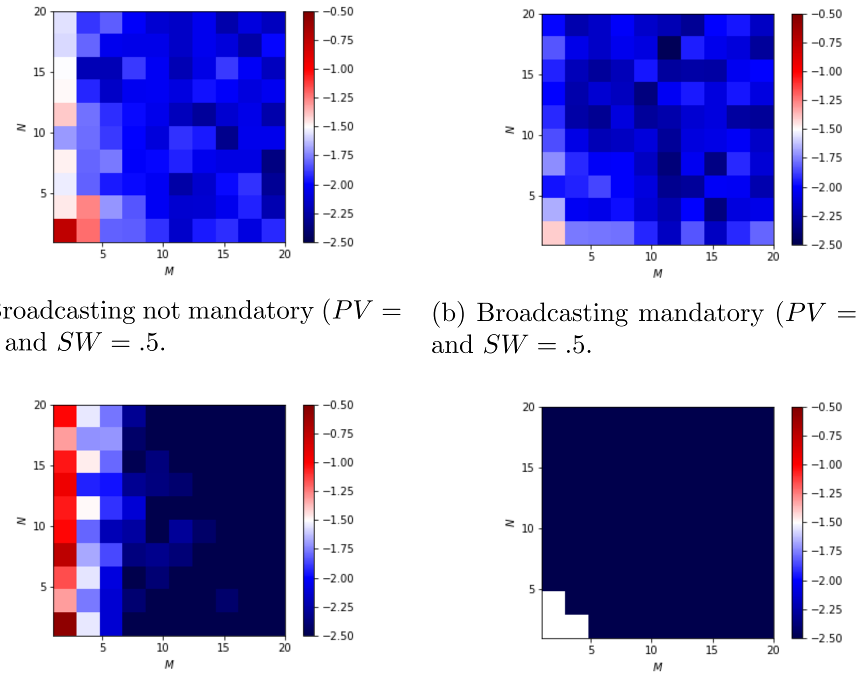

(a) Broadcasting not mandatory ($PV = \frac{1}{\sqrt{N}}$) and $SW = .5$.

(b) Broadcasting mandatory ($PV = 1$) and $SW = .5$.

(c) Broadcasting not mandatory ($PV = \frac{1}{\sqrt{N}}$) and $SW = 1$.

(d) Broadcasting mandatory ($PV = 1$) and $SW = 1$.

Figure 4: Log(distortion) as a function of the number of members in each group (N) and the number of groups (M). "Distortion" is the absolute value of the difference between the global position and the average private opinion on the 100^{th} time step.

3.2 Representativeness of group positions

The second opportunity for distortion is that within each group, a relatively small, and usually unequal, number of agents broadcast their positions. In the case of mandatory broadcasting, each of the group positions is calculated from a larger and equal number of observations.

Figure 4b shows how the level of distortion changes as a function of N and M when broadcasting is mandatory ($PV = 1$). As expected, information distortion is reduced with increasing N and M, and reduced overall.

3.3 Attention to the perceived consensus

Figures 2c and 2d show the effect of N and M, respectively, when $SW = 1$—agents do not incorporate the perceived consensus into their judgments. As expected, PCL has no effect since it is not factored into the global position. For the reasons explained above, increasing N has no effect, while increasing M tends to reduce distortion.

Figure 4c shows how distortion varies as a function of both N and M when $SW = 1$. Unlike in Figure 4a, where $SW = .5$, i.e. the perceived consensus is fed through agents' judgments, increasing the number of agents within a group actually reduces distortion. This may be by creating more accurate group positions. While it only matters that the group positions are representative of the generative distribution, when M is finite, increasing N in the absence of the feedback loop may ensure this is the case.

Figure 4d shows distortion when broadcasting is mandatory and $SW = 1$. As expected, the distortion is extraordinarily low. The white sections correspond to a distortion of 0 ($log(0) = -inf$). In the case where $M = 1$, $SW = 1$ and broadcasting is mandatory, the global position reduces to the average private information (we speculate that in this case with larger N, the non-zero distortion levels are simply due to floating point rounding and are actually just very small numbers).

4 Conclusion and Future Work

When agents only see a short history, the aggregate position is volatile, and does not represent the state of the world. When agents get to see a longer history, the aggregate position becomes more stable. Therefore, we can infer that when agents can see more information, they are not as eager to update or change their positions. We can infer that when this model is applied to model behavior, individuals who have seen more information are less likely to change their opinions on a topic as quickly. When an agent has a longer history of seeing a topic, their opinion will be more stable.

Further ideas worth exploring include assigning higher weights to individuals who have extreme positions, and binarizing positions. Additionally, in our current model, individuals have opinions drawn from a random distribution. We are interested in investigating a scenario where individuals will instead 'inherit' beliefs from the agents who leave groups, so that we can create a generational model where new agents learn from previous ones.

In this paper, we have presented a model which incorporates group behavior, averaged opinions, and aggregation. We also modelled the effect that group consensus had on influencing individual behavior. In this model, we show that longer group history results in a more stable opinions.